642-885 Dumps

Deploying Cisco Service Provider Advanced Routing (SPADVOUTE)

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

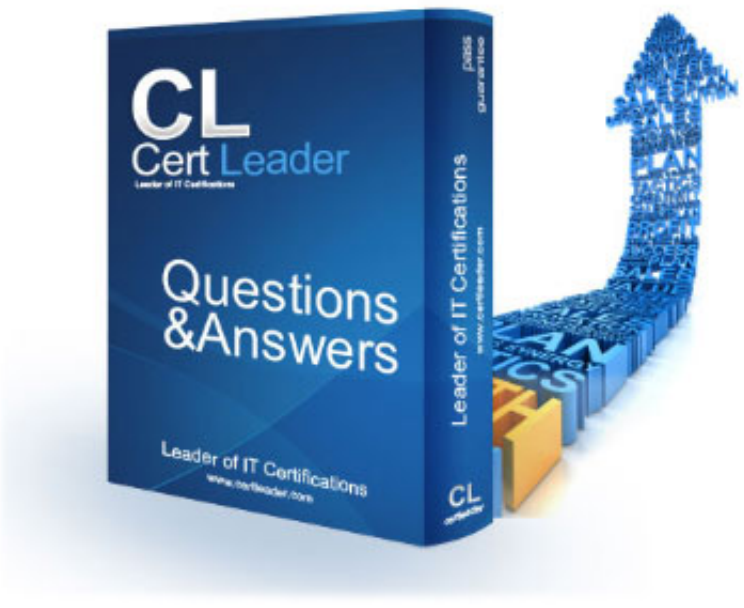

NEW QUESTION 1
Which command configures a Source Specific Multicast on a Cisco IOS XR router?

A. configuremulticast-routing address-family ipv4 interface all enableexitrouter igmp version 3 commit
B. configuremulticast-routing address-family ipv4 interface all enableexitrouter igmp version 2 commit
C. configuremulticast-routing address-family ipv4 interface all enableexitrouter igmp version 1commit
D. configure interface all enable exitrouter igmp version 3 commit

Answer: A

NEW QUESTION 2
When implementing interdomain multicast routing, which mechanism can be used to advertise multicast sources in one domain to the other domains, allowing the RPs to build interdomain multicast distribution trees?

A. Multiprotocol BGP
B. PIM
C. MSDP
D. Auto RP
E. BSR
F. MLD

Answer: C

Explanation: Multicast Source Discovery Protocol
Multicast Source Discovery Protocol (MSDP) is a mechanism to connect multiple PIM sparse-mode domains.
MSDP allows multicast sources for a group to be known to all rendezvous point(s) (RPs) in different domains.
Each PIM-SM domain uses its own RPs and need not depend on RPs in other domains. An RP in a PIM-SM domain has MSDP peering relationships with MSDP-enabled routers in other domains.
Each peering relationship occurs over a TCP connection, which is maintained by the underlying routing system.
MSDP speakers exchange messages called Source Active (SA) messages. When an RP learns about a local active source, typically through a PIM register message, the MSDP process encapsulates the register in an SA message and forwards the information to its peers. The message contains the source and group information for the multicast flow, as well as any encapsulated data. If a neighboring RP has local joiners for the multicast group, the RP installs the S, G route, forwards the encapsulated data contained in the SA
message, and sends PIM joins back towards the source. This process describes how a multicast path can be built between domains.

NEW QUESTION 3
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 4
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 5
A junior network engineer has just configured a new IBGP peering between two Cisco ASR9K PE routers in the network using the loopback interface of the router, but the IBGP neighborship is not able to be established. Which two verification steps will be helpful in troubleshooting this problem? (Choose two.)

A. Verify that the network command under router BGP is configured correct on each router for announcing the router's loopback interface in BGP
B. Verify that the ibgp-multihop command under the BGP neighbor is configured correctly on each router
C. Verify that the loopback interfaces are reachable over the IGP

D. Verify that the update-source loopback command under the BGP neighbor is configured correctly on each router
E. Verify that the ttl-security command under the BGP neighbor is configured correctly on each router to enable the router to send the BGP packets using a proper TTL value
F. Verify that the UDP port 179 traffic is not being blocked by an ACL or firewall between the two IBGP peers

Answer: CD

NEW QUESTION 6
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 7
A CRS router that runs Cisco IOS XR has dual routing processors installed. Which solution should be implemented to prevent OSPF adjacency flapping if the primary routing processor fails?

A. NSR
B. OSPF Fast Timers
C. OSPF RE Sync
D. router msdp
E. NSF

Answer: A

NEW QUESTION 8
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

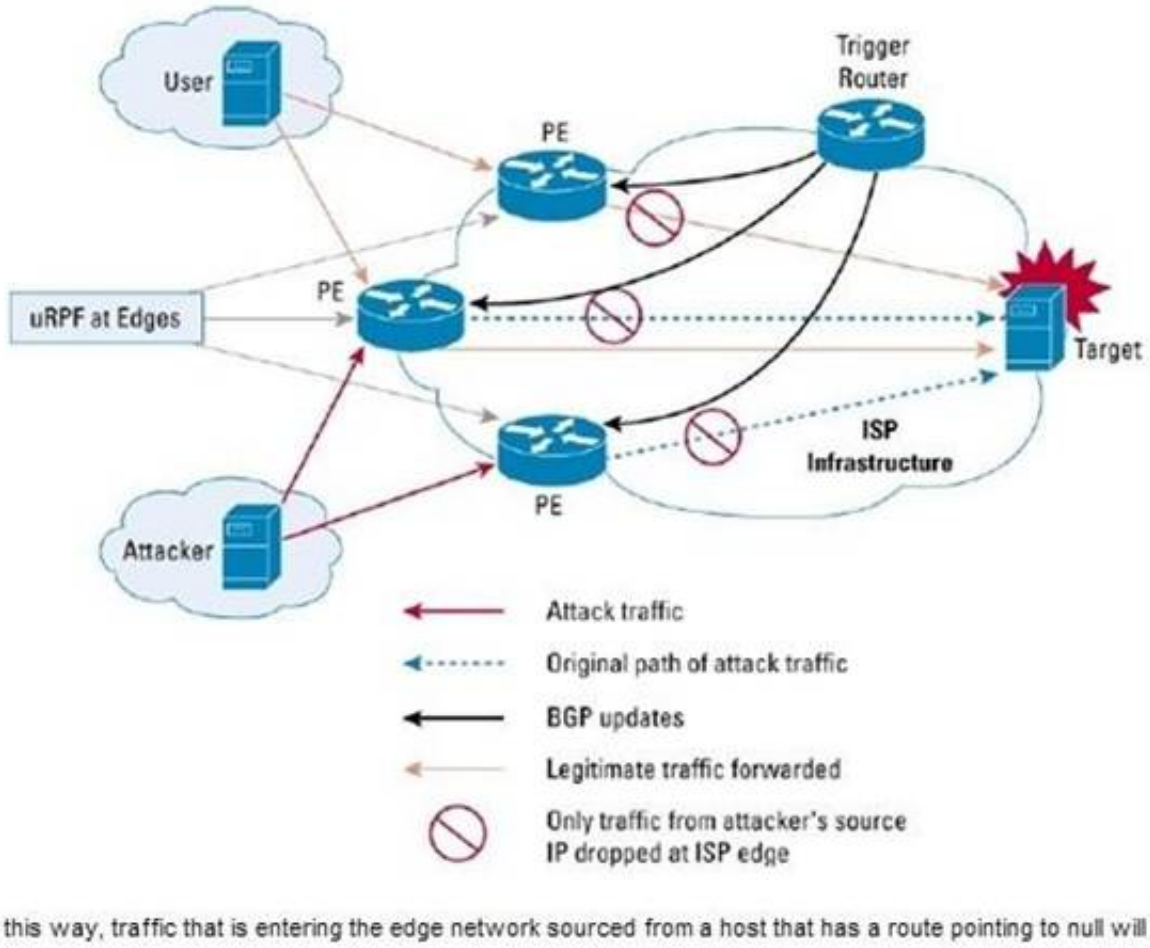

In this way, traffic that is entering the edge network sourced from a host that has a route pointing to null will result in a uRPF drop.

NEW QUESTION 9
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 10
Which type of BGP session behaves like an EBGP session during session establishment but behaves like an IBGP session when propagating routing updates where the local preference, multi-exit discriminator, and next-hop attributes are not changed?

A. BGP sessions between a route reflector and its clients
B. BGP sessions between a route reflector and its non-client IBGP peers
C. BGP sessions between a route reflector and another route reflector
D. Intra-confederation IBGP sessions
E. Intra-confederation EBGP sessions

Answer: E

Explanation: http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.7/routing/configuration/guide/rc37bgp.html#wp1191371
BGP Routing Domain Confederation
One way to reduce the iBGP mesh is to divide an autonomous system into multiple subautonomous systems and group them into a single confederation. To the outside world, the confederation looks like a single autonomous system. Each autonomous system is fully meshed within itself and has a few connections to other autonomous systems in the same confederation. Although the peers in different autonomous systems have eBGP sessions, they exchange routing information as if they were iBGP peers. Specifically, the next hop, MED, and local preference information is preserved. This feature allows you to retain a single IGP for all of the autonomous systems.

NEW QUESTION 10
Which two functions are supported for BGP extension MP-BGP for IP multicasting? (Choose two.)

A. A network can support incongruent unicast and multicast topologies.
B. A network can support congruent unicast and multicast topologies.
C. MP-BGP is an enhanced BGP that carries routing information for multiple network layer protocols and IP multicast routes.
D. MP-BGP carries single sets of routes for unicast routing and multicast routing.
E. MP-BGP is useful when a link dedicated to multicast and unicast traffic is desired.

Answer: AC

NEW QUESTION 15

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 20
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 24
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 27
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 32
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 36
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

Explanation: http://www.cisco.com/en/US/docs/ios/ipmulti/configuration/guide/imc_msdp_im_pim_sm.ht ml
When MSDP is enabled, an RP in a PIM-SM domain maintains MSDP peering relationships with MSDPenabled routers in other domains. This peering relationship occurs over a TCP connection, where primarily a list of sources sending to multicast groups is exchanged. MSDP uses TCP (port 639) for its peering connections. As with BGP, using point-to-point TCP peering means that each peer must be explicitly configured. The TCP connections between RPs, moreover, are achieved by the underlying routing system. The receiving RP uses the source lists to establish a source path. If the multicast sources are of interest to a domain that has receivers, multicast data is delivered over the normal, source- tree building mechanism provided by PIM-SM. MSDP is also used to announce sources sending to a group. These announcements must originate at the RP of the domain.

NEW QUESTION 37
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 38
Which two statements correctly describe the BGP ttl-security feature? (Choose two.)

A. This feature protects the BGP processes from CPU utilization-based attacks from EBGP neighbors which can be multiple hops away
B. This feature prevents IBGP sessions with non-directly connected IBGP neighbors
C. This feature will cause the EBGP updates from the router to be sent using a TTL of 1
D. This feature needs to be configured on each participating BGP router
E. This feature is used together with the ebgp-multihop command

Answer: AD

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

NEW QUESTION 40
Refer to the exhibit.

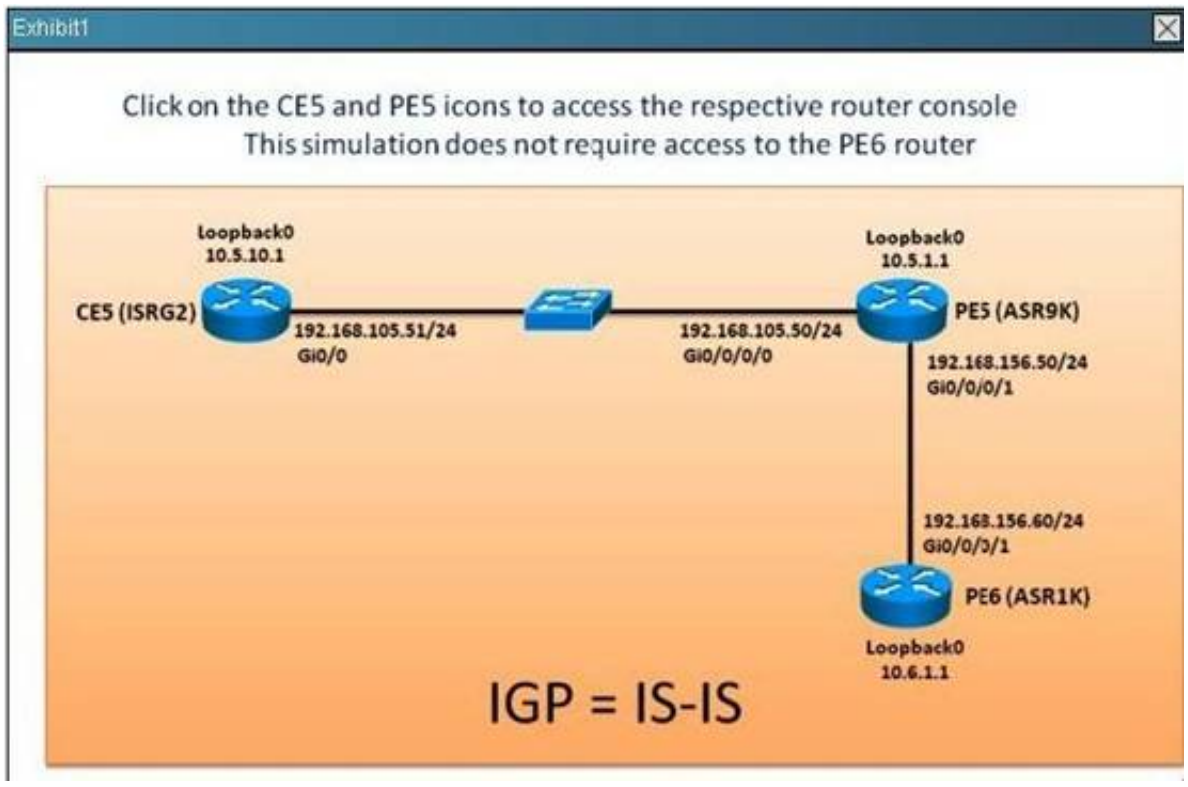

Which three statements are correct regarding the various multicast groups? (Choose three.)

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 43
Which multicast implementation is preferred for traffic that is required by a small number of receivers across a large distributed network?

A. DVMRP
B. PIM-DM
C. PIM-SM
D. IGMP

Answer: C

NEW QUESTION 45
Which configuration for implementing 6PE on an IS-IS-enabled Cisco IOS XR router is correct?

A. interface GigabitEthernet0/0/0/0 ipv6 address 2001:DB8:DD11::1/64 router isis ipv6-tunnet 49.0000.0000.00010.00address-family ipv6 unicast single-topologyredistribute bgp 200interface GigabitEthernet0/0/0/0 address-family ipv6 unicast router bgp 200bgp router-id 209.165.202.129 address-family ipv4 unicast address-family ipv6 unicast redistribute isis ipv6-tun neighbor 209.165.202.130remote-as 200address-family ipv4 unicast address-family ipv6 labeled-unicast
B. interface GigabitEthernet0/0/0/0 ipv6 address 2001:DB8:DD11::1/64 router isis ipv6-tunnet 49.0000.0000.00010.00address-family ipv6 unicast single-topologyrouter bgp 200bgp router-id 209.165.202.129 address-family ipv4 unicast address-family ipv6 unicast redistribute isis ipv6-tun neighbor 209.165.202.130remote-as 200address-family ipv4 unicast address-family ipv6 labeled-unicast
C. interface GigabitEthernet0/0/0/0 ipv6 address 2001:DB8:DD11::1/64 router isis ipv6-tunnet 49.0000.0000.00010.00address-family ipv6 unicast single-topologyinterface GigabitEthernet0/0/0/0 address-family ipv6 unicast router bgp 200bgp router-id 209.165.202.129 address-family ipv4 unicast address-family ipv6 unicast redistribute staticneighbor 209.165.202.130remote-as 200address-family ipv4 unicast address-family ipv6 labeled-unicast
D. interface GigabitEthernet0/0/0/0 ipv6 address 2001:DB8:DD11::1/64 router isis ipv6-tunnet 49.0000.0000.00010.00address-family ipv6 unicast single-topologyinterface GigabitEthernet0/0/0/0 address-family ipv6 unicast router bgp 200bgp router-id 209.165.202.129 address-family ipv4 unicast address-family ipv6 unicastredistribute connected redistribute isis ipv6-tun neighbor 209.165.202.130remote-as 200address-family ipv4 unicast address-family ipv6 labeled-unicast
E. interface GigabitEthernet0/0/0/0 ipv6 address 2001:DB8:DD11::1/64 router isis ipv6-tunnet 49.0000.0000.00010.00address-family ipv6 unicast single-topologyinterface GigabitEthernet0/0/0/0 address-family ipv6 unicast router bgp 200bgp router-id 209.165.202.129 address-family ipv4 unicast address-family ipv6 unicast redistribute connected redistribute isis ipv6-tun neighbor 209.165.202.130remote-as 200address-family ipv4 unicast

Answer: D

NEW QUESTION 50
DRAG DROP

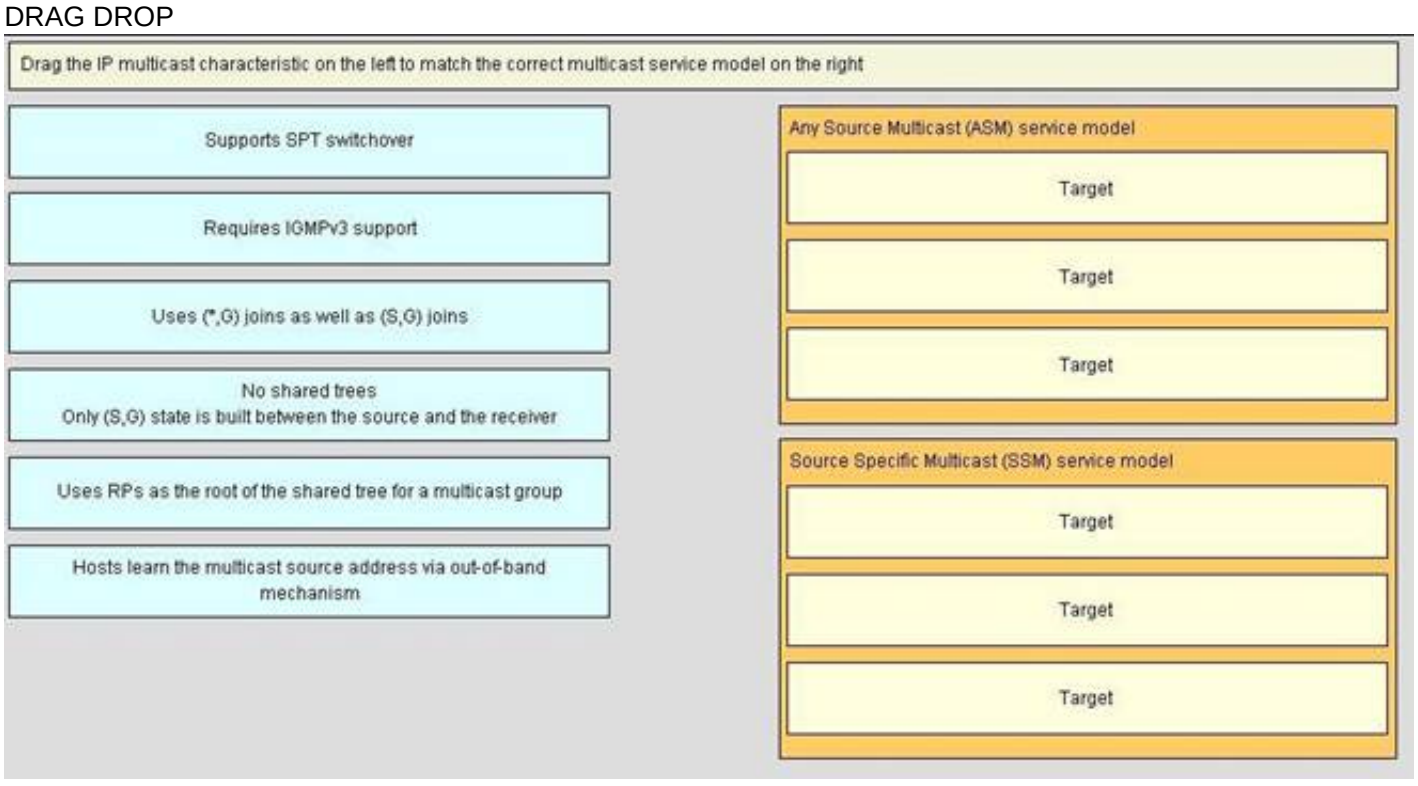

Answer:

Explanation: Any Source Multicast - Uses RP's as the root of the shared tree for a multicast group, ONly (S,G) state is build between the source and the recevier, Spport SPT Switchover Source Specific Multicast - Uses (*,G) joins as well as (S,G) Joins , Requires IGMPV3 Support,
Hosts learn the multicast source address via out-of-banf mechanism
i) Dense Mode Flood-and-Prune Protocols (DVMRP / MOSPF / PIM-DM)
In dense mode protocols, all routers in the network are aware of all trees, their sources and receivers.
Protocols such as DVMRP and PIM dense mode flood “active source” information across the whole network and build trees by creating “Prune State” in parts of the topology where traffic for a specific tree is unwanted.
They are also called flood-and-prune protocols. In MOSPF, information about receivers is flooded throughout the network to support the building of trees.
Dense mode protocols are undesirable because every tree built in some part of the network will always cause resource utilization (with convergence impact) on all routers in the network (or within the administrative scope, if configured). We will not be discussing these protocols in the rest of this paper.
ii) Sparse Mode Explicit Join Protocols (PIM-SM/PIM-BiDir)
With sparse mode explicit join protocols we do not create a group-specific forwarding state in the network unless a receiver has sent an explicit IGMP/MLD membership report (or “join”) for a group. This variant of ASM is known to scale well and is the multicast paradigm we will mainly be discussing. This is the basis for PIMSparse Mode, which most multicast deployments have used to this point. This is also the basis for PIM-BiDir, which will be increasingly deployed for MANY (sources) TO MANY (receivers) applications.

These protocols are called sparse mode because they efficiently support IP multicast delivery trees with a “sparse” receiver population – creating control plane state only on routers in the path between sources and receivers, and in PIM-SM/BiDir, the Rendezvous Point (RP). They never create state in other parts of the network. State in a router is only built explicitly when it receives a join from a downstream router or receiver, hence the name “explicit join protocols”.
Both PIM-SM and PIM-BiDir employ “SHARED TREES”, which allow traffic from any source to be forwarded to a receiver. The forwarding state on a shared tree is referred to as (*,G) forwarding state, where the * is a wild card for ANY SOURCE. Additionally, PIM- SM supports the creation of forwarding state that relates to traffic from a specific source. These are known as SOURCE TREES, and the associated state is referred to as (S, G) forwarding state SSM is the model used when the receiver (or some proxy) sends (S,G) “joins” to indicate that it wants to receive traffic sent by source S to group G. This is possible with IGMPv3/MLDv2 “INCLUDE” mode membership reports. We therefore refer to this model as the Source-Specific Multicast (SSM) model. SSM mandates the use of an explicit-join protocol between routers. The standard protocol for this is PIM-SSM, which is simply the subset of PIM-SM used to create (S,G) trees. There are no shared trees (*,G) state in SSM. Multicast receivers can thus “join” an ASM group G, or “join” (or more accurately “subscribe” to) an SSM (S, G) channel. To avoid having to repeat the term “ASM group or SSM channel”, we will use the term (multicast) flow in the text, implying that the flow could be an ASM group or an SSM channel

NEW QUESTION 53
Which multicast routing protocol is most optimal for supporting many-to-many multicast applications?

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

Answer: B

Explanation: PIM-Bidirectional Operations
PIM Bidirectional (BIDIR) has one shared tree from sources to RP and from RP to receivers. This is unlike the PIM-SM, which is unidirectional by nature with multiple source trees - one per (S, G) or a shared tree from receiver to RP and multiple SG trees from RP to sources.
Benefits of PIM BIDIR are as follows:
• As many sources for the same group use one and only state (*, G), only minimal states are required in each router.
• No data triggered events.
• Rendezvous Point (RP) router not required. The RP address only needs to be a routable address and need not exist on a physical device.

NEW QUESTION 57
Refer to the EBGP configuration on a PE IOS-XR router exhibit.
After the EBGP configuration, no routes are accepted from the EBGP peer, nor are any routes advertised to the EBGP peer.

```
router bgp 65001
 address-family ipv4 unicast
  network 172.16.1.0/24
  network 192.168.1.0/24
 !
 neighbor 10.1.1.1
  remote-as 65002
 !
```

What could be the problem?

A. The update-source neighbor configuration command must also be configured
B. The next-hop-self neighbor configuration command must also be configured
C. EBGP neighbors must have an inbound and outbound route policy configured
D. An access list is blocking IP protocol 179 packets between the two EBGP peers
E. The maximum-prefix neighbor configuration command must also be configured

Answer: C

NEW QUESTION 59
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 64
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

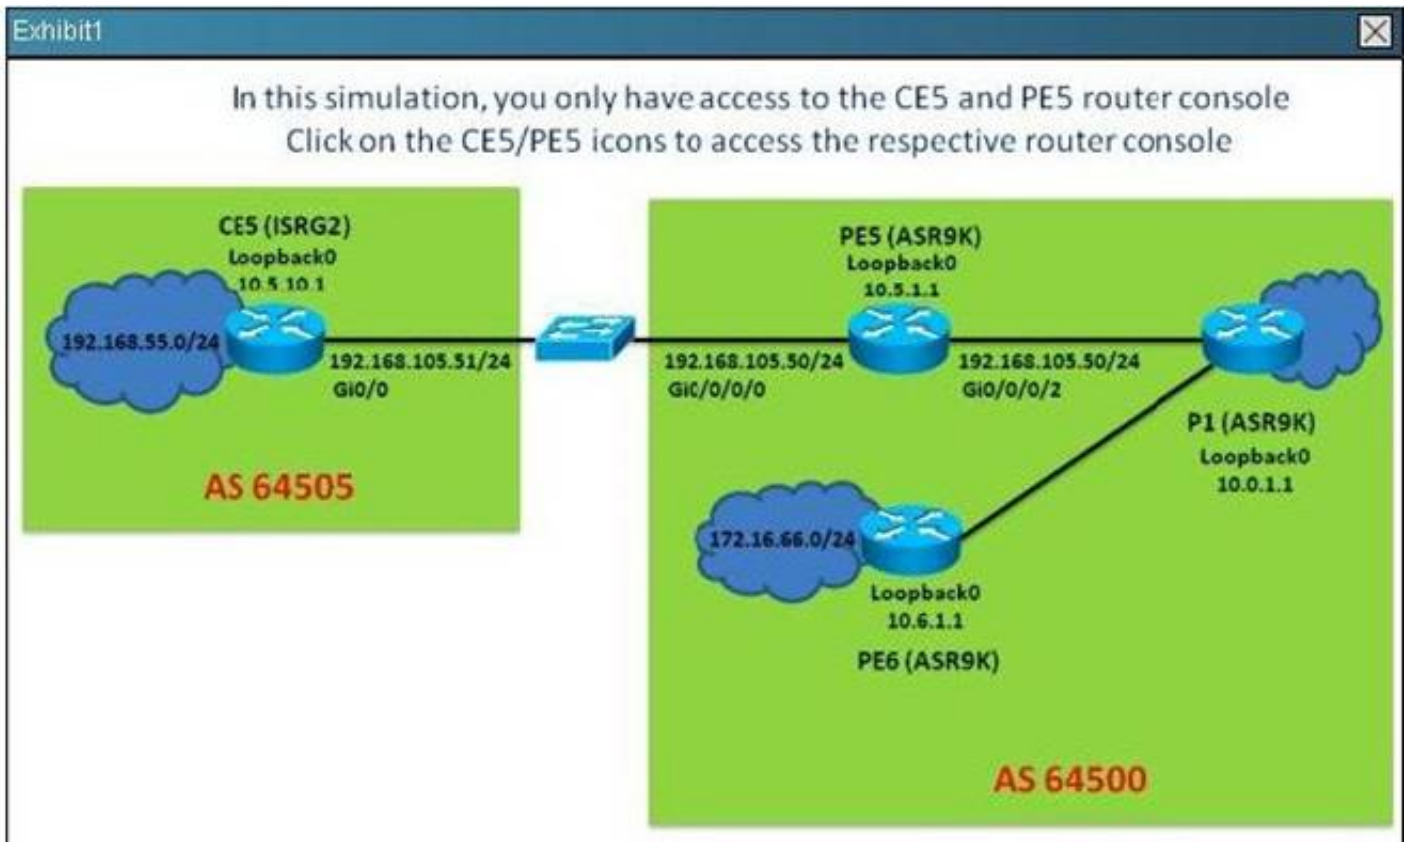

CE5

CE5#

On the PE5 router, which statementis correct regarding the learned BGP prefixes?

A. The 209.165.201.0/27 prefix is received from the 10.0.1.1 IBGP peer which is a route reflector
B. The 172.16.66.0/24 prefix BGP next-hop points to the route reflector
C. All prefixes learned on PE5 has the default local preference value
D. The 209.165.202.128/27 prefix is originated by the 10.0.1.1 IBGP peer

Answer: C

Explanation: #show ip bgp -- check i tag for PE5

NEW QUESTION 68
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 74
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 77
In Cisco IOS-XR, the maximum-prefix command, to control the number of prefixes that can be installed from a BGP neighbor, is configured under which configuration mode?

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

Answer: D

Explanation: http://www.cisco.com/en/US/tech/tk365/technologies_configuration_example09186a00801 0a28a.shtml

NEW QUESTION 82
Refer to the exhibit.

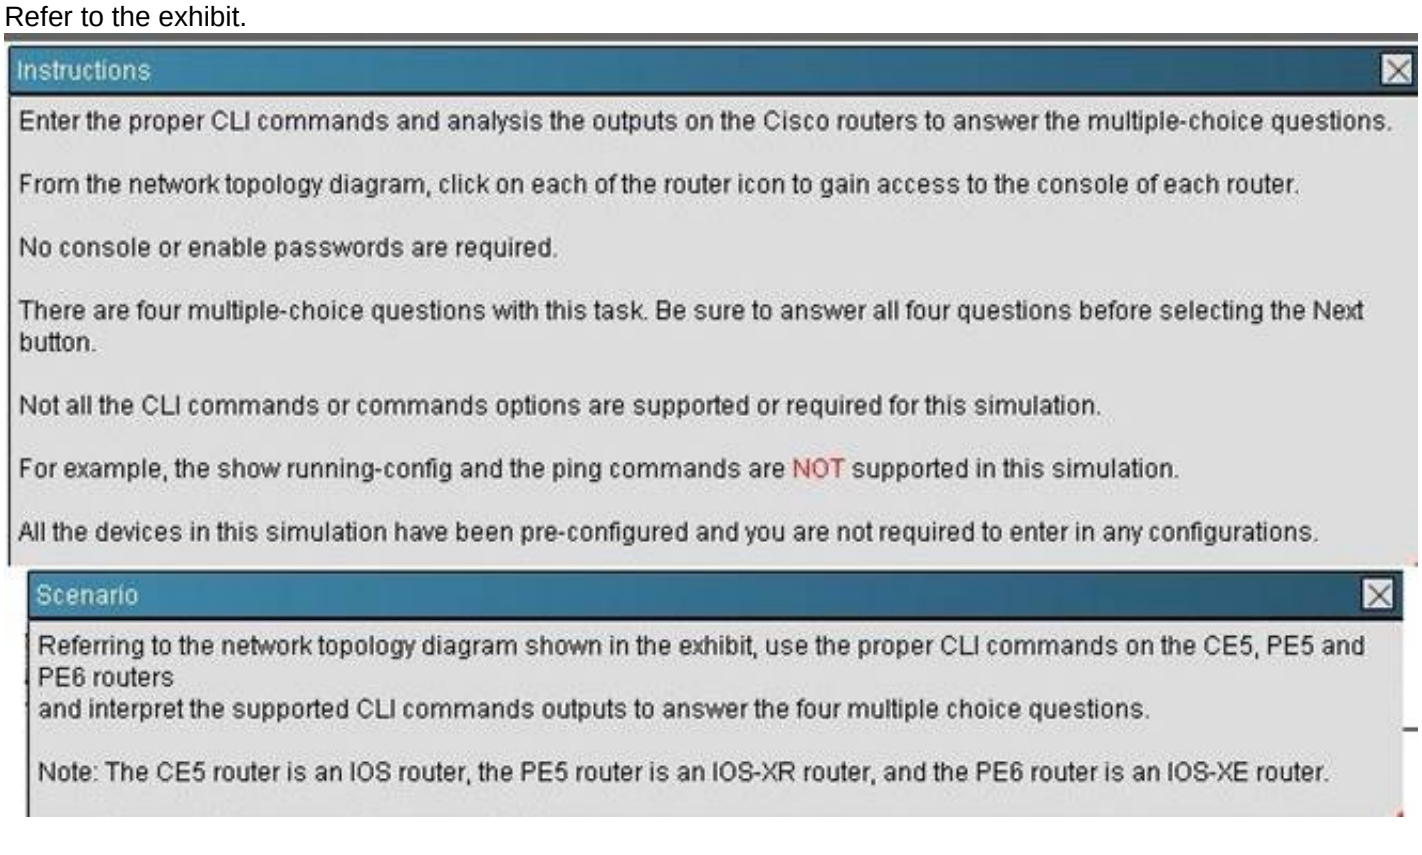

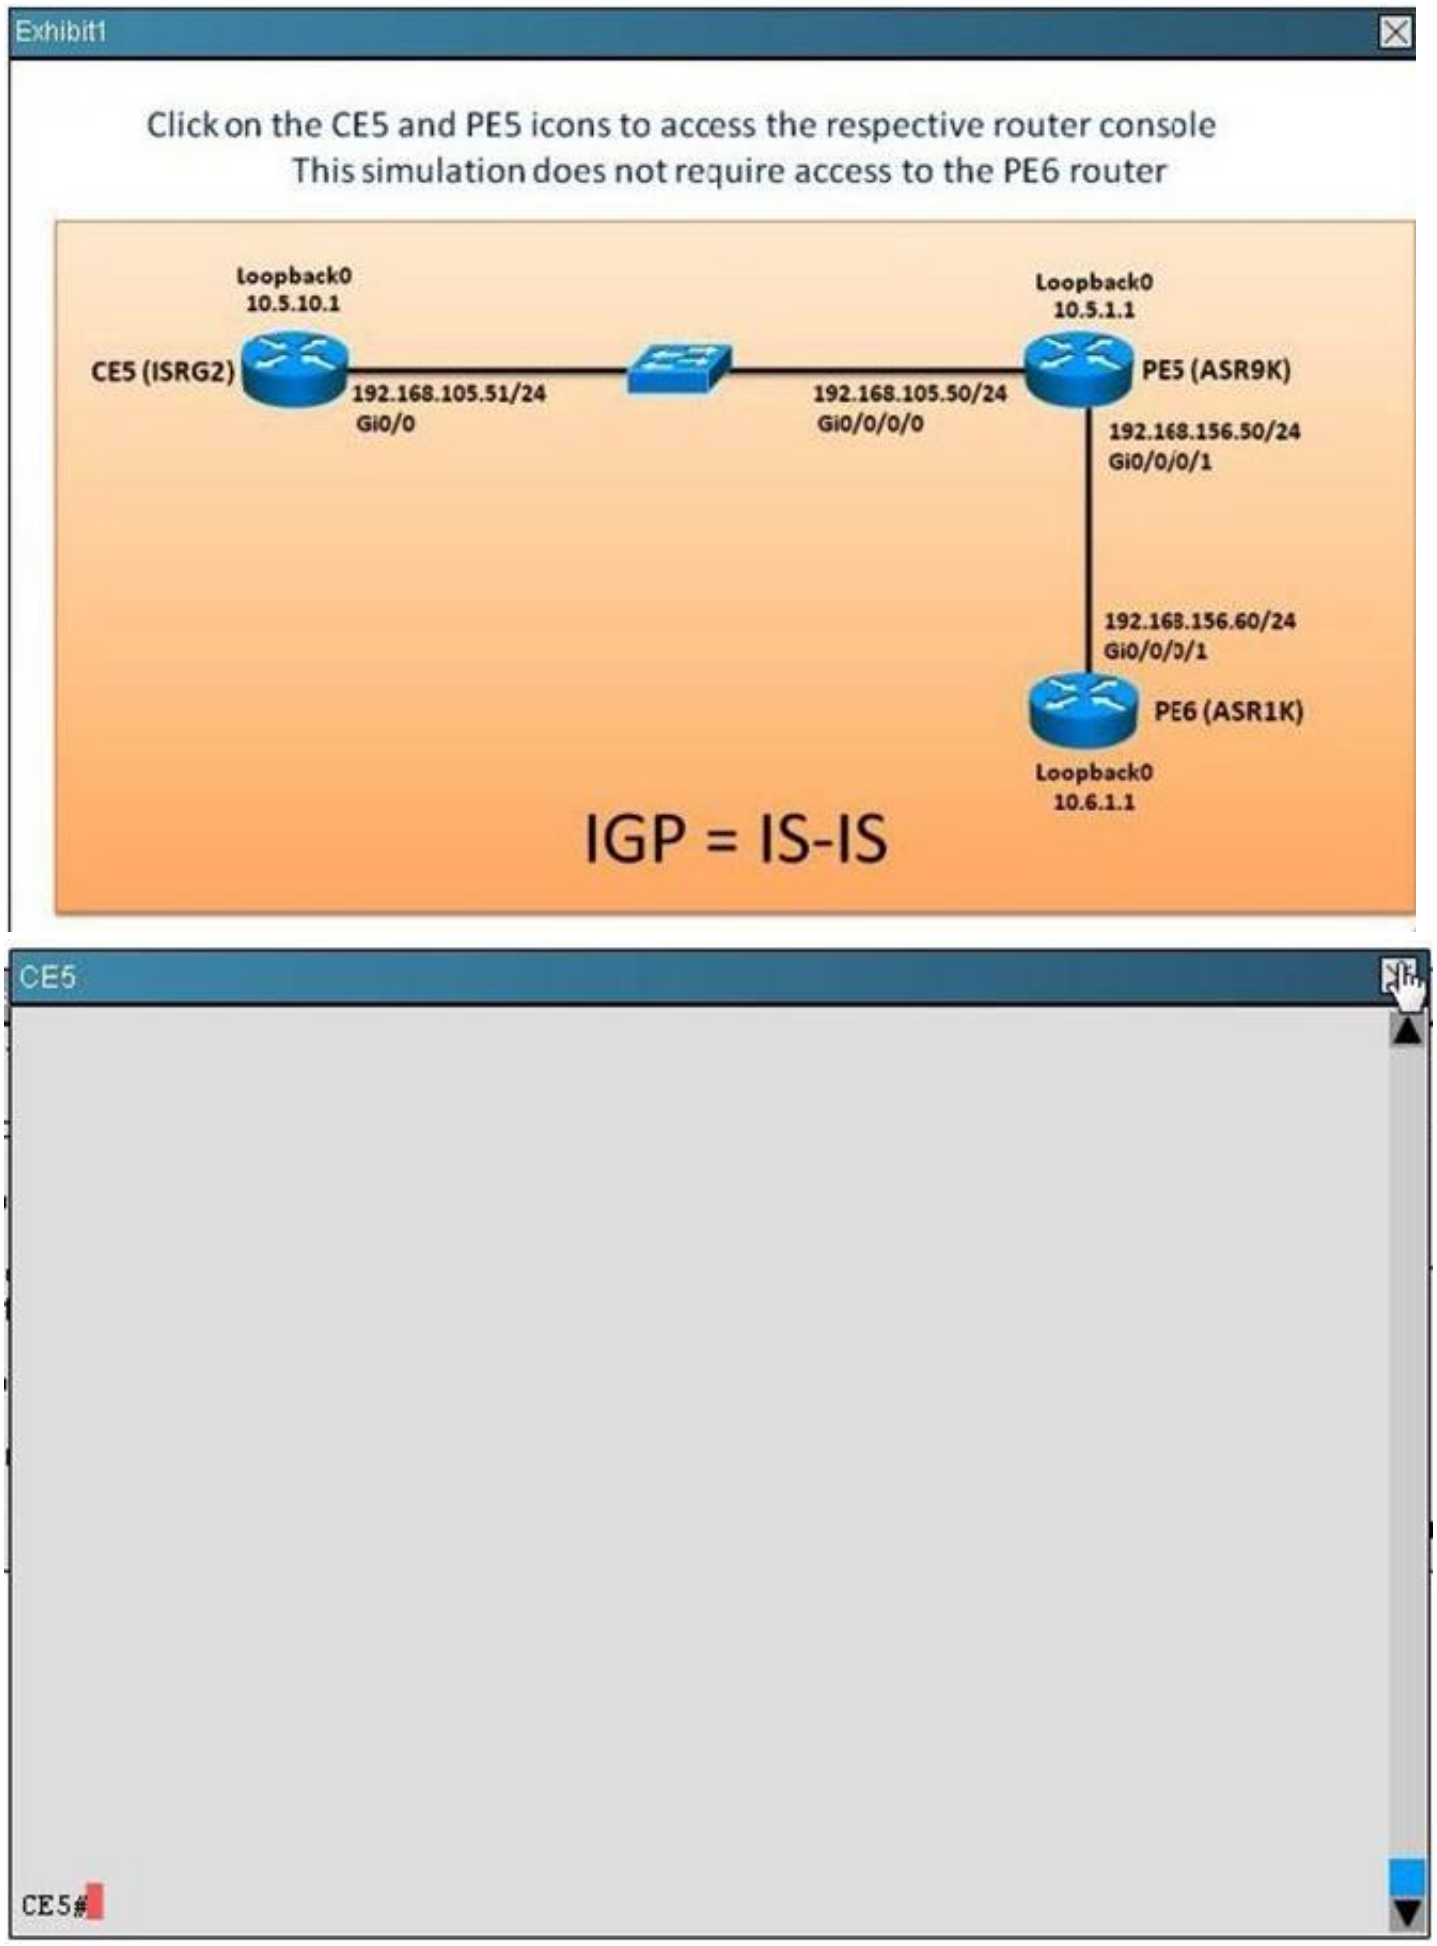
Exhibit1
Click on the CE5 and PE5 icons to access the respective router console
This simulation does not require access to the PE6 router
Loopback0
10.5.10.1
CE5 (ISRG2)
192.168.105.51/24
Gi0/0
192.168.105.50/24
Gi0/0/0/0
Loopback0
10.5.1.1
PE5 (ASR9K)
192.168.156.50/24
Gi0/0/0/1
192.168.156.60/24
Gi0/0/0/1
PE6 (ASR1K)
Loopback0
10.6.1.1
IGP = IS-IS
CE5
CE5#

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 86
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 88
Which two specific characteristics categorize traceroute in an IPv6 routing environment? (Choose two.)

A. Traceroute can show the path to reach any destination IPv6 address.
B. Traceroute returns an error for a link-local IPv6 address.
C. Traceroute is based on ICMPv6 Type 1 (Destination Unreachable) reply packets to determine the network path.
D. Traceroute is based on ICMPv6 Type 3 (Time Exceeded) reply packets to determine the network path.
E. Traceroute is based on ICMPv6 Type 2 (Packet Too Big) reply packets to determine the network path.
F. Traceroute for IPv6 implements a backwards compatibility option to provide a detailed report in environments running dual-stack.

Answer: AD

NEW QUESTION 91
The 224.192.16.1 multicast IP address maps to which multicast MAC address?

A. 01-00-5E-C0-10-01
B. 01-00-5E-40-10-01
C. 01-00-5E-00-10-01
D. 01-00-5E-C0-16-01

Answer: B

Explanation: Least significant 23 bits of IP address and pre-pend 01-00-5E
224 ignore
192 less 128 becomes 64 = 40
16 = 10
1 = 01
01-00-5E-40-10-01

NEW QUESTION 94
Which type of DNS record is used for IPv6 forward lookups?

A. A records
B. AAAA records
C. PTR records
D. MX records

Answer: B

NEW QUESTION 95
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 100
Which two BGP mechanisms are used to prevent routing loops when using a design with redundant route reflectors? (Choose two.)

A. Cluster-list
B. AS-Path
C. Originator ID
D. Community
E. Origin

Answer: AC

Explanation: http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.7/routing/configuration/guide/rc37bgp.html
As the iBGP learned routes are reflected, routing information may loop. The route reflector model has the following mechanisms to avoid routing loops:
•Originator ID is an optional, nontransitive BGP attribute. It is a 4-byte attributed created by a route reflector.
The attribute carries the router ID of the originator of the route in the local autonomous system. Therefore, if a misconfiguration causes routing information to come back to the originator, the information is ignored.
•Cluster-list is an optional, nontransitive BGP attribute. It is a sequence of cluster IDs that the route has passed. When a route reflector reflects a route from its clients to nonclient peers, and vice versa, it appends the local cluster ID to the cluster-list. If the cluster-list is empty, a new cluster-list is created. Using this attribute, a route reflector can identify if routing information is looped back to the same cluster due to misconfiguration. If the local cluster ID is found in the cluster-list, the advertisement is ignored.

NEW QUESTION 104
Which types of multicast distribution tree can PIM-SM use?

A. Only shared tree rooted at the source
B. Only shared tree rooted at the RP
C. Only shortest path tree rooted at the RP
D. Shared tree rooted at the source and shortest path tree switchover
E. Shared tree rooted at the RP and shortest path tree switchover
F. Shared tree rooted at the first-hop router and shortest path tree rooted at the RP

Answer: E

NEW QUESTION 107
Which two features are used to provide high availability multicast? (Choose two.)

A. BFD
B. NSF/SSO
C. PIM NSR
D. PIM triggered join
E. IGMP triggered report
F. MSDP

Answer: BD

Explanation: Triggered joins are sent when the primary or the secondary RPF information changes. No RPF change prunes are sent for MoFRR streams.
mofrr
To perform a fast convergence (multicast-only fast reroute, or MoFRR) of specified routes/flows when a failure is detected on one of multiple equal-cost paths between the router and the source, use the mofrr command under PIM configuration mode.

mofrr rib acl_name no rib acl_name

NEW QUESTION 111
Which protocol can be used to secure multicast in a group multicast solution where group key management is needed for secure key exchange?

A. DOI
B. ISAKMP
C. GDOI
D. IPsec

Answer: C

NEW QUESTION 114
Which Cisco IOS XR command setssuccessfully configure a value of 20 for the advertisement-interval?

A. RP/0/RSP0/CPU0:routerconfig)# router bgp 65512 RP/0/RSP0/CPU0:router(config-bgp)# session-group test RP/0/RSP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 20 RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit RP/0/RSP0/CPU0:router(config-bgp)# neighbor-group test RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# advertisement-interval 25 RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# exit RP/0/RSP0/CPU0:router(config-bgp)# exit RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.1.1RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 65513 RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group test RP/0/RSP0/CPU0:router(config-bgp-nbr)# use neighbor-group test
B. RP/0/RSP0/CPU0:routerconfig)# router bgp 65512 RP/0/RSP0/CPU0:router(config-bgp)# session-group test RP/0/RSP0/CPU0:router(config-bgp-sngrp)# ebgp-multihop 2 RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit RP/0/RSP0/CPU0:router(config-bgp)# neighbor-group test RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# advertisement-interval 20 RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# exit RP/0/RSP0/CPU0:router(config-bgp)# exit RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.1.1RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 65513 RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group test RP/0/RSP0/CPU0:router(config-bgp-nbr)# use neighbor-group test
C. RP/0/RSP0/CPU0:routerconfig)# router bgp 65512 RP/0/RSP0/CPU0:router(config-bgp)# session-group test RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit RP/0/RSP0/CPU0:router(config-bgp)# neighbor-group test RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# exit RP/0/RSP0/CPU0:router(config-bgp)# exit RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.1.1RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 65513 RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group test RP/0/RSP0/CPU0:router(config-bgp-nbr)# use neighbor-group test
D. RP/0/RSP0/CPU0:routerconfig)# router bgp 65512 RP/0/RSP0/CPU0:router(config-bgp)# session-group test RP/0/RSP0/CPU0:router(config-bgp-sngrp)# advertisement-interval 25 RP/0/RSP0/CPU0:router(config-bgp-sngrp)# exit RP/0/RSP0/CPU0:router(config-bgp)# neighbor-group test RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# advertisement-interval 20 RP/0/RSP0/CPU0:router(config-bgp-nbrgrp)# exit RP/0/RSP0/CPU0:router(config-bgp)# exit RP/0/RSP0/CPU0:router(config-bgp)# neighbor 192.168.1.1RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 65513 RP/0/RSP0/CPU0:router(config-bgp-nbr)# use session-group test RP/0/RSP0/CPU0:router(config-bgp-nbr)# use neighbor-group test

Answer: A

NEW QUESTION 116
Refer to the Cisco IOS DHCPv6 configuration shown in the exhibit.

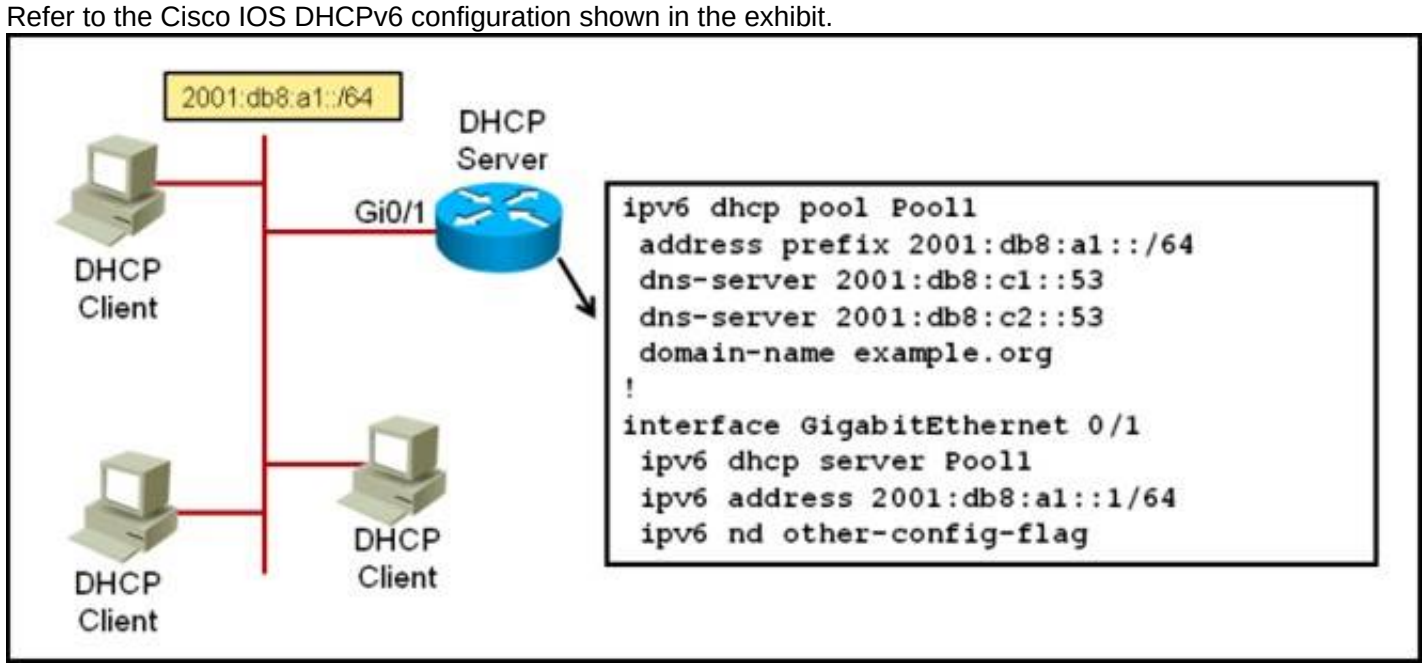

Which statement is correct?

A. The configuration is missing a command under interface Gi0/1 to indicate to the attached hosts to use stateful DHCPv6 to obtain their IPv6 addresses
B. The IPv6 router advertisements indicate to the attached hosts on the Gi0/1 interface to get other information besides their IPv6 address via stateless auto configuration
C. The IPv6 DHCPv6 server pool configuration is misconfigured
D. The DNS server address can also be imported from another upstream DHCPv6 server

Answer: A

Explanation: Server Configuration
In Global Configuration Mode ipv6 unciast-routing
ipv6 dhcp pool <pool name>
address prefix <specify address prefix> lifetime <infinite> <infinite> dns-server <specify the dns server address>
domain-name <specify the domain name> exit
In Interface Configuration Mode
ipv6 address <specify IPv6 Address>
ipv6 dhcp server <server name>rapid-commit Client Configuration
In Global Configuration Mode enable
configure terminal ipv6 unicast-routing
In Interface Configuration Mode ipv6 address dhcp rapid commit ipv6 enable

exit

NEW QUESTION 118
Which two methods represent IPv6 tunneling implementations? (Choose two.)

A. IPv6 over GRE tunneling
B. manually configured tunnels
C. automatic tunnels
D. 6to4 tunneling
E. IPv6 over an IPv4 tunnel over MPLS

Answer: BC

NEW QUESTION 123
A network engineer of an ISP using Cisco IOS XR routers wants to limit the number of prefixes that BGP peers can accept. To accomplish this task, the command maximum- prefix 1000 is used. Which two results of this configuration are expected? (Choose two.)

A. A warning message displays by default when 750 prefixes are received.
B. A warning message displays by default when 850 prefixes are received.
C. A BGP peer resets when it receives 1001 prefixes.
D. A BGP peer resets when it receives 1000 prefixes.
E. A BGP peer ceases when it receives 1001 prefixes.
F. A BGP peer ceases when it receives 1000 prefixes.
G. The BGP peer tries to reestablish the session after one minute.

Answer: AE

NEW QUESTION 126
Which two options are thecommon methods for implementing Site of Origin on Cisco IOS XE routers for loop avoidance in multihome BGP customers? (Choose two.)

A. Configure the route-map in command on the CE BGP neighbor.
B. Configure Site of Origin directly on the CE BGP neighbor command.
C. Configure site-map on VRF interface and redistribution of iBGP.
D. Configure site-map on VRF interface and network command.
E. Configure the route-map out command on the P router.

Answer: AB

NEW QUESTION 127
On Cisco IOS-XR, which BGP configuration group allows you to define address-family independent commands and address-family dependent commands for each address family?

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

Answer: A

Explanation: •Commands relating to a peer group found in Cisco IOS Release 12.2 have been removed from Cisco IOS XR software. Instead, the af-group, session-group, and neighbor-group configuration commands are added to support the neighbor in Cisco IOS XR software:
–The af-group command is used to group address family-specific neighbor commands within an IPv4 or IPv6 address family. Neighbors that have the same address family configuration are able to use the address family group name for their address family- specific configuration. A neighbor inherits the configuration from an address family group by way of the use command. If a neighbor is configured to use an address family group, the neighbor will (by default) inherit the entire configuration from the address family group. However, a neighbor will not inherit all ofthe configuration from the address family group if items are explicitly configured for the neighbor.
–The session-group command allows you to create a session group from which neighbors can inherit address family-independent configuration. A neighbor inherits the configuration from a session group by way of the use command. If a neighbor is configured to use a session group, the neighbor (by default) inherits the session group's entire configuration. A neighbor does not inherit all the configuration from a session group if a configuration is done directly on that neighbor.
–The neighbor-group command helps you apply the same configuration to one or more neighbors. Neighbor groups can include session groups and address family groups. This additional flexibility can create a complete configuration for a neighbor. Once a neighbor group is configured, each neighbor can inherit the configuration through the use command. If a neighbor is configured to use a neighbor group, the neighbor (by default) inherits the neighbor group's entire BGP configuration.
–However, a neighbor will not inherit all of the configuration from the neighbor group if items are explicitly configured for the neighbor. In addition, some part of the neighbor group's configuration could be hidden if a session group or address family group was also being used

NEW QUESTION 130
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 131
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

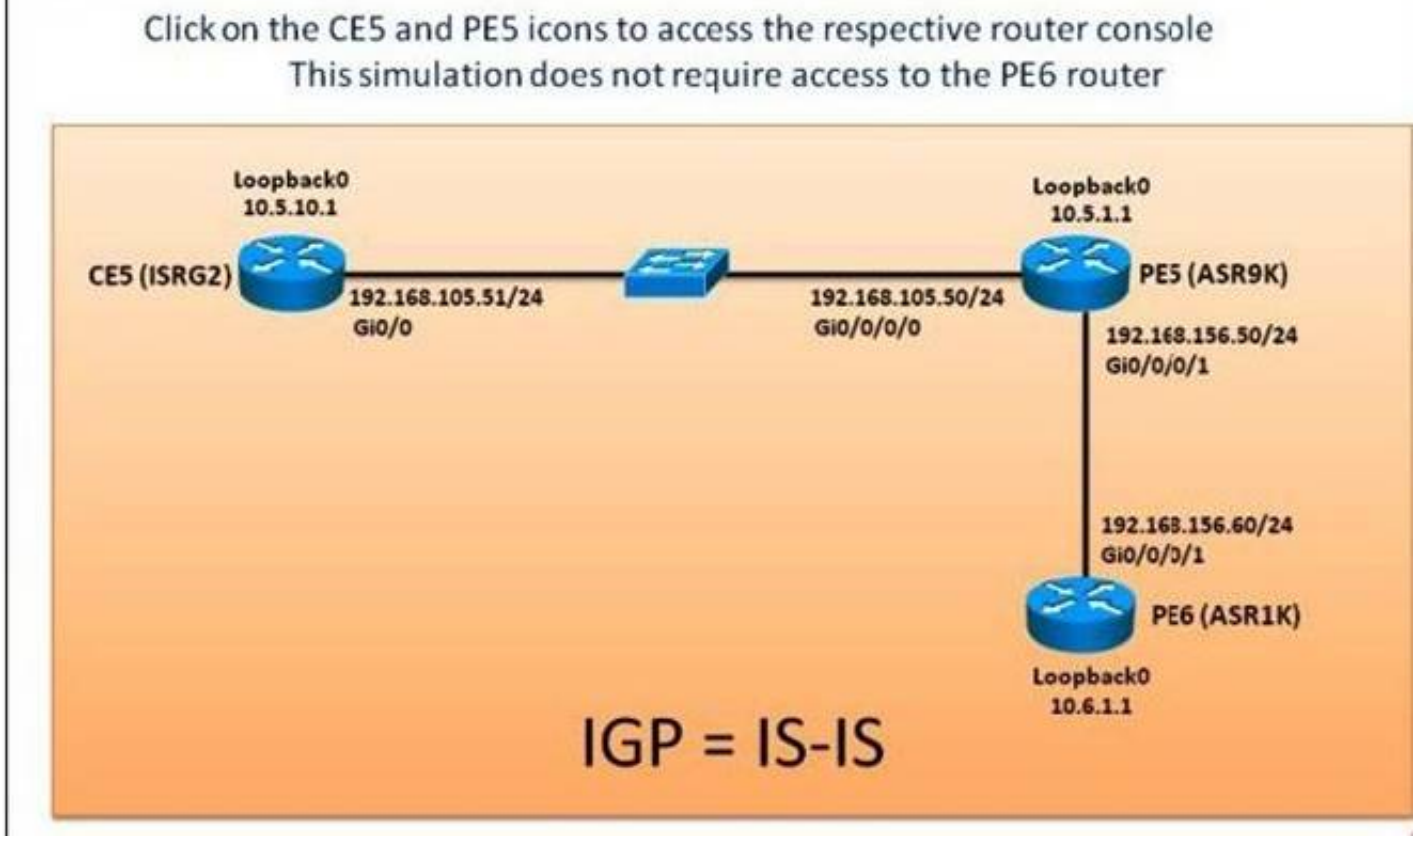

Which router is configured as the RPforthe 234.1.1.1 multicast group and which Is the multicast source that is currently sending traffic to the 234.1.1.1 multicast group? (Choose two.)

A. CE5
B. PE5
C. PE6
D. 10.5.10.1
E. 10.5.1.1
F. 192.168.156.60

Answer: CE

Explanation: #show ip mroute234.1.1.1
#show ip route

NEW QUESTION 133
Refer to the exhibit.

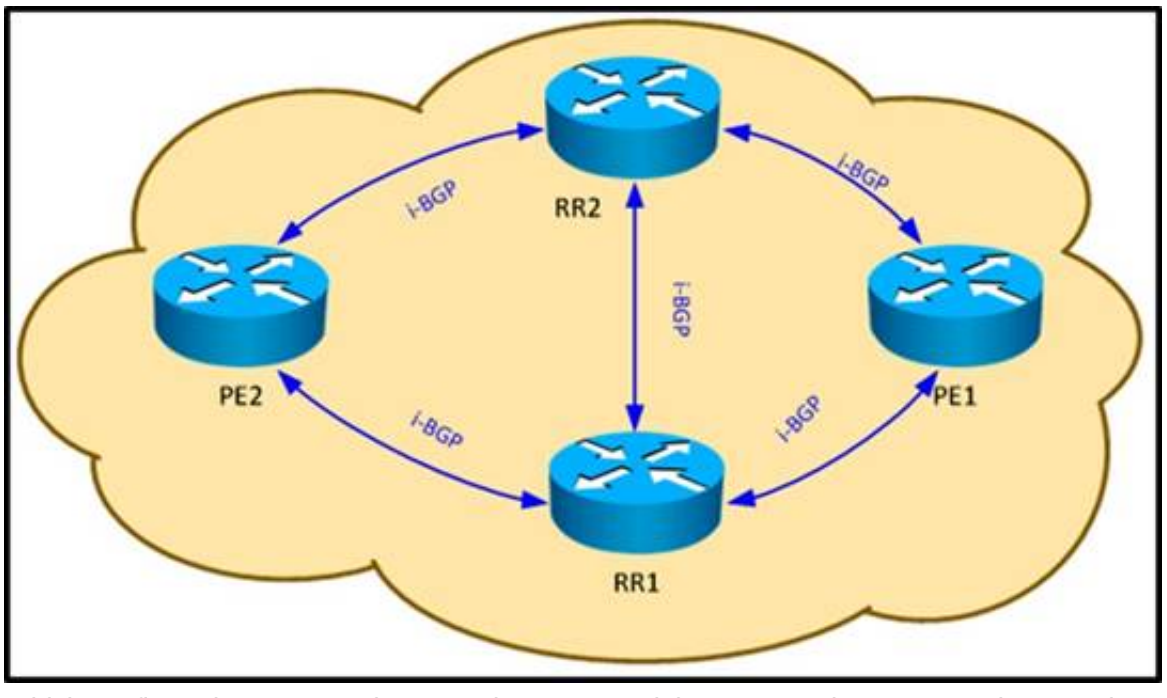

Which configuration ensures that RR2 does not send the same updates to PE2 that RR1 learns via PE1?

A. RR1 and RR2 should have different router IDs.
B. RR1 and RR2 should have different originator IDs.
C. RR1 and RR2 should have the same router IDs.
D. RR1 and RR2 should have the same cluster IDs.

Answer: D

NEW QUESTION 136
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 138
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 140
Which three statements are correct regarding PIM-SM? (Choose three.)

A. There are three ways to configure the RP: Static RP, Auto-RP, or BSR
B. PIM-SM only uses the RP rooted shared tree and has no option to switch over to the shortest path tree
C. Different RPs can be configured for different multicast groups to increase RP scalability
D. Candidate RPs and RP mapping agents are configured to enable Auto-RP
E. PIM-SM uses the implicit join model

Answer: ACD

NEW QUESTION 143
Refer to the exhibit.

```
RP/0/0/CPU0:R1# sh ip bgp nei | i time
Thu Jun 26 17:55:20.919 UTC
   Hold time is 90, keepalive interval is 30 seconds
   Configured hold time: 180, keepalive: 60, min acceptable holt time: 3
   Minimum time between advertisement runs is 30 secs
!
RP/0/0/CPU0:R3# sh ip bgp nei | i time
Thu Jun 26 17:55:34.109 UTC
   Hold time is 90, keepalive interval is 30 seconds
   Configured hold time: 90, keepalive: 30, min acceptable holt time: 3
   Minimum time between advertisement runs is 30 secs
```

Based on the output of two eBGP adjacent neighbors, which command can be used to set up the default BGP timers?

A. RP/0/0/CPU0:R1(config-bgp)#timers bgp 60 30
B. RP/0/0/CPU0:R2(config-bgp)#timers bgp 30 60
C. RP/0/0/CPU0:R2(config-bgp-nbr)#timers bgp 180 60
D. RP/0/0/CPU0:R2(config-bgp)#timers bgp 60 180
E. RP/0/0/CPU0:R1(config-bgp)#timers bgp 60 180

Answer: D

NEW QUESTION 146
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 150
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.

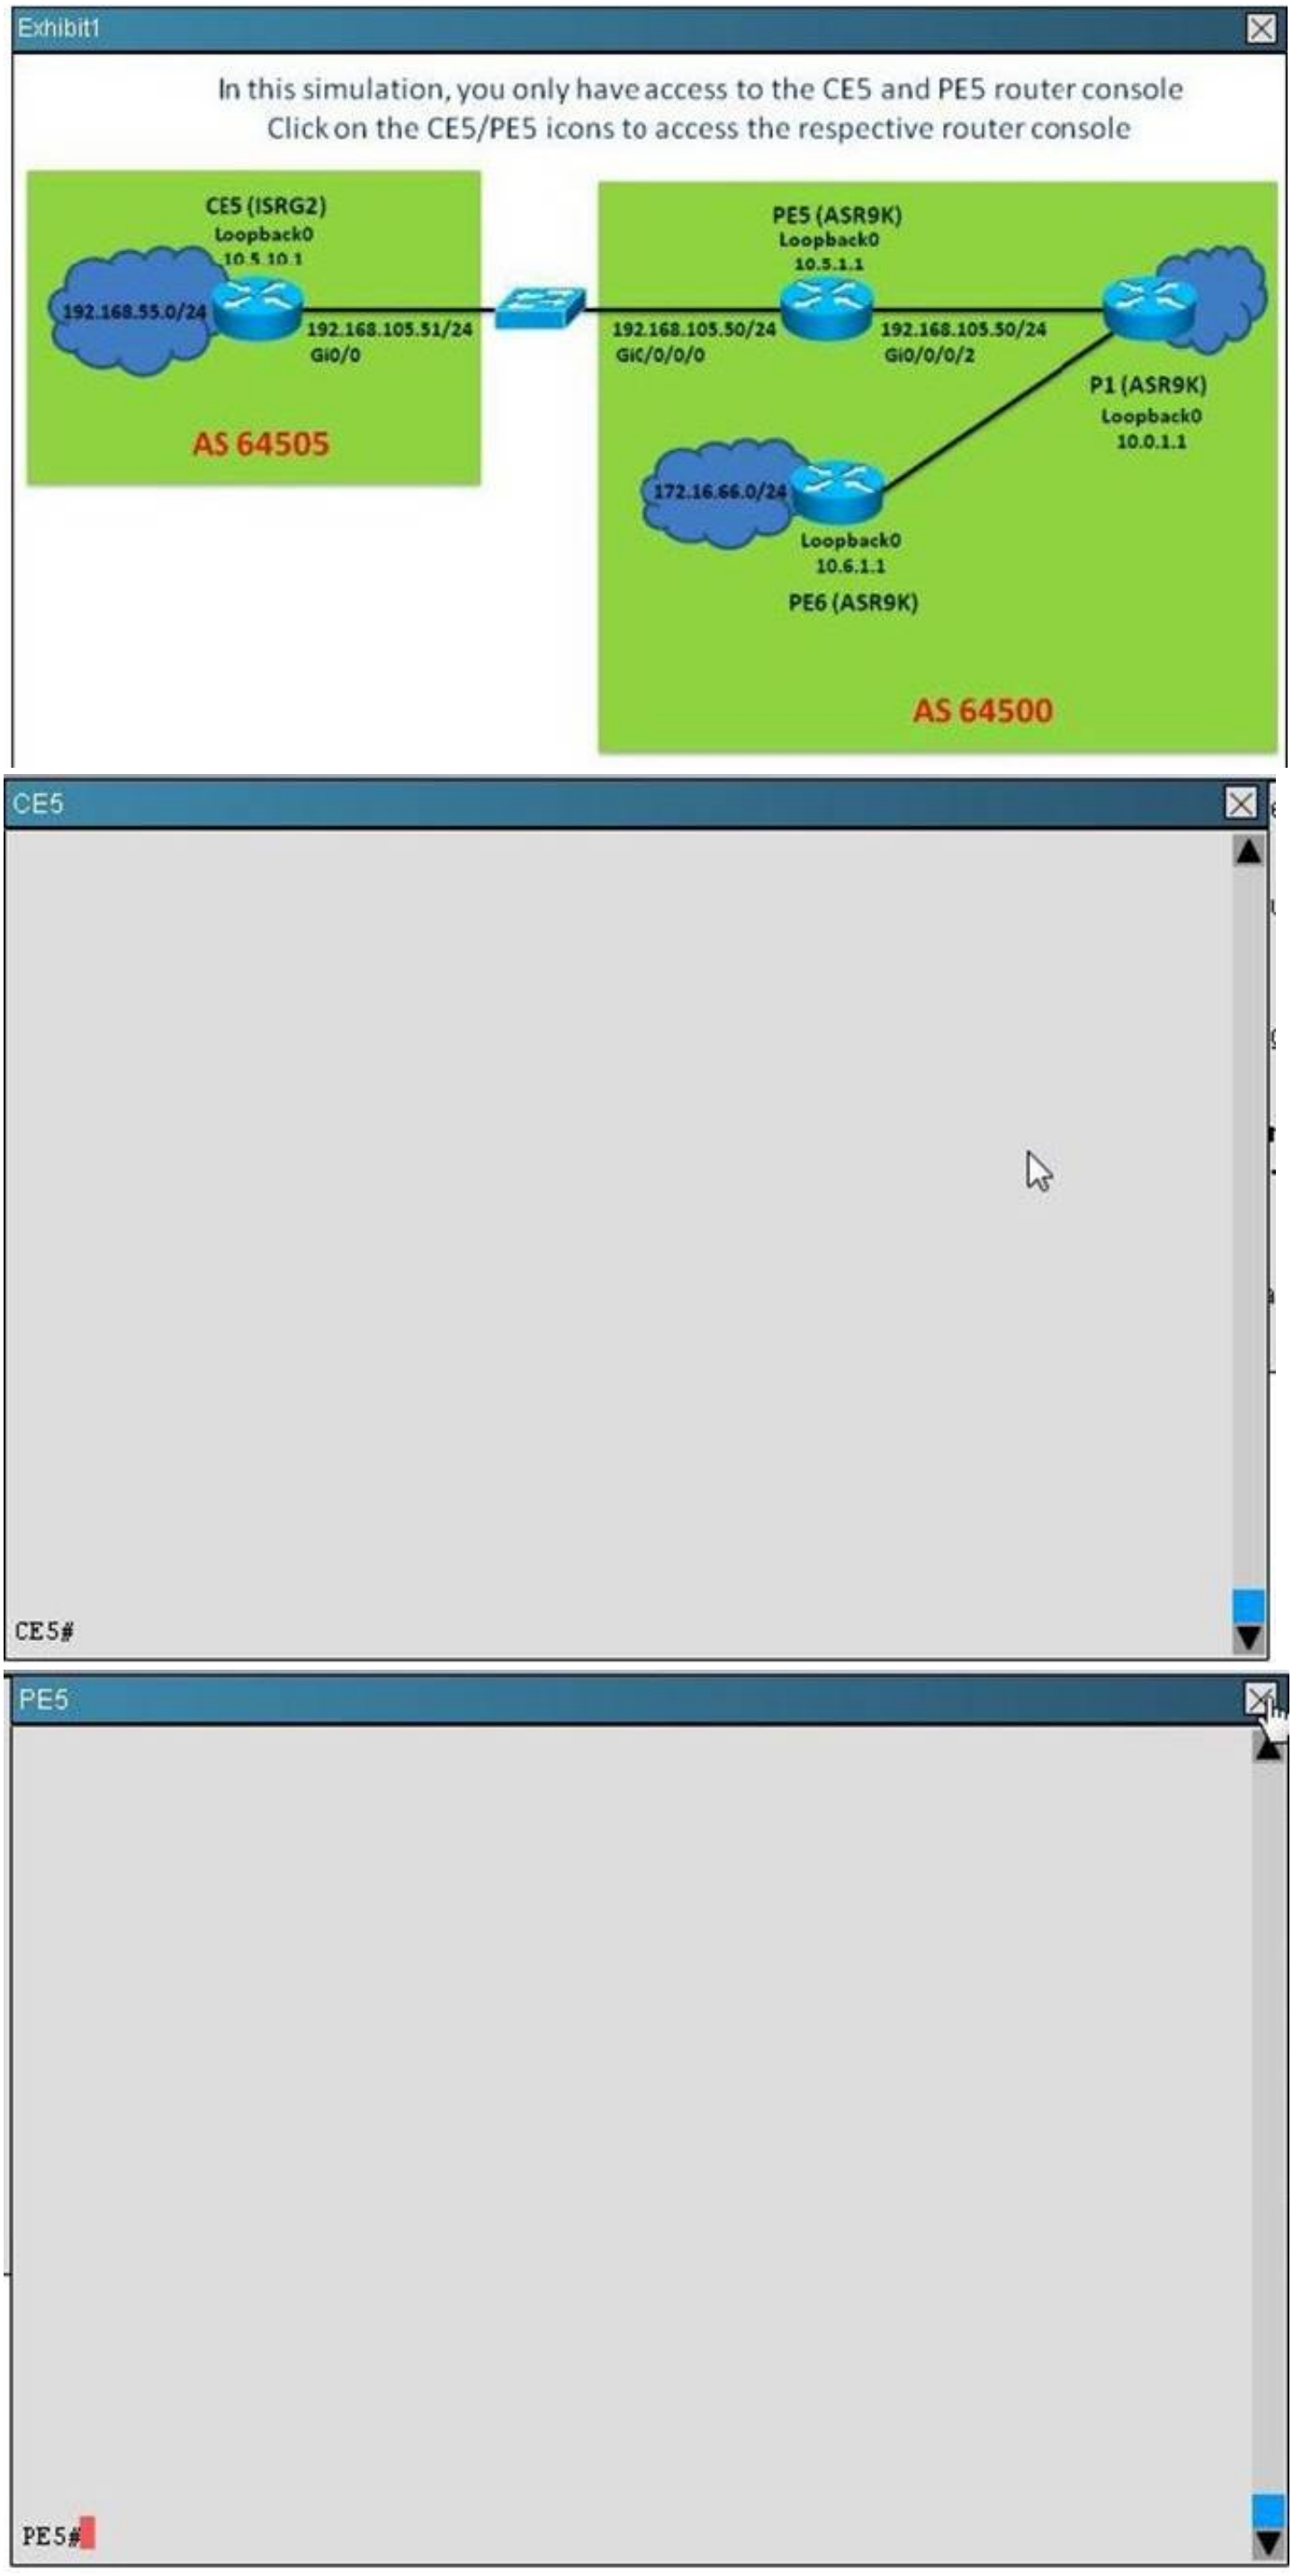

Which two statements regarding the BGP peerlngs are correct? (Choose two)

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 153
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 158
A network engineer must deploy an iBGP-based cloud region configuration by means of templates to reduce the overall BGP CLI required. Which three commands represent a basic configuration for a BGP peer session template on a regular Cisco IOS instance? (Choose three.)

A. template peer-session session-template-name
B. remote-as as-number
C. neighbor-family config template
D. peer-family config template
E. as-override
F. timers keepalive-interval hold-time

Answer: ABF

NEW QUESTION 162
When implementing Anycast RP, the RPs are also required to establish which kind of peering with each other?

A. BGP
B. Multiprotocol BGP
C. MSDP
D. Bidirectional PIM
E. PIM SSM

Answer: C

Explanation: http://www.cisco.com/en/US/docs/ios/solutions_docs/ip_multicast/White_papers/anycast.ht ml
Using Anycast RP is an implementation strategy that provides load sharing and redundancy in Protocol Independent Multicast sparse mode (PIM-SM) networks. Anycast RP allows two or more rendezvous points (RPs) to share the load for source registration and the ability to act as hot backup routers for each other. Multicast Source Discovery Protocol (MSDP) is the key protocol that makes Anycast RP possible.

NEW QUESTION 166
When configuring BFD, the multiplier configuration option is used to determine which value?

A. The retry interval
B. The number of BFD packets that can be lost before the BFD peer is declared "down"
C. The minimum interval between packets accepted from the BFD peers
D. The number of BFD echo packets that will be originated by the router
E. The number of routing protocols that will use BFD for fast peer failure detection

Answer: B

NEW QUESTION 170
Refer to the configuration exhibit, taken from a Cisco IOS-XR router.

```
!
router static
address-family ipv4 unicast
192.0.2.1/32 Null0
!
route-policy RTBH
if tag is 666 then
set next-hop 192.0.2.1
endif
end-policy
!
router bgp 65123
address-family ipv4 unicast
redistribute static route-policy RTBH
!
!When attacks are detected from 209.165.201.144/28
!
router static
address-family ipv4 unicast
209.165.201.144/28 null0 tag 666
!
```

Which configuration change is required to properly enable this router as the signaling router for implementing source-based RTBH filtering?

A. Set community (no-export) in the route policy
B. Pass in the route policy
C. Set local-preference 1000 in the route policy
D. The 192.0.2.1/32 static route should be tagged as 666 (tag 666)

Answer: A

NEW QUESTION 175
You noticed a recent change to the BGP configuration on a PE router, the bgp scan time has been changed from the default value to 30s. Which three effects will this change have? (Choose three.)

A. The BGP table will be examined and verified more frequently
B. The BGP keepalive messages will be sent to the BGP peers at a faster rate
C. The BGP table will be modified more quickly in the event that a next-hop address becomes unreachable
D. The CPU load of the router will increase
E. The minimum time interval between sending EBGP and IBGP routing updates will decrease
F. The BGP convergence time will increase

Answer: ACD

NEW QUESTION 179
Refer to the exhibit.

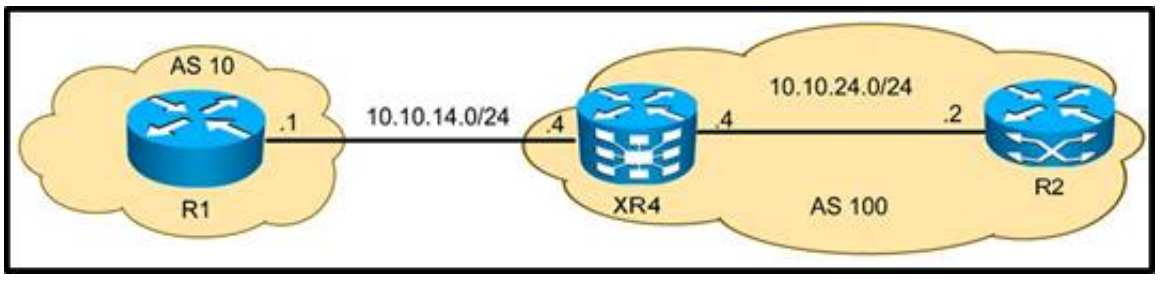

XR4 must protect itself from a DOS attack against its BGP process from R1 by using the TTL security feature. Which configuration achieves this goal?

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 182
.....

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