Cisco

Exam Questions 642-885

Deploying Cisco Service Provider Advanced Routing (SPADVOUTE)

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
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 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
Each router (RTA, RTB, and RTC) has one iBGP adjacency with the route reflector router RTD. Router RTC has an iBGP route advertised by RTA, but the same route is missing from RTB. Thenetwork engineer verifies that route filtering does not deny the route advertisement. Which action corrects the problem?

A. RTD(config-router)#neighbor 192.168.1.1 route-reflector-client RTD(config-router)#neighbor 192.168.1.1 description RTA RTD(config-router)#neighbor 192.168.1.2 route-reflector-client RTD(config-router)#neighbor 192.168.1.2 description RTB
B. RTC(config-router)#neighbor 192.168.1.4 route-reflector-client RTC(config-router)#neighbor 192.168.1.4 description RTD
C. RTA(config-router)#neighbor 192.168.1.4 route-reflector-client RTA(config-router)#neighbor 192.168.1.4 description RTDRTB(config-router)#neighbor 192.168.1.4 route-reflector-client RTB(config-router)#neighbor 192.168.1.4 description RTD
D. RTB(config-router)#neighbor 192.168.1.3 route-reflector-client RTB(config-router)#neighbor 192.168.1.3 description RTC
E. RTB(config-router)#neighbor 192.168.1.3 route-reflector-client RTB(config-router)#bgp cluster-id 192.168.1.2RTB(config-router)#no bgp client-to-client reflection

Answer: A

NEW QUESTION 5
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 6
Refer to the exhibit.

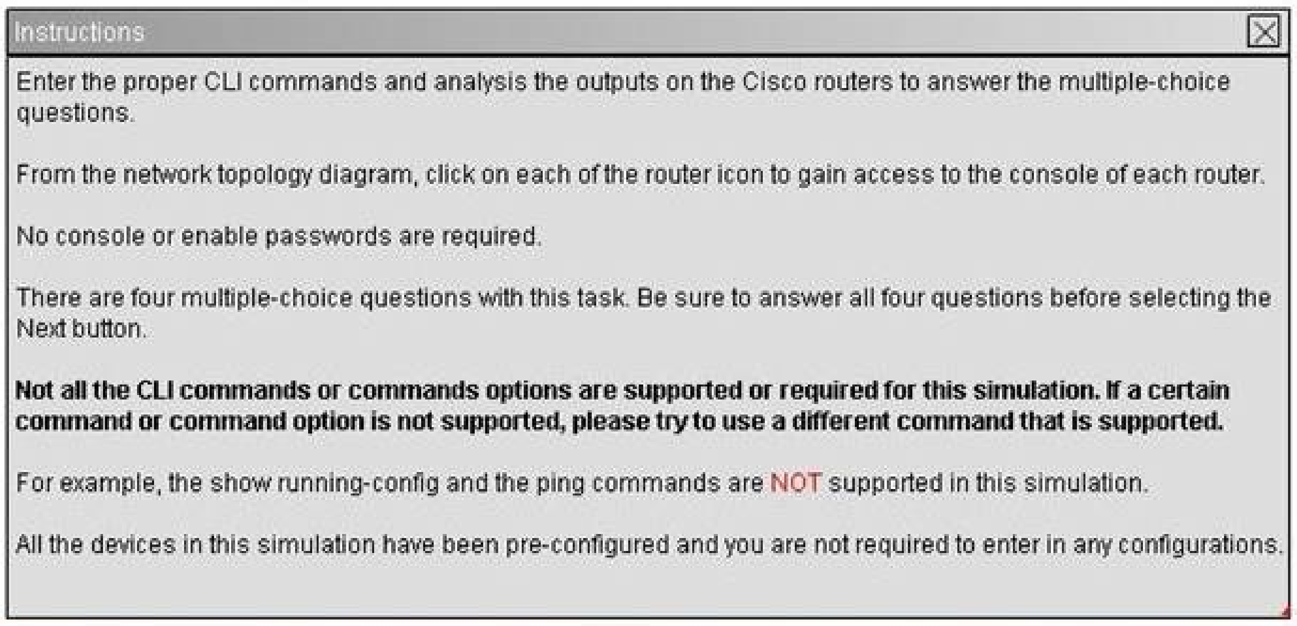

Instructions

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.

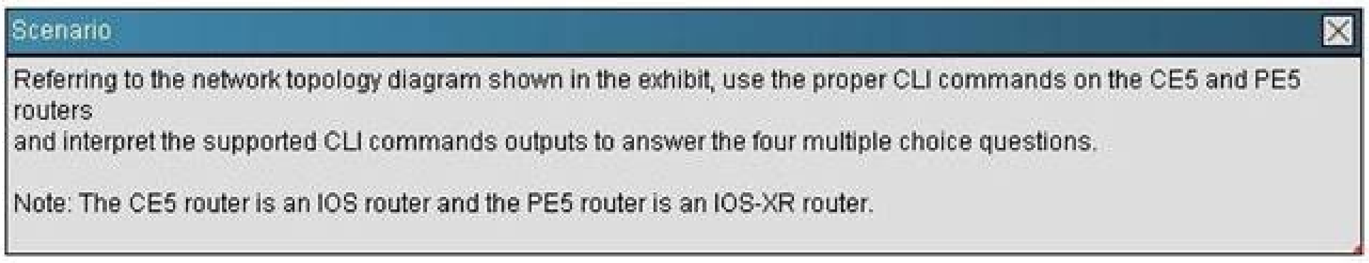

Scenario

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.

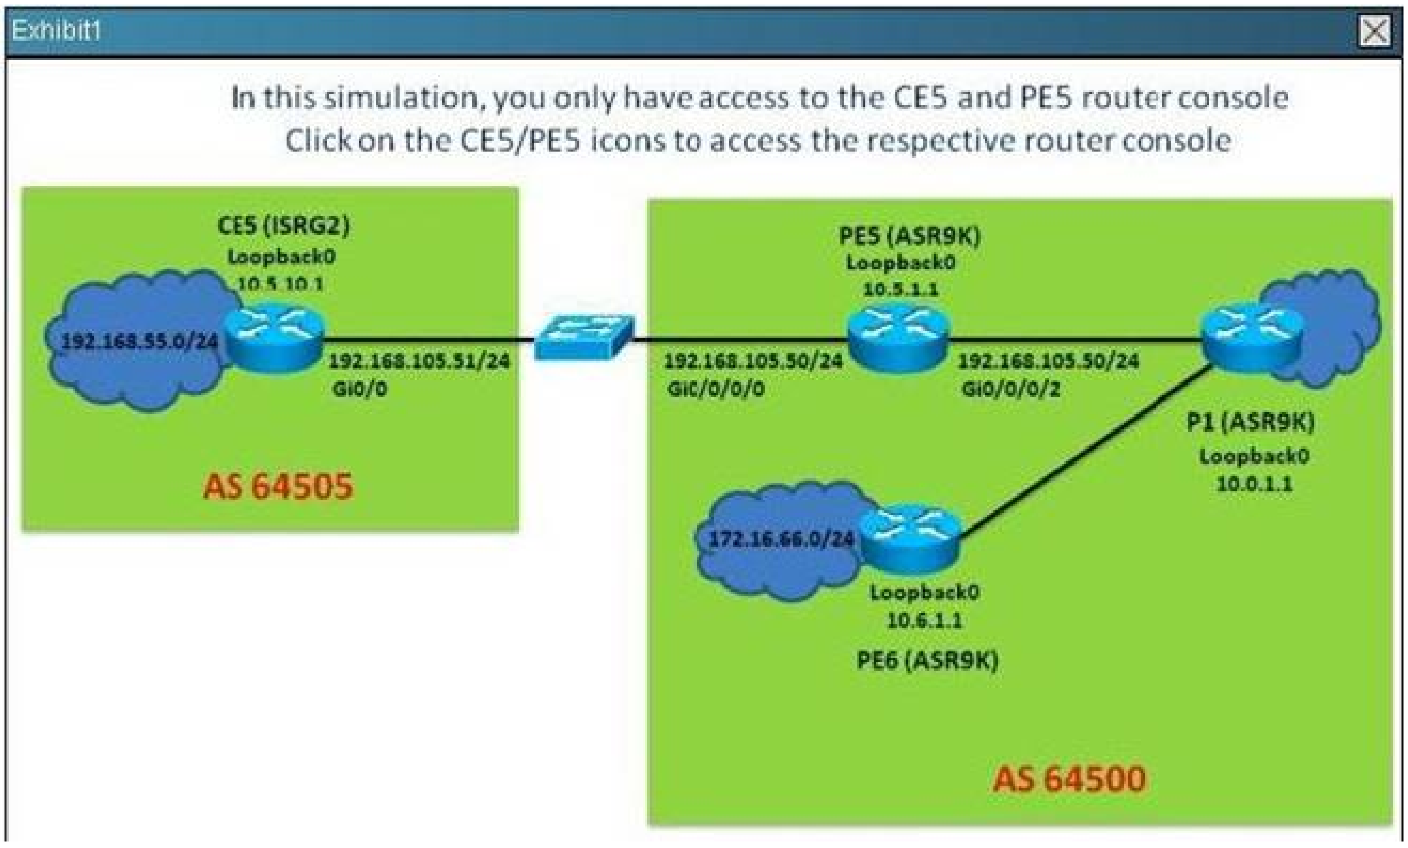

Exhibit1

In this simulation, you only have access to the CE5 and PE5 router console
Click on the CE5/PE5 icons to access the respective router console

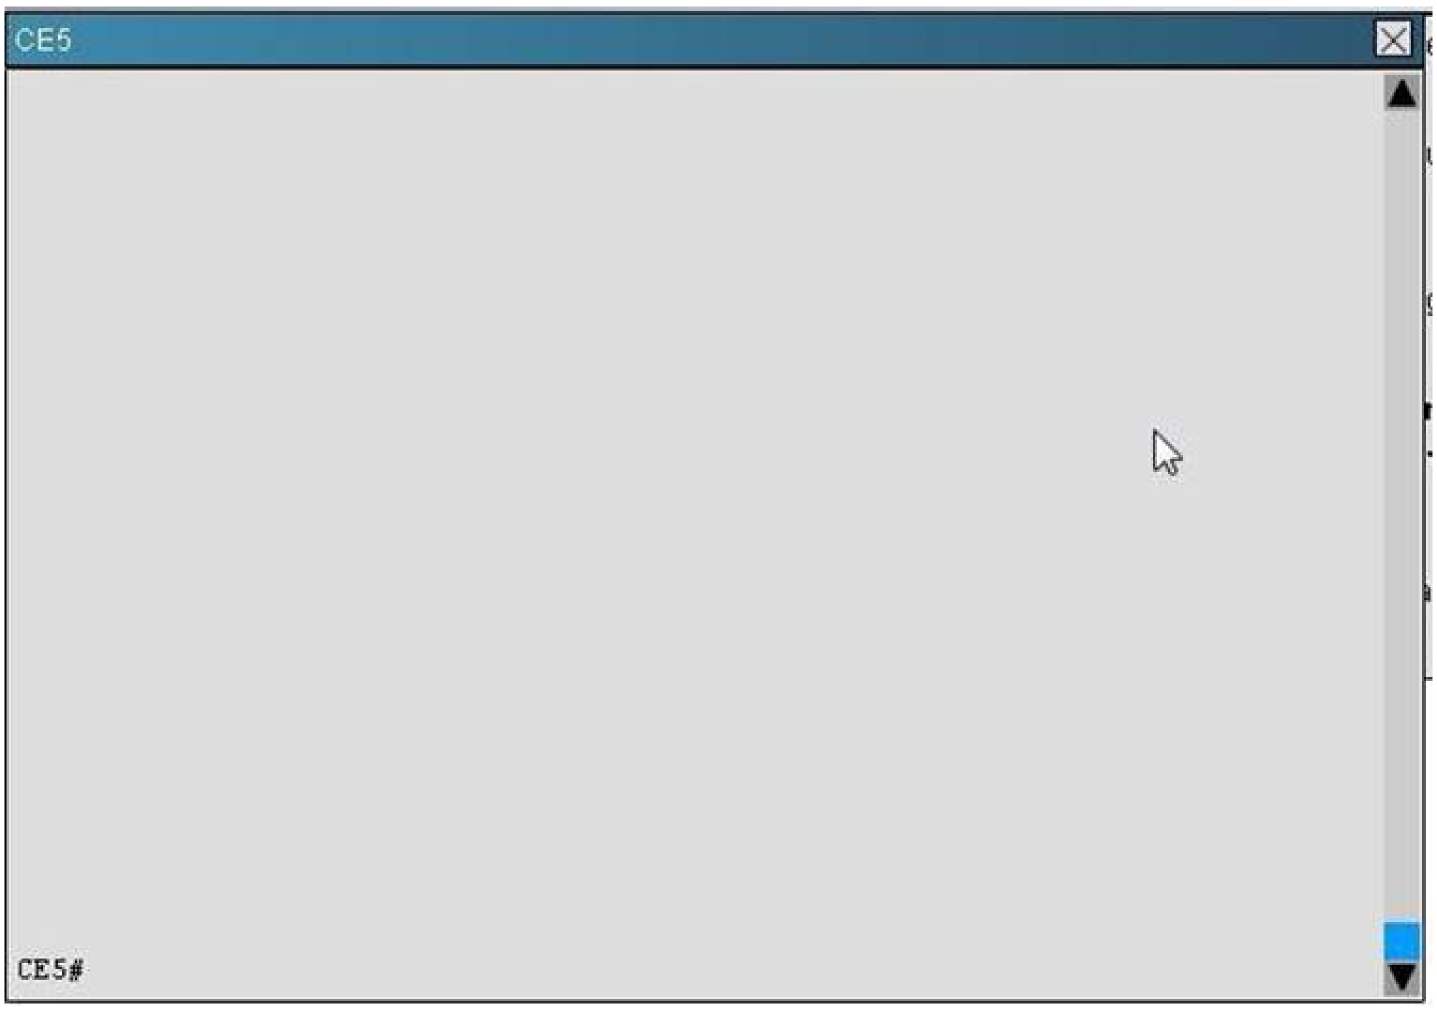

CE5

```
CE5#
```

```
PE5#
```

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
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 8
After configuring the tunnel interface as shown in the exhibit, no IPv6 traffic is passed over the IPv4 network.

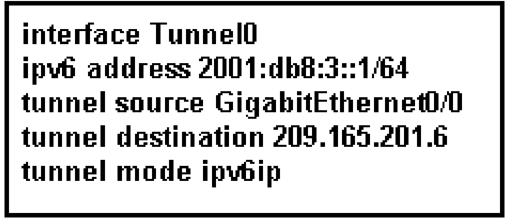

```
interface Tunnel0
ipv6 address 2001:db8:3::1/64
tunnel source GigabitEthernet0/0
tunnel destination 209.165.201.6
tunnel mode ipv6ip
```

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 9
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 10
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 10
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 15
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 18
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 19
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 21
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 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 26
Which of the following can be used by dual-stack service providers supporting IPv4/IPv6
customers with dual-stack hosts using public IPv6 addresses and private IPv4 addresses?

A. NAT64
B. 6RD
C. 6to4 tunnels
D. Carrier-grade NAT

Answer: D

Explanation: Carrier Grade NAT is a large-scale NAT, capable of providing private-IPv4-to-public-IPv4 translation in the order of millions of translations. Carrier Grade NAT can support several hundred thousand subscribers with the bandwidth throughput of at least 10Gb/s full-duplex. With IPv4 addresses reaching depletion, Carrier Grade NAT is vital in providing private IPv4 connectivity to the public IPv4 internet. In addition, Carrier Grade NAT is not limited to IPv4 NAT; it can also translate between IPv4 and IPv6 addresses.

NEW QUESTION 30
An engineer is enabling multicast routing across an entire core infrastructure. Which two commands enable multicast routing on Cisco IOS XE instances? (Choose two.)

A. ip multicast-routing
B. ip multicast-routing vrf global
C. interface type slot/path_id ip pim sparse-mode
D. interface type slot/path_id ip cgmp
E. interface type slot/path_id ip pim dense-mode
F. ip mroute-cache

Answer: AC

NEW QUESTION 31
Refer to the Cisco IOS-XR configuration exhibit.

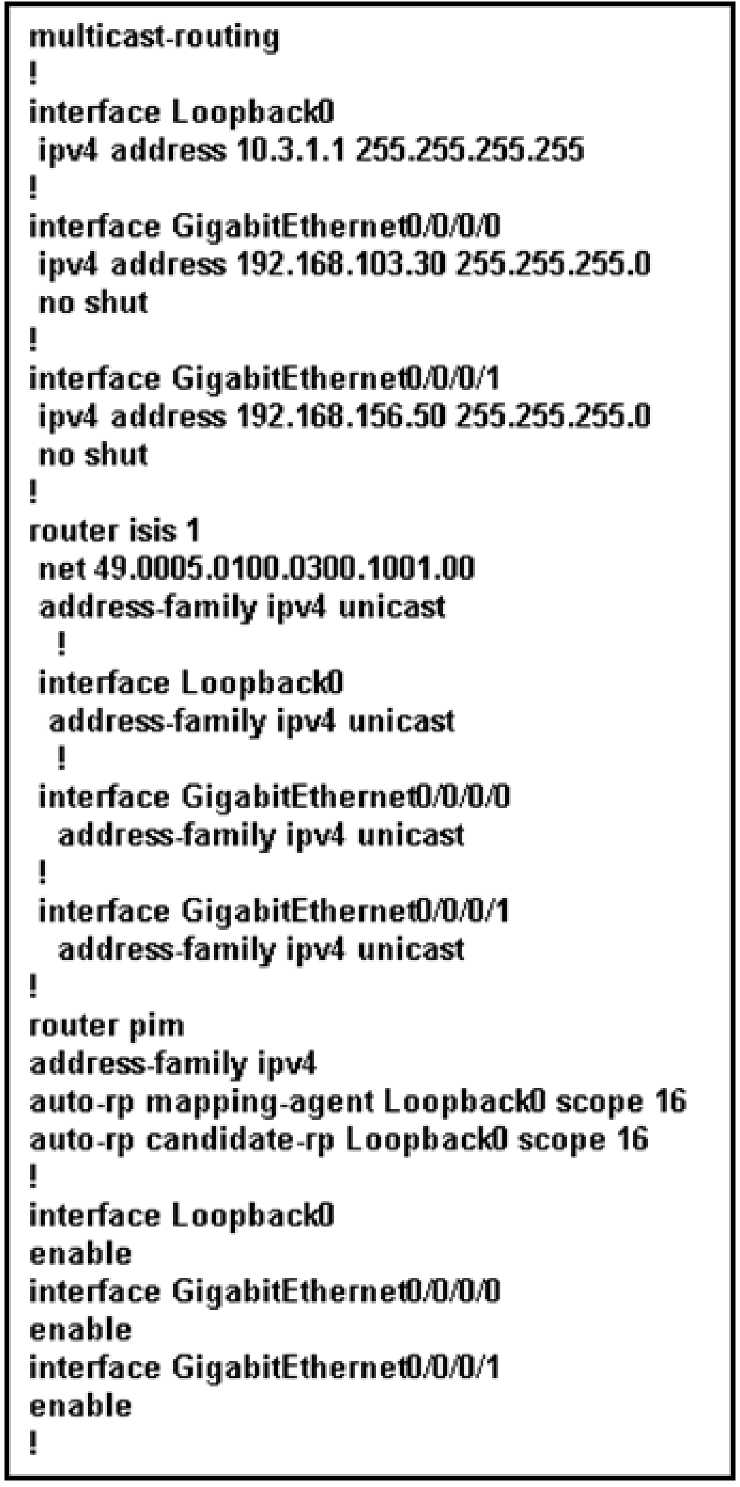

```
multicast-routing
!
interface Loopback0
 ipv4 address 10.3.1.1 255.255.255.255
!
interface GigabitEthernet0/0/0/0
 ipv4 address 192.168.103.30 255.255.255.0
 no shut
!
interface GigabitEthernet0/0/0/1
 ipv4 address 192.168.156.50 255.255.255.0
 no shut
!
router isis 1
 net 49.0005.0100.0300.1001.00
 address-family ipv4 unicast
   !
 interface Loopback0
  address-family ipv4 unicast
   !
 interface GigabitEthernet0/0/0/0
   address-family ipv4 unicast
 !
 interface GigabitEthernet0/0/0/1
   address-family ipv4 unicast
!
router pim
address-family ipv4
auto-rp mapping-agent Loopback0 scope 16
auto-rp candidate-rp Loopback0 scope 16
!
interface Loopback0
enable
interface GigabitEthernet0/0/0/0
enable
interface GigabitEthernet0/0/0/1
enable
!
```

The Cisco IOS-XR router is unable to establish any PIM neighbor relationships. What is wrong with the configuration?

A. The configuration is missing:interface gi0/0/0/0 ip pim sparse-mode interface gi0/0/0/1 ip pim sparse-mode interface loopback0 ip pim sparse-mode
B. The configuration is missing: multicast-routingaddress-family ipv4 interface gi0/0/0/0 enableinterface gi0/0/0/1 enable
C. The auto-rp scoping configurations should be set to 1 not 16
D. The RP address has not been configured using the rp-address router PIM configuration command
E. PIM defaults to dense mode operations only, so PIM sparse mode must be enabled using the pim sparse-mode router PIM configuration command

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 40
Refer to the Cisco IOS-XR show output exhibit.

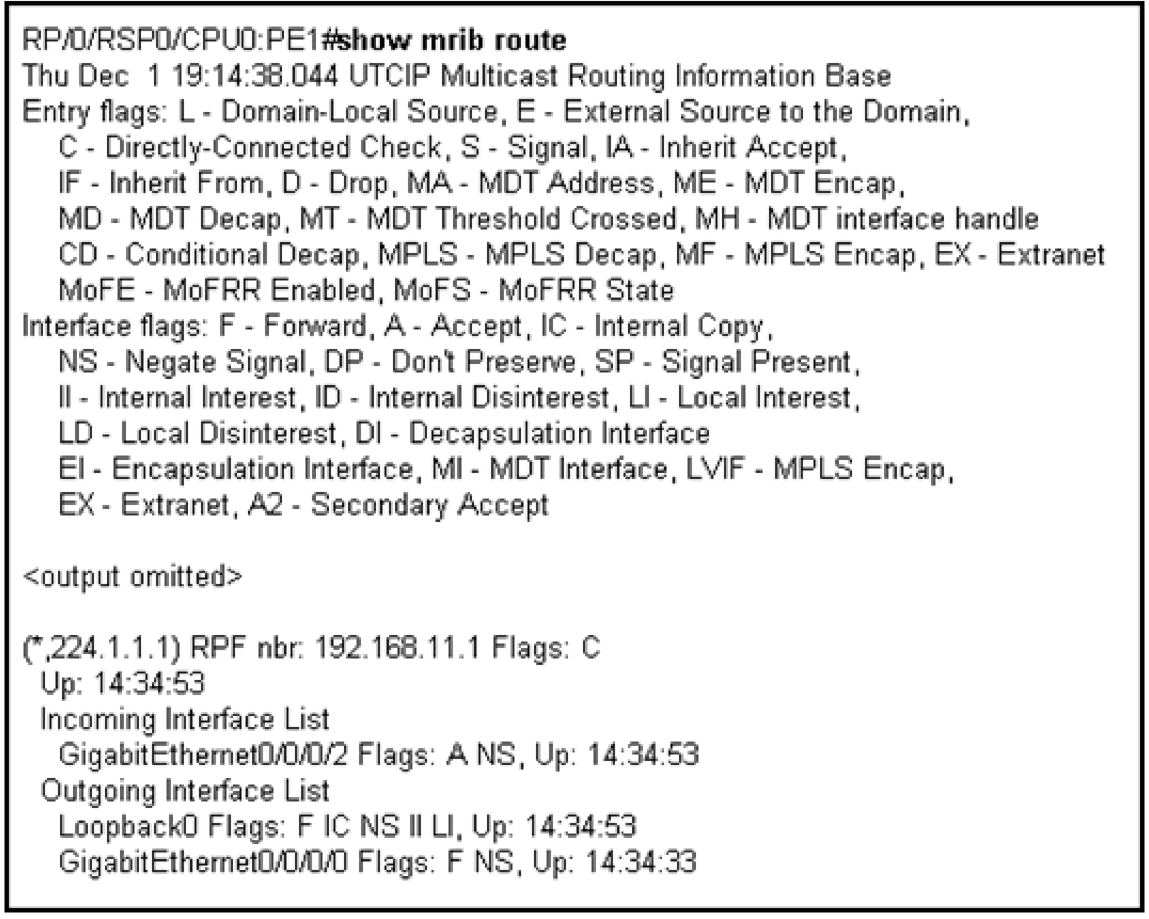

```
RP/0/RSP0/CPU0:PE1#show mrib route
Thu Dec  1 19:14:38.044 UTCIP Multicast Routing Information Base
Entry flags: L - Domain-Local Source, E - External Source to the Domain,
    C - Directly-Connected Check, S - Signal, IA - Inherit Accept,
    IF - Inherit From, D - Drop, MA - MDT Address, ME - MDT Encap,
    MD - MDT Decap, MT - MDT Threshold Crossed, MH - MDT interface handle
    CD - Conditional Decap, MPLS - MPLS Decap, MF - MPLS Encap, EX - Extranet
    MoFE - MoFRR Enabled, MoFS - MoFRR State
Interface flags: F - Forward, A - Accept, IC - Internal Copy,
    NS - Negate Signal, DP - Don't Preserve, SP - Signal Present,
    II - Internal Interest, ID - Internal Disinterest, LI - Local Interest,
    LD - Local Disinterest, DI - Decapsulation Interface
    EI - Encapsulation Interface, MI - MDT Interface, LVIF - MPLS Encap,
    EX - Extranet, A2 - Secondary Accept

<output omitted>

(*,224.1.1.1) RPF nbr: 192.168.11.1 Flags: C
  Up: 14:34:53
  Incoming Interface List
    GigabitEthernet0/0/0/2 Flags: A NS, Up: 14:34:53
  Outgoing Interface List
    Loopback0 Flags: F IC NS II LI, Up: 14:34:53
    GigabitEthernet0/0/0/0 Flags: F NS, Up: 14:34:33
```

Which two statements are correct? (Choose two.)

A. The RPF neighbor 192.168.11.1 is the path towards the RP for the 224.1.1.1 multicast group
B. The RP for the 224.1.1.1 multicast group is reachable over the Gi0/0/0/0 interface
C. This router is the RP for the 224.1.1.1 multicast group
D. Incoming 224.1.1.1 multicast group traffic will be sent out through the Gi0/0/0/0 interface
E. Incoming 224.1.1.1 multicast group traffic will be sent out through the Gi0/0/0/2 interface

Answer: AD

NEW QUESTION 44
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 48
DRAG DROP

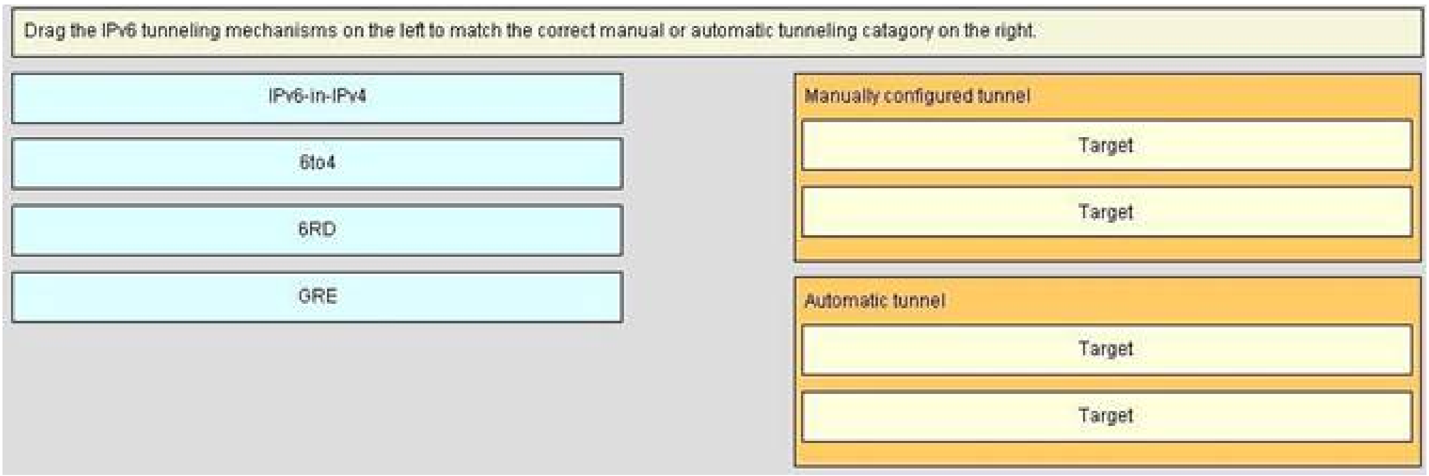

Drag the IPv6 tunneling mechanisms on the left to match the correct manual or automatic tunneling catagory on the right.

Answer:

Explanation: IPv6-in-IPv4 and GRE are manual and 6RDand 6to4

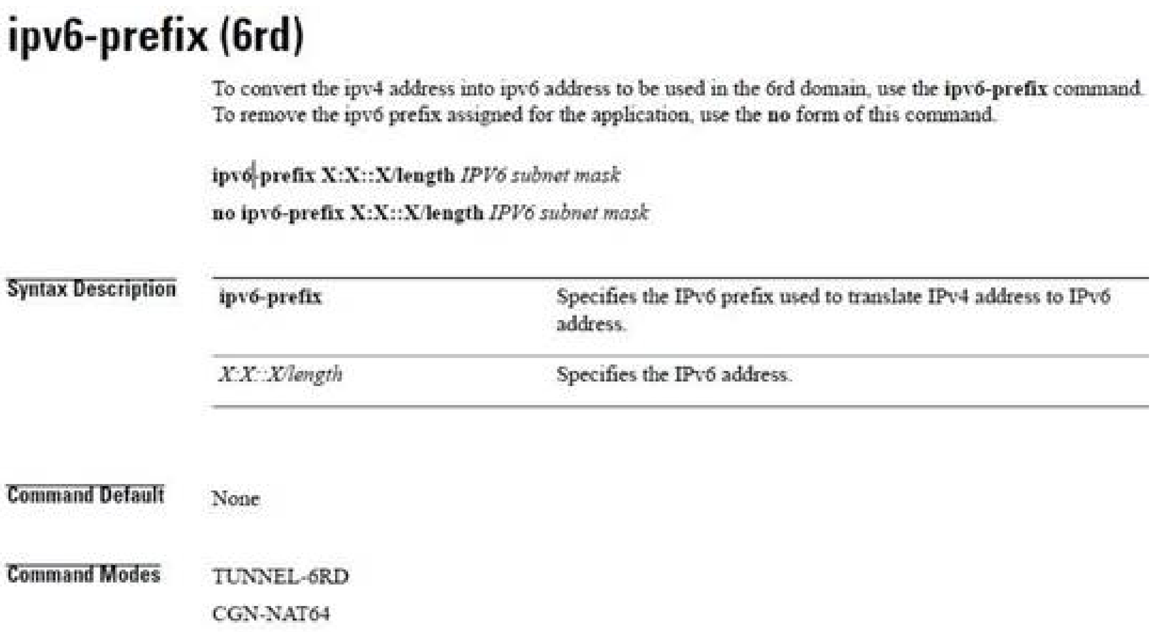

ipv6-prefix (6rd)

To convert the ipv4 address into ipv6 address to be used in the 6rd domain, use the **ipv6-prefix** command. To remove the ipv6 prefix assigned for the application, use the **no** form of this command.

ipv6-prefix X:X::X/length *IPV6 subnet mask*

no ipv6-prefix X:X::X/length *IPV6 subnet mask*

Syntax Description		
	ipv6-prefix	Specifies the IPv6 prefix used to translate IPv4 address to IPv6 address.
	X:X::X/length	Specifies the IPv6 address.

Command Default None

Command Modes TUNNEL-6RD
CGN-NAT64

Download this chapter Implementing Tunnels Download the complete book
Interface and Hardware Component Configuration Guide, Cisco IOS XE Release 3S (PDF - 1 MB) Feedback
Contents Implementing Tunnels
Finding Feature Information Restrictions for Implementing Tunnels
Information About Implementing Tunnels Tunneling Versus Encapsulation
Tunnel ToS
Generic Routing Encapsulation
GRE Tunnel IP Source and Destination VRF Membership GRE IPv4 Tunnel Support for IPv6 Traffic
EoMPLS over GRE
Provider Edge to Provider Edge Generic Routing EncapsulationTunnels Provider to Provider Generic Routing Encapsulation Tunnels
Provider Edge to Provider Generic Routing Encapsulation Tunnels Features Specific to Generic Routing Encapsulation
Features Specific to Ethernet over MPLS
Features Specific to Multiprotocol Label Switching Virtual Private Network Overlay Tunnels for IPv6
IPv6 Manually Configured Tunnels Automatic 6to4 Tunnels
ISATAP Tunnels Path MTU Discovery
QoS Options for Tunnels How to Implement Tunnels Determining the Tunnel Type
Configuring an IPv4 GRE Tunnel GRE Tunnel Keepalive
What to Do Next
Configuring GRE on IPv6 Tunnels What to Do Next
Configuring GRE Tunnel IP Source and Destination VRF Membership What to Do Next
Manually Configuring IPv6 Tunnels What to Do Next
Configuring 6to4 Tunnels What to Do Next
Configuring ISATAP Tunnels
Verifying Tunnel Configuration and Operation Configuration Examples for Implementing Tunnels Example: Configuring a GRE IPv4 Tunnel Example: Configuring GRE on IPv6 Tunnels
Example: Configuring GRE Tunnel IP Source and Destination VRF Membership Example: Configuring EoMPLS over GRE
Example: Manually Configuring IPv6 Tunnels Example: Configuring 6to4 Tunnels Example: Configuring ISATAP Tunnels
Configuring QoS Options on Tunnel Interfaces Examples Policing Example
Additional References
Feature Information for Implementing Tunnels Implementing Tunnels
Last Updated: September 17, 2012

This module describes the various types of tunneling techniques. Configuration details and examples are provided for the tunnel types that use physical or virtual interfaces. Many tunneling techniques are implemented using technology-specific commands, and links are provided to the appropriate technology modules.

Tunneling provides a way to encapsulate arbitrary packets inside a transport protocol.
Tunnels are
implemented as virtual interfaces to provide a simple interface for configuration purposes.
The tunnel interface
is not tied to specific "passenger" or "transport" protocols, but rather is an architecture to provide the services necessary to implement any standard point-to-point encapsulation scheme. Note
Cisco ASR 1000 Series Aggregation Services Routers support VPN routing and forwarding (VRF)-aware generic routing encapsulation (GRE) tunnel keepalive features. Finding Feature Information
Restrictions for Implementing Tunnels Information About Implementing Tunnels How to Implement Tunnels
Configuration Examples for Implementing Tunnels Additional References
Feature Information for Implementing Tunnels Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see Bug Search Tool and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table at the end of this module.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Restrictions for Implementing Tunnels

It is important to allow the tunnel protocol to pass through a firewall and access control list (ACL) check.

Multiple point-to-point tunnels can saturate the physical link with routing information if the bandwidth is not configured correctly on a tunnel interface.

A tunnel looks like a single hop link, and routing protocols may prefer a tunnel over a multihop physical path.

The tunnel, despite looking like a single hop link, may traverse a slower path than a multihop link. A tunnel is as robust and fast, or as unreliable and slow, as the links that it actually traverses. Routing protocols that make their decisions based only on hop counts will often prefer a tunnel over a set of physical links. A tunnel might appear to be a one-hop, point-to-point link and have the lowest-cost path, but the tunnel may actually cost more in terms of latency when compared to an alternative physical topology. For example, in the topology shown in the figure below, packets from Host 1 will appear to travel across networks w, t, and z to get to Host 2 instead of taking the path w, x, y, and z because the tunnel hop count appears shorter. In fact, the packets going through the tunnel will still be traveling across Router A, B, and C, but they must also travel to Router D before coming back to Router C. Figure 1
Tunnel Precautions: Hop Counts

A tunnel may have a recursive routing problem if routing is not configured accurately. The best path to a tunnel destination is via the tunnel itself; therefore

recursive routing causes the tunnel interface to flap. To avoid recursive routing problems, keep the control-plane routing separate from the tunnel routing by using the following methods:
Use a different autonomous system number or tag. Use a different routing protocol.
Ensure that static routes are used to override the first hop (watch for routing loops). The following error is displayed when there is recursive routing to a tunnel destination:
%TUN-RECURDOWN Interface Tunnel 0 temporarily disabled due to recursive routing Information About Implementing Tunnels Tunneling Versus Encapsulation Tunnel ToS
Generic Routing Encapsulation EoMPLS over GRE
Overlay Tunnels for IPv6
IPv6 Manually Configured Tunnels Automatic 6to4 Tunnels
ISATAP Tunnels Path MTU Discovery
QoS Options for Tunnels
Tunneling Versus Encapsulation
To understand how tunnels work, you must be able to distinguish between concepts of encapsulation and tunneling. Encapsulation is the process of adding headers to data at each layer of a particular protocol stack.
The Open Systems Interconnection (OSI) reference model describes the functions of a network. To send a data packet from one host (for example, a PC) to another on a network, encapsulation is used to add a header in front of the data packet at each layer of the protocol stack in descending order. The header must contain a data field that indicates the type of data encapsulated at the layer immediately above the current layer. As the packet ascends the protocol stack on the receiving side of the network, each encapsulation header is removed in reverse order.
Tunneling encapsulates data packets from one protocol within a different protocol and transports the packets on a foreign network. Unlike encapsulation, tunneling allows a lower-layer protocol and a same-layer protocol to be carried through the tunnel. A tunnel interface is a virtual (or logical) interface. Tunneling consists of three main components: Passenger protocol--The protocol that you are encapsulating. For example, IPv4 and IPv6 protocols. Carrier protocol--The protocol that encapsulates. For example, generic routing encapsulation (GRE) and Multiprotocol Label Switching (MPLS).
Transport protocol--The protocol that carries the encapsulated protocol. The main transport protocol is IP.
The figure below illustrates IP tunneling terminology and concepts: Figure 2
IP Tunneling Terminology and Concepts Tunnel ToS
Tunnel type of service (ToS) allows you to tunnel network traffic and group all packets in the same ToS byte value. The ToS byte values and Time-to-Live (TTL) hop-count value can be set in the encapsulating IP header of tunnel packets for an IP tunnel interface on a router. Tunnel ToS feature is supported for Cisco Express Forwarding (formerly known as CEF), fast switching, and process switching.
The ToS and TTL byte values are defined in RFC 791. RFC 2474, and RFC 2780 obsolete the use of the ToS byte as defined in RFC 791. RFC 791 specifies that bits 6 and 7 of the ToS byte (the first two least significant bits) are reserved for future use and should be set to
1. For Cisco IOS XE Release 2.1, the Tunnel ToS feature does not conform to this standard and allows you touse the whole ToS byte value, including bits 6 and 7, and to decide to which RFC standard the ToS byte of your packets should conform.
Generic Routing Encapsulation
GRE is defined in RFC 2784. GRE is a carrier protocol that can be used with many different underlying transport protocols and can carry many passenger protocols. RFC
2784 also covers the use of GRE with IPv4 as the transport protocol and the passenger protocol. Cisco software supports GRE as the carrier protocol with many combinations of passenger and transport protocols.
GRE tunnels are described in the following sections: GRE Tunnel IP Source and Destination VRF Membership GRE IPv4 Tunnel Support for IPv6 Traffic
GRE Tunnel IP Source and Destination VRF Membership
The GRE Tunnel IP Source and Destination VRF Membership feature allows you to configure the source and destination of a tunnel to belong to any VPN routing and forwarding (VRFs) tables. A VRF table stores routing data for each VPN. The VRF table defines the VPN membership of a customer site that is attached to the network access server (NAS). Each VRF table comprises an IP routing table, a derived Cisco Express Forwarding table, and guidelines and routing protocol parameters that control the information that is included in the routing table.
Prior to Cisco IOS XE Release 2.2, GRE IP tunnels required the IP tunnel destination to be in the global routing table. The implementation of this feature allows you to configure a tunnel source and destination to belong to any VRF. As with existing GRE tunnels, the tunnel becomes disabled if no route to the tunnel destination is defined.
GRE IPv4 Tunnel Support for IPv6 Traffic
IPv6 traffic can be carried over IPv4 GRE tunnels by using the standard GRE tunneling technique to provide the services necessary to implement a standard point-to-point encapsulation scheme. GRE tunnels are links between two points, with a separate tunnel for each point. GRE tunnels are not tied to a specific passenger or transport protocol, but in case of IPv6 traffic, IPv6 is the passenger protocol, GRE is the carrier protocol, and IPv4 is the transport protocol.
The primary use of GRE tunnels is to provide a stable connection and secure communication between two edge devices or between an edge device and an end system. The edge device and the end system must have a dual-stack implementation.
GRE has a protocol field that identifies the passenger protocol. GRE tunnels allow intermediate system to intermediate system (IS-IS) or IPv6 to be specified as the passenger protocol, therebyallowing both IS-IS and IPv6 traffic to run over the same tunnel. If GRE does not have a protocol field, it becomes impossible to distinguish whether the tunnel is carrying IS-IS or IPv6 packets.
EoMPLS over GRE
Ethernet over MPLS (EoMPLS) is a tunneling mechanism that allows you to tunnel Layer 2 traffic through a Layer 3 MPLS network. EoMPLS is also known as Layer 2 tunneling.
EoMPLS effectively facilitates Layer 2 extension over long distances. EoMPLS over GRE helps you to create the GRE tunnel as hardware-based switched, and encapsulates EoMPLS frames within the GRE tunnel. The GRE connection is established between the two core routers, and then the MPLS label switched path (LSP) is tunneled over.
GRE encapsulation is used to define a packet that has header information added to it prior to being forwarded.
De-encapsulation is the process of removing the additional header information when the packet reaches the destination tunnel endpoint.
When a packet is forwarded through a GRE tunnel, two new headers are added to the front of the packet and hence the context of the new payload changes. After encapsulation, what was originally the data payload and separate IP header are now known as the GRE payload. A GRE header is added to the packet to provide information on the protocol type and the recalculated checksum. A new IP header is also added to the front of the GRE header. This IP header contains the destination IP address of the tunnel. The GRE header is added to packets such as IP, Layer 2 VPN, and Layer 3 VPN before the header enters into the tunnel. All routers along the path that receives the encapsulated packet use the new IP header to determine how the packet can reach the tunnel endpoint.
In IP forwarding, on reaching the tunnel destination endpoint, the new IP header and the GRE header are removed from the packet and the original IP header is used to forward the packet to the final destination.
The EoMPLS over GRE feature removes the new IP header and GRE header from the packet at the tunnel destination, and the MPLS label is used to forward the packet to the appropriate Layer 2 attachment circuit or Layer 3 VRF.
The scenarios in the following sections describe the L2VPN and L3VPN over GRE deployment on provider edge (PE) or provider (P) routers:
Provider Edge to Provider Edge Generic Routing EncapsulationTunnels Provider to Provider Generic Routing Encapsulation Tunnels
Provider Edge to Provider Generic Routing Encapsulation Tunnels Features Specific to Generic Routing Encapsulation
Features Specific to Ethernet over MPLS
Features Specific to Multiprotocol Label Switching Virtual Private Network Provider Edge to Provider Edge Generic Routing EncapsulationTunnels
In the Provider Edge to Provider Edge (PE) GRE tunnels scenario, a customer does not transition any part of the core to MPLS but prefers to offer EoMPLS and basic MPLS VPN services. Therefore, GRE tunneling of MPLS traffic is done between PEs.
Provider to Provider Generic Routing Encapsulation Tunnels
In the Provider to Provider (P) GRE tunnels scenario, Multiprotocol Label Switching (MPLS) is enabled between Provider Edge (PE) and P routers but the network core can either have non-MPLS aware routers or IP encryption boxes. In this scenario, GRE tunneling of the MPLS labeled packets is done between P routers.

Provider Edge to Provider Generic Routing Encapsulation Tunnels in a Provider Edge to Provider GRE tunnels scenario, a network has MPLS-aware P to P nodes. GRE tunneling is done between a PE to P non-MPLS network segment. Features Specific to Generic Routing Encapsulation You should understand the following configurations and information for a deployment scenario:
Tunnel endpoints can be loopbacks or physical interfaces.
Configurable tunnel keepalive timer parameters per endpoint and a syslog message must be generated when the keepalive timer expires.
Bidirectional forwarding detection (BFD) is supported for tunnel failures and for the Interior Gateway Protocol (IGP) that use tunnels.
IGP load sharing across a GRE tunnel is supported. IGP redundancy across a GRE tunnel is supported. Fragmentation across a GRE tunnel is supported. Ability to pass jumbo frames is supported.
All IGP control plane traffic is supported.
IP ToS preservation across tunnels is supported.
A tunnel should be independent of the endpoint physical interface type; for example, ATM, Gigabit, Packet over SONET (POS), and TenGigabit.
Up to 100 GRE tunnels are supported. Features Specific to Ethernet over MPLS
Any Transport over MPLS (AToM) sequencing. IGP load sharing and redundancy.
Port mode Ethernet over MPLS (EoMPLS). Pseudowire redundancy.
Support for up to to 200 EoMPLS virtual circuits (VCs).
Tunnel selection and the ability to map a specific pseudowire to a GRE tunnel. VLAN mode EoMPLS.
Features Specific to Multiprotocol Label Switching Virtual Private Network Support for the PE role with IPv4 VRF.
Support for all PE to customer edge (CE) protocols.
Load sharing through multiple tunnels and also equal cost IGP paths with a single tunnel. Support for redundancy through unequal cost IGP paths with a single tunnel.
Support for the IP precedence value being copied onto the expression (EXP) bits field of the Multiprotocol Label Switching (MPLS) label and then onto the precedence bits on the outer IPv4 ToS field of the generic routing encapsulation (GRE) packet.
See the section, "Example: Configuring EoMPLS over GRE" for a sample configuration sequence of EoMPLS over GRE. For more details on EoMPLS over GRE, see the Deploying and Configuring MPLS Virtual Private Networks
In IP Tunnel Environments document. Overlay Tunnels for IPv6
The figure below illustrates how overlay tunneling encapsulates IPv6 packets in IPv4 packets for delivery across an IPv4 infrastructure (a core network or the Internet). By using overlay tunnels, you can communicate with isolated IPv6 networks without upgrading the IPv4 infrastructure between them. Overlay tunnels can be configured between border routers or between a border router and a host; however, both tunnel endpoints must support, IPv4 and IPv6 protocol stacks. IPv6 supports the following types of overlay tunneling mechanisms:
6to4 GRE
Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) IPv4-compatible
Manual Figure 3
Overlay Tunnels Note
If the basic IPv4 packet header does not have optional fields, overlay tunnels can reduce the maximum transmission unit (MTU) of an interface by 20 octets. A network that uses overlay tunnels is difficult to troubleshoot. Therefore, overlay tunnels that connect isolated IPv6 networks should not be considered as the final IPv6 network architecture. The use of overlay tunnels is considered as a transition technique for a network that supports either both IPv4 and IPv6 protocol stacks or just the IPv6 protocol stack.
Consult the table below to determine which type of tunnel you want to configure to carry IPv6 packets over an IPv4 network.
Table 1
Suggested Usage of Tunnel Types to Carry IPv6 Packets over an IPv4 Network Tunneling Type
Suggested Usage Usage Notes 6to4
Point-to-multipoint tunnels that can be used to connect isolated IPv6 sites. Sites use addresses that begin with the 2002::/16 prefix.
GRE/IPv4
Simple point-to-point tunnels that can be used within a site or between sites.
Tunnels can carry IPv6, Connectionless Network ServiceCLNS, and many other types of packets.
ISATAP
Point-to-multipoint tunnels that can be used to connect systems within a site. Sites can use any IPv6 unicast addresses.
Manual
Simple point-to-point tunnels that can be used within a site or between sites. Tunnels can carry IPv6 packets only.
Individual tunnel types are discussed in detail in the following concepts, and we recommend that you review and understand the information on the specific tunnel type that you want to implement. Consult the table below for a summary of the tunnel configuration parameters that you may find useful.
Table 2
Overlay Tunnel Configuration Parameters by Tunneling Type Overlay Tunneling Type
Overlay Tunnel Configuration Parameter Tunnel Mode
Tunnel Source Tunnel Destination
Interface Prefix/Address 6to4
ipv6ip 6to4
An IPv4 address or a reference to an interface on which IPv4 is configured.
Not required. These are all point-to-multipoint tunneling types. The IPv4 destination address is calculated, on a per-packet basis, from the IPv6 destination.
An IPv6 address. The prefix must embed the tunnel source IPv4 address.
GRE/IPv4
gre ip
An IPv4 address. An IPv6 address. ISATAP
ipv6ip isatap
Not required. These are all point-to-multipoint tunneling types. The IPv4 destination address is calculated on a per-packet basis from the IPv6 destination.
An IPv6 prefix in modified eui-64 format. The IPv6 address is generated from the prefix and the tunnel source IPv4 address.
Manual ipv6ip
An IPv4 address. An IPv6 address.
IPv6 Manually Configured Tunnels
A manually configured tunnel is equivalent to a permanent link between two IPv6 domains over an IPv4 backbone. The primary use of a manually configured tunnel is to stabilize connections that require secure communication between two edge routers, or between an end system and an edge router. The manual configuration tunnel also stabilizes connection between remote IPv6 networks.
An IPv6 address is manually configured on a tunnel interface. Manually configured IPv4 addresses are assigned to the tunnel source and destination. The host or router at each end of a configured tunnel must support both the IPv4 and IPv6 protocol stacks. Manually configured tunnels can be configured between border routers or between a border router and a host. Cisco Express Forwarding switching can be used for manually configured IPv6 tunnels. Switching can be disabled if process switching is required.
Automatic 6to4 Tunnels
An automatic 6to4 tunnel allows isolated IPv6 domains to be connected over an IPv4 network to remote IPv6 networks. The key difference between automatic 6to4 tunnels and manuallyconfigured tunnels is that the tunnel is not point-to-point; it is point-to-multipoint. In automatic 6to4 tunnels, routers are not configured in pairs because they treat the IPv4 infrastructure as a virtual nonbroadcast multiaccess (NBMA) links. The IPv4 address embedded in the IPv6 address is used to find the other end of the automatic tunnel.
An automatic 6to4 tunnel may be configured on a border router in an isolated IPv6 network, which creates a tunnel on a per-packet basis on a border router in another IPv6 network over an IPv4 infrastructure. The tunnel destination is determined by the IPv4 address of the border router extracted from the IPv6 address

that starts with the prefix 2002::/16, where the format is 2002:border-router-IPv4-address ::/48.The embedded IPv4 addresses are 16 bits and can be used to number networks within the site. The border router at each end of a 6to4 tunnel must support both IPv4 and IPv6 protocol stacks. 6to4 tunnels are configured between border routers or between a border router and a host.
The simplest deployment scenario for 6to4 tunnels is to interconnect multiple IPv6 sites, each of which has at least one connection to a shared IPv4 network. This IPv4 network could either be the Internet or a corporate backbone. The key requirement is that each site have a globally unique IPv4 address; the Cisco software uses this address to construct a globally unique 6to4/48 IPv6 prefix. A tunnel with appropriate entries in a Domain Name System (DNS) that maps hostnames and IP addresses for both IPv4 and IPv6 domains, allows the applications to choose the required address IPv6 traffic can be carried over IPv4 GRE tunnels by using the standard GRE tunneling technique to provide the services necessary to implement a standard point-to-point encapsulation scheme. GRE tunnels are links between two points, with a separate tunnel for each point. GRE tunnels are not tied to a specific passenger or transport protocol, but in case of IPv6 traffic, IPv6 is the passenger protocol, GRE is the carrier protocol, and IPv4 is the transport protocol.
The primary use of GRE tunnels is to provide a stable connection and secure communication between two edge devices or between an edge device and an end system. The edge device and the end system must have a dual-stack implementation. GRE has a protocol field that identifies the passenger protocol. GRE tunnels allow intermediate system to intermediate system (IS-IS) or IPv6 to be specified as the passenger protocol, thereby allowing both IS-IS and IPv6 traffic to run over the same tunnel. If GRE does not have a protocol field, it becomes impossible to distinguish whether the tunnel is carrying IS-IS or IPv6 packets.

NEW QUESTION 52
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 54
Refer to the exhibit.

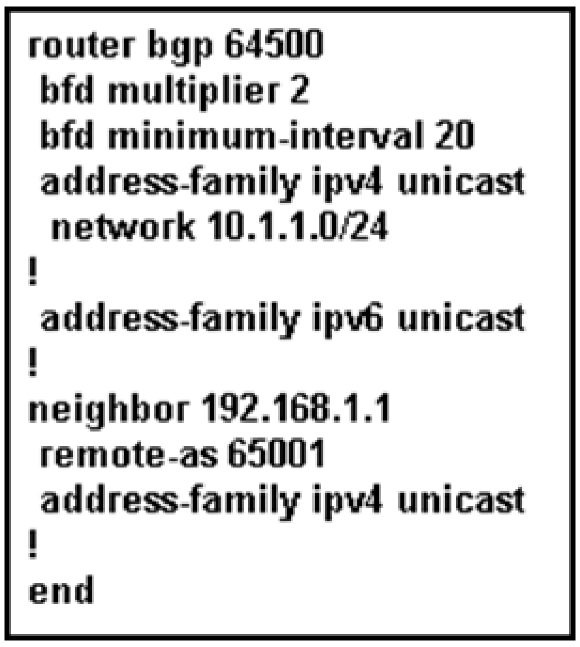

```
router bgp 64500
 bfd multiplier 2
 bfd minimum-interval 20
 address-family ipv4 unicast
  network 10.1.1.0/24
!
 address-family ipv6 unicast
!
neighbor 192.168.1.1
 remote-as 65001
 address-family ipv4 unicast
!
end
```

Which configuration is missing to complete the configuration task of enabling BFD with the 192.168.1.1 EBGP peer?

A. bfd fast-detect also needs to be enabled globally under router bgp 64500 RP/0/RSP0/CPU0:P1(config-bgp)#bfd fast-detect
B. bfd fast-detect also needs to be enabled for the address-family under address-family ipv4 unicastRP/0/RSP0/CPU0:P1(config-bgp-af)#bfd fast-detect
C. bfd fast-detect also needs to be enabled for the 192.168.1.1 neighbor under neighbor 192.168.1.1RP/0/RSP0/CPU0:P1(config-bgp-nbr)#bfd fast-detect
D. bfd fast-detect also needs to be enabled for the 192.168.1.1 neighbor address-family under neighbor 192.168.1.1 address-family ipv4 unicastRP/0/RSP0/CPU0:P1(config-bgp-nbr-af)#bfd fast-detect
E. bfd fast-detect also needs to be enabled globally on the router RP/0/RSP0/CPU0:P1(config)#bfd fast-detect

Answer: C

NEW QUESTION 59
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 60
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 62
What must occur before an (S,G) entry can be populated in the multicast routing table?

A. The (*,G) entry must have timed out
B. The (*,G) entry OIL must be null
C. The router must be directly connected to the multicast source
D. The parent (*,G) entry must be created first

Answer: D

NEW QUESTION 67
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 69
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.

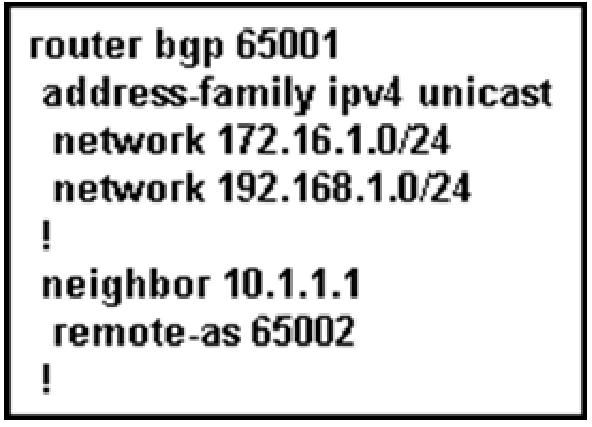

```
router bgp 65001
 address-family ipv4 unicast
  network 172.16.1.0/24
  network 192.168.1.0/24
 !
 neighbor 10.1.1.1
  remote-as 65002
 !
```

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 72
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 73
What is enabled by default on Cisco IOS-XR routers and cannot be disabled?

A. SSH server
B. Multicast routing
C. IPv4 and IPv6 CEF

D. IPv6 routing
E. CDP
F. BFD

Answer: C

Explanation: Before using the BGP policy accounting feature, you must enable BGP on the router (CEF is enabled by default).

NEW QUESTION 74
Which command set is used to implement an IPv6 PIM with the global scope embedded RP address of 2001:DB8::1 on a Cisco IOS XE router?

A. ipv6 unicast-routing ipv6 multicast-routingipv6 pim rp-address 2001:DB8::1 bidir
B. ipv6 multicast-routingipv6 pim rp-address 2001:DB8::1
C. ipv6 unicast-routing ipv6 multicast-routingipv6 pim rp-address FF7E:0120:2001:DB8:1111::4321
D. ipv6 unicast-routing ipv6 multicast-routing int Lo0ipv6 mld join-group FF7E:0120:2001:DB8:1111::4321
E. ipv6 unicast-routing ipv6 multicast-routing int Lo0ipv6 mld join-group FF75:0120:2001:DB8:1111::4321

Answer: D

NEW QUESTION 77
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 81
Which informationdoes the multicast supported router need to forward the multicast traffic over the source or shared tree?

A. source address
B. multicast address
C. destination address
D. mGRE headers
E. MDT Data

Answer: C

NEW QUESTION 85
In secure multicast, which protocol is used to distribute secure keys to a multicast group?

A. ISAKMP
B. RSA
C. IPsec
D. GDOI
E. SKIP

Answer: D

NEW QUESTION 86
Refer to the exhibit.

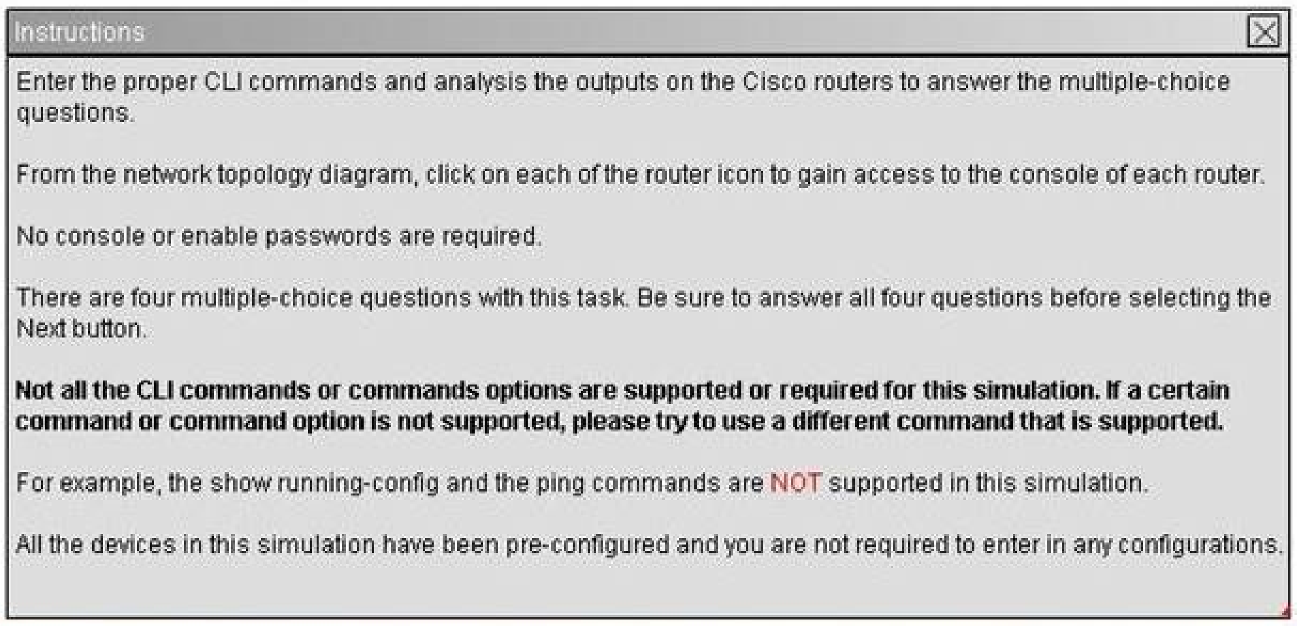

Instructions

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.

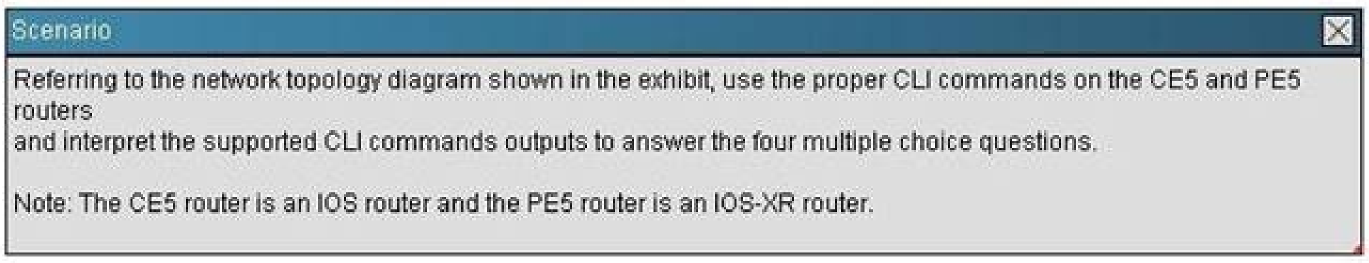

Scenario

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.

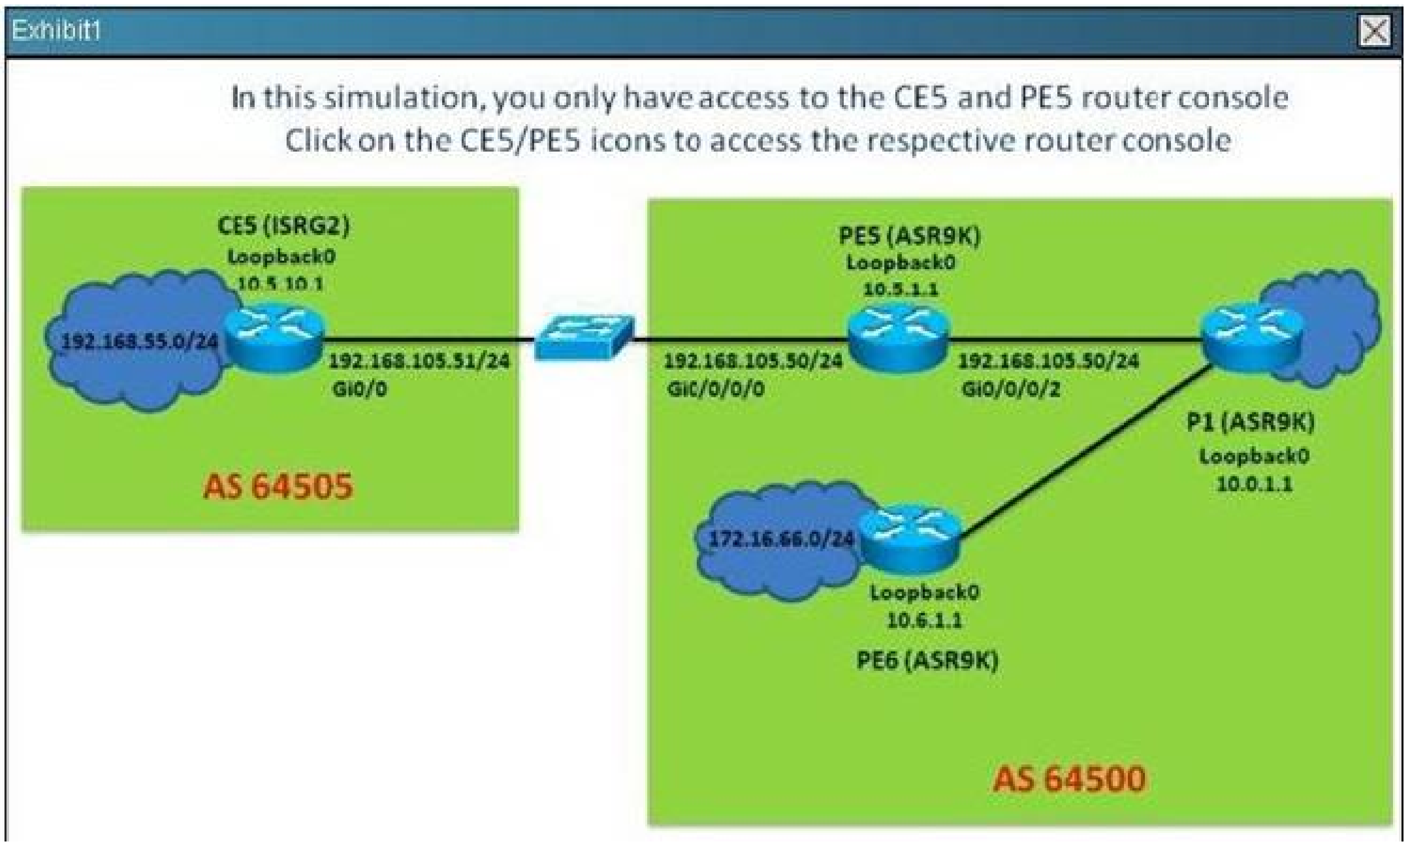

Exhibit1

In this simulation, you only have access to the CE5 and PE5 router console
Click on the CE5/PE5 icons to access the respective router console

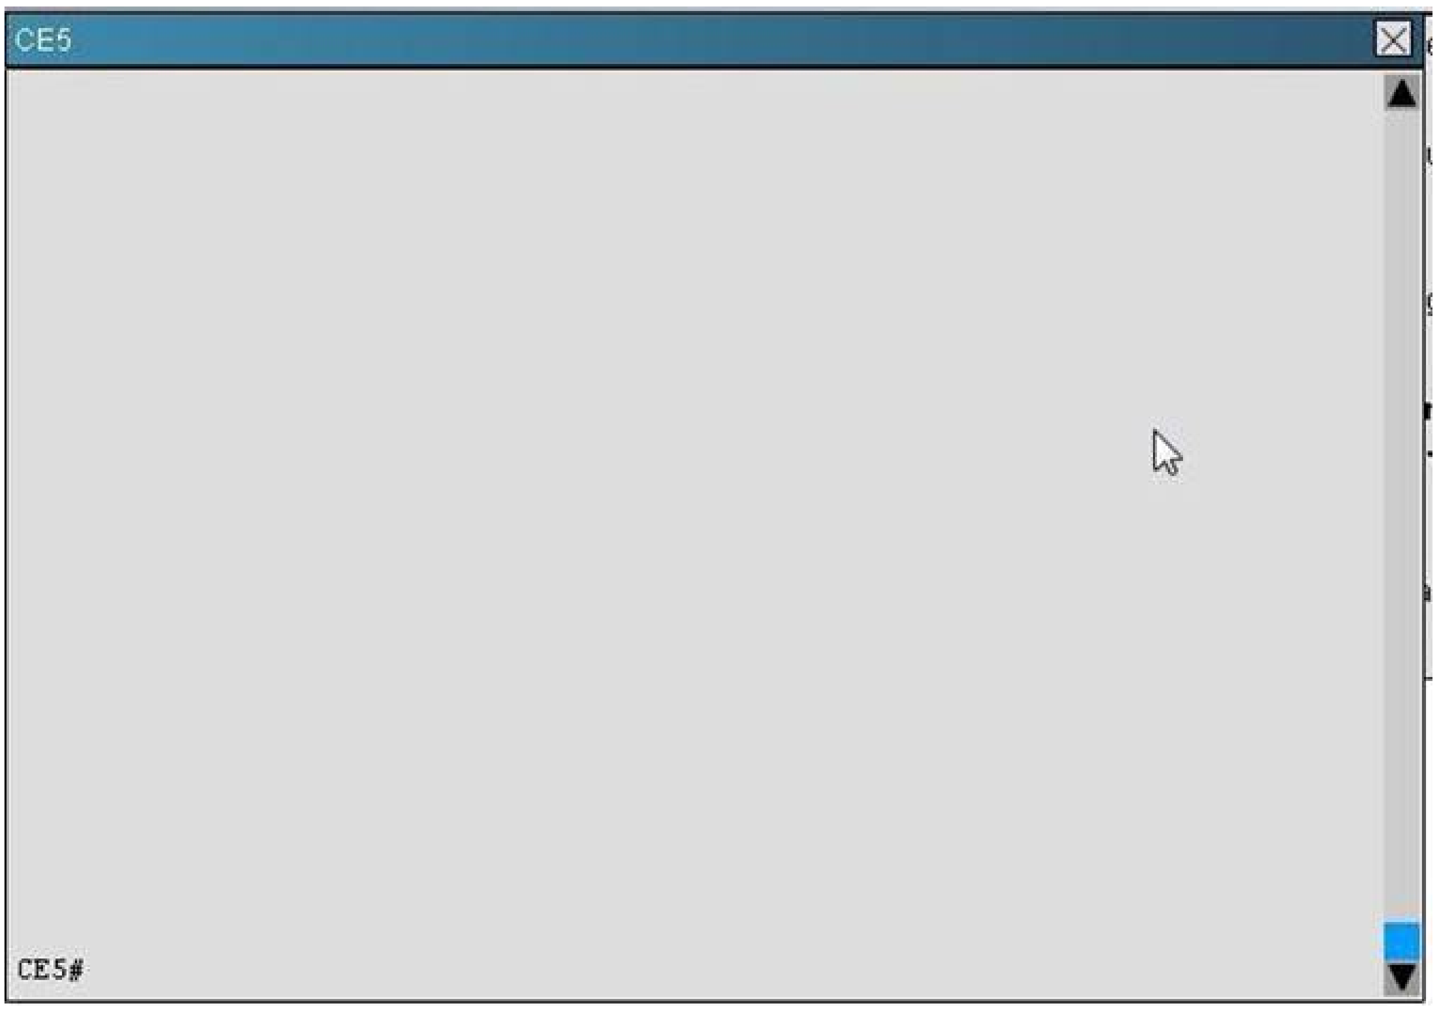

CE5

```
CE5#
```

```
PE5#
```

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 89
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 93
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 94
Refer to the exhibit.

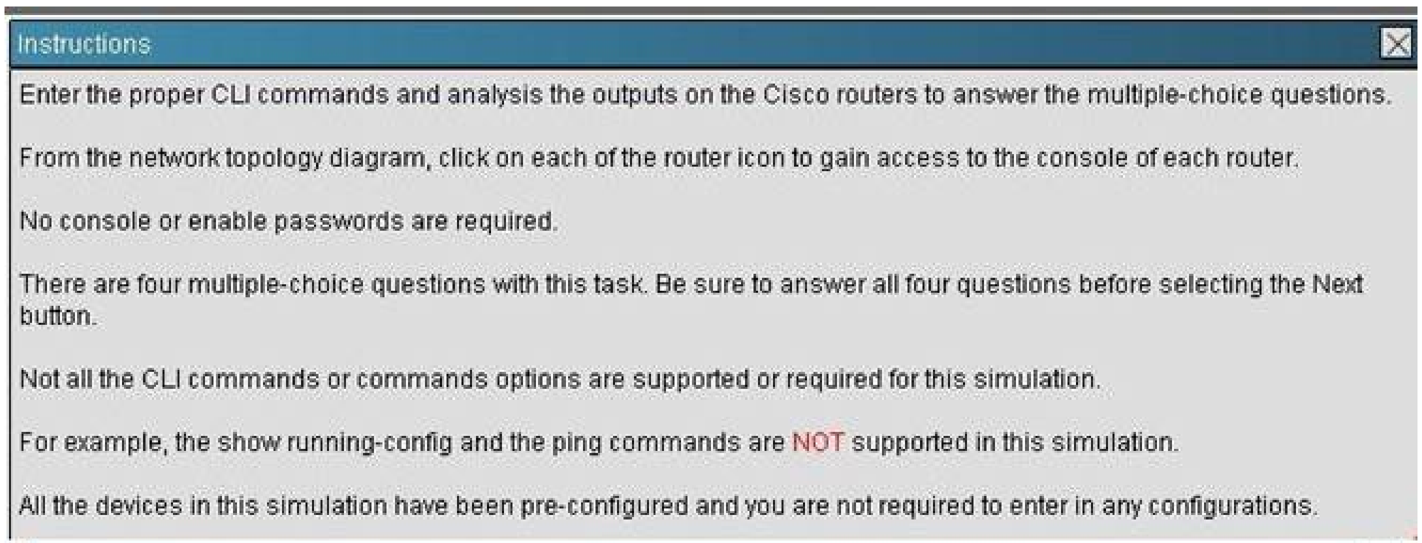

Enter the proper CLI commands and analysis the outputs on the Cisco routers to answer the multiple-choice questions.

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.

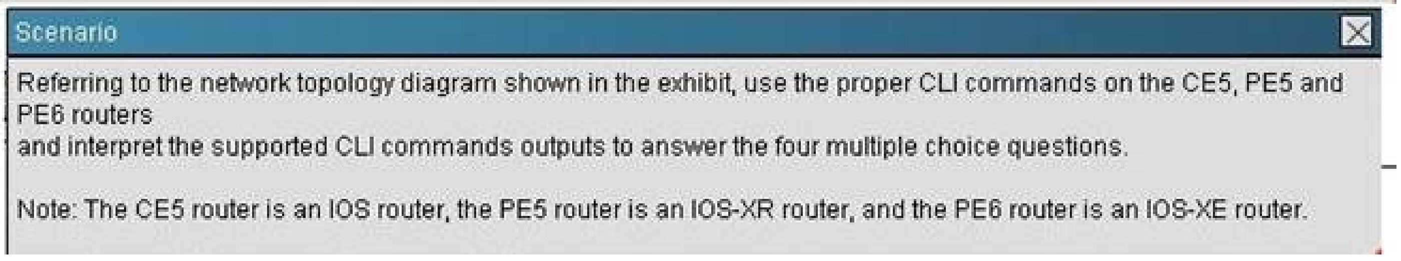

Referring to the network topology diagram shown in the exhibit, use the proper CLI commands on the CE5, PE5 and PE6 routers
and interpret the supported CLI commands outputs to answer the four multiple choice questions.

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.

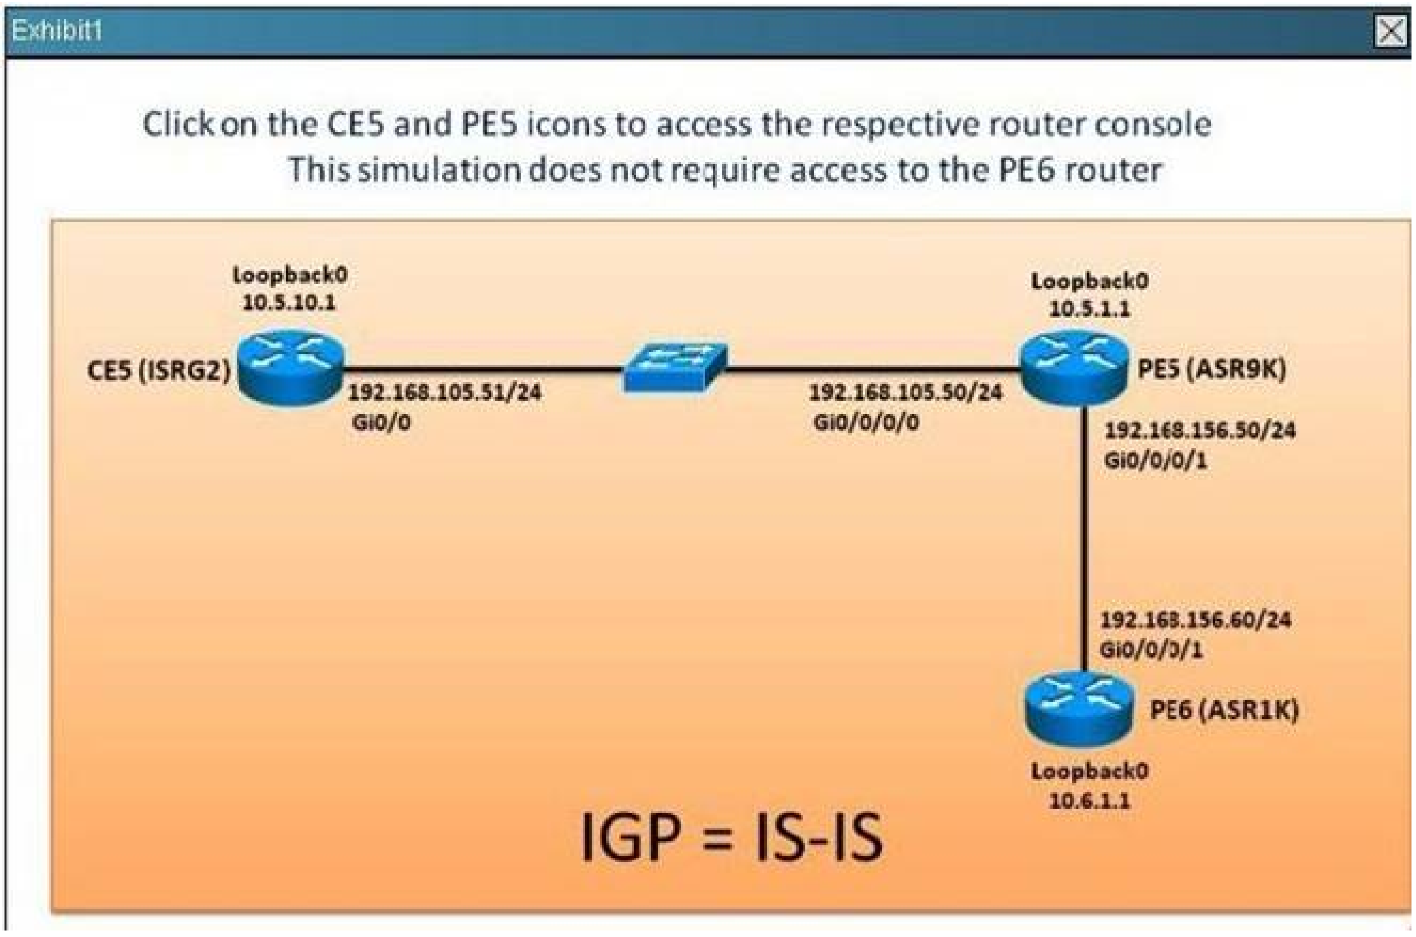

Click on the CE5 and PE5 icons to access the respective router console
This simulation does not require access to the PE6 router

```
CE5#
```

```
PE5#
```

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 99
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 102
The bsr-border router PIM interface configuration command is used for what purpose?

A. To enable the router as the candidate RP
B. To enable the router as the candidate BSR
C. To enable the router as the BSR mapping agent
D. To set up an administrative boundary to prevent BSR messages from being sent out through an interface
E. To define a boundary to restrict the RP discovery and announcement messages from being sent outside the PIM-SM domain

Answer: D

NEW QUESTION 104
When enabling interdomain multicast routing, which two statements are correct? (Choose two.)

A. Multiprotocol BGP is used instead of PIM SM to build the intradomain and interdomain multicast distribution trees
B. Use MSDP to enable the RPs from different domains to exchange information about active multicast sources
C. MSDP SA packets are sent between the multiprotocol BGP peers
D. Noncongruent unicast and multicast topologies can be supported using multiprotocol BGP

Answer: BD

Explanation: http://prakashkalsaria.wordpress.com/2010/08/11/mbgp-msdp/

MSDP In the PIM-SM model, multicast sources and receivers must register with their local RP. Actually, the router closest to the sources or receivers registers with the RP, but the key point to note is that the RP knows about all the sources and receivers for any particular
group. RPs in other domains have no way of knowing about sources located in other domains. MSDP is an elegant way to solve this problem.
MSDP is a mechanism that allows RPs to share information about active sources. RPs know about the receivers in their local domain. When RPs in remote domains hear about the active sources, they can pass on that information to their local receivers and multicast data can then be forwarded between the domains. A useful feature of MSDP is that it allows each domain to maintain an independent RP that does not rely on other domains, but it does enable RPs to forward traffic between domains. PIM-SM is used to forward the traffic between the multicast domains.
The RP in each domain establishes an MSDP peering session using a TCP connection with the RPs in other domains or with border routers leading to the other domains. When the RP learns about a new multicast source within its own domain (through the normal PIM register mechanism), the RP encapsulates the first data packet in a Source-Active (SA) message and sends the SA to all MSDP peers. The SA is forwarded by each receiving peer using a modified RPF check, until the SA reaches every MSDP router in the interconnected networks—theoretically the entire multicast internet. If the receiving MSDP peer is an RP, and the RP has a (*, G) entry for the group in the SA (there is an interested receiver), the RP creates (S, G) state for the source and joins to the shortest path tree for the source. The encapsulated data is decapsulated and forwarded down the shared tree of that RP. When the packet is received by the last hop router of the receiver, the last hop router also may join the shortest path tree to the source. The MSDP speaker periodically sends SAs that include all sources within the own domain of the RP
http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.2/routing/configuration/guide/rc32bgp.html
Multiprotocol BGP
Multiprotocol BGP is an enhanced BGP that carries routing information for multiple network layer protocols and IP multicast routes. BGP carries two sets of routes, one set for unicast routing and one set for multicast routing.
The routes associated with multicast routing are used by the Protocol Independent Multicast (PIM) feature to build data distribution trees.
Multiprotocol BGP is useful when you want a link dedicated to multicast traffic, perhaps to limit which resources are used for which traffic. Multiprotocol BGP allows you to have a unicast routing topology different from a multicast routing topology providing more control over your network and resources.
In BGP, the only way to perform interdomain multicast routing was to use the BGP infrastructure that was in place for unicast routing. Perhaps you want all multicast traffic exchanged at one network access point (NAP).
If those routers were not multicast capable, or there were differing policies for which you wanted multicast traffic to flow, multicast routing could not be supported without multiprotocol BGP.
Note It is possible to configure BGP peers that exchange both unicast and multicast network layer reachability information (NLRI), but you cannot connect multiprotocol BGP clouds with a BGP cloud. That is, you cannot redistribute multiprotocol BGP routes into BGP.

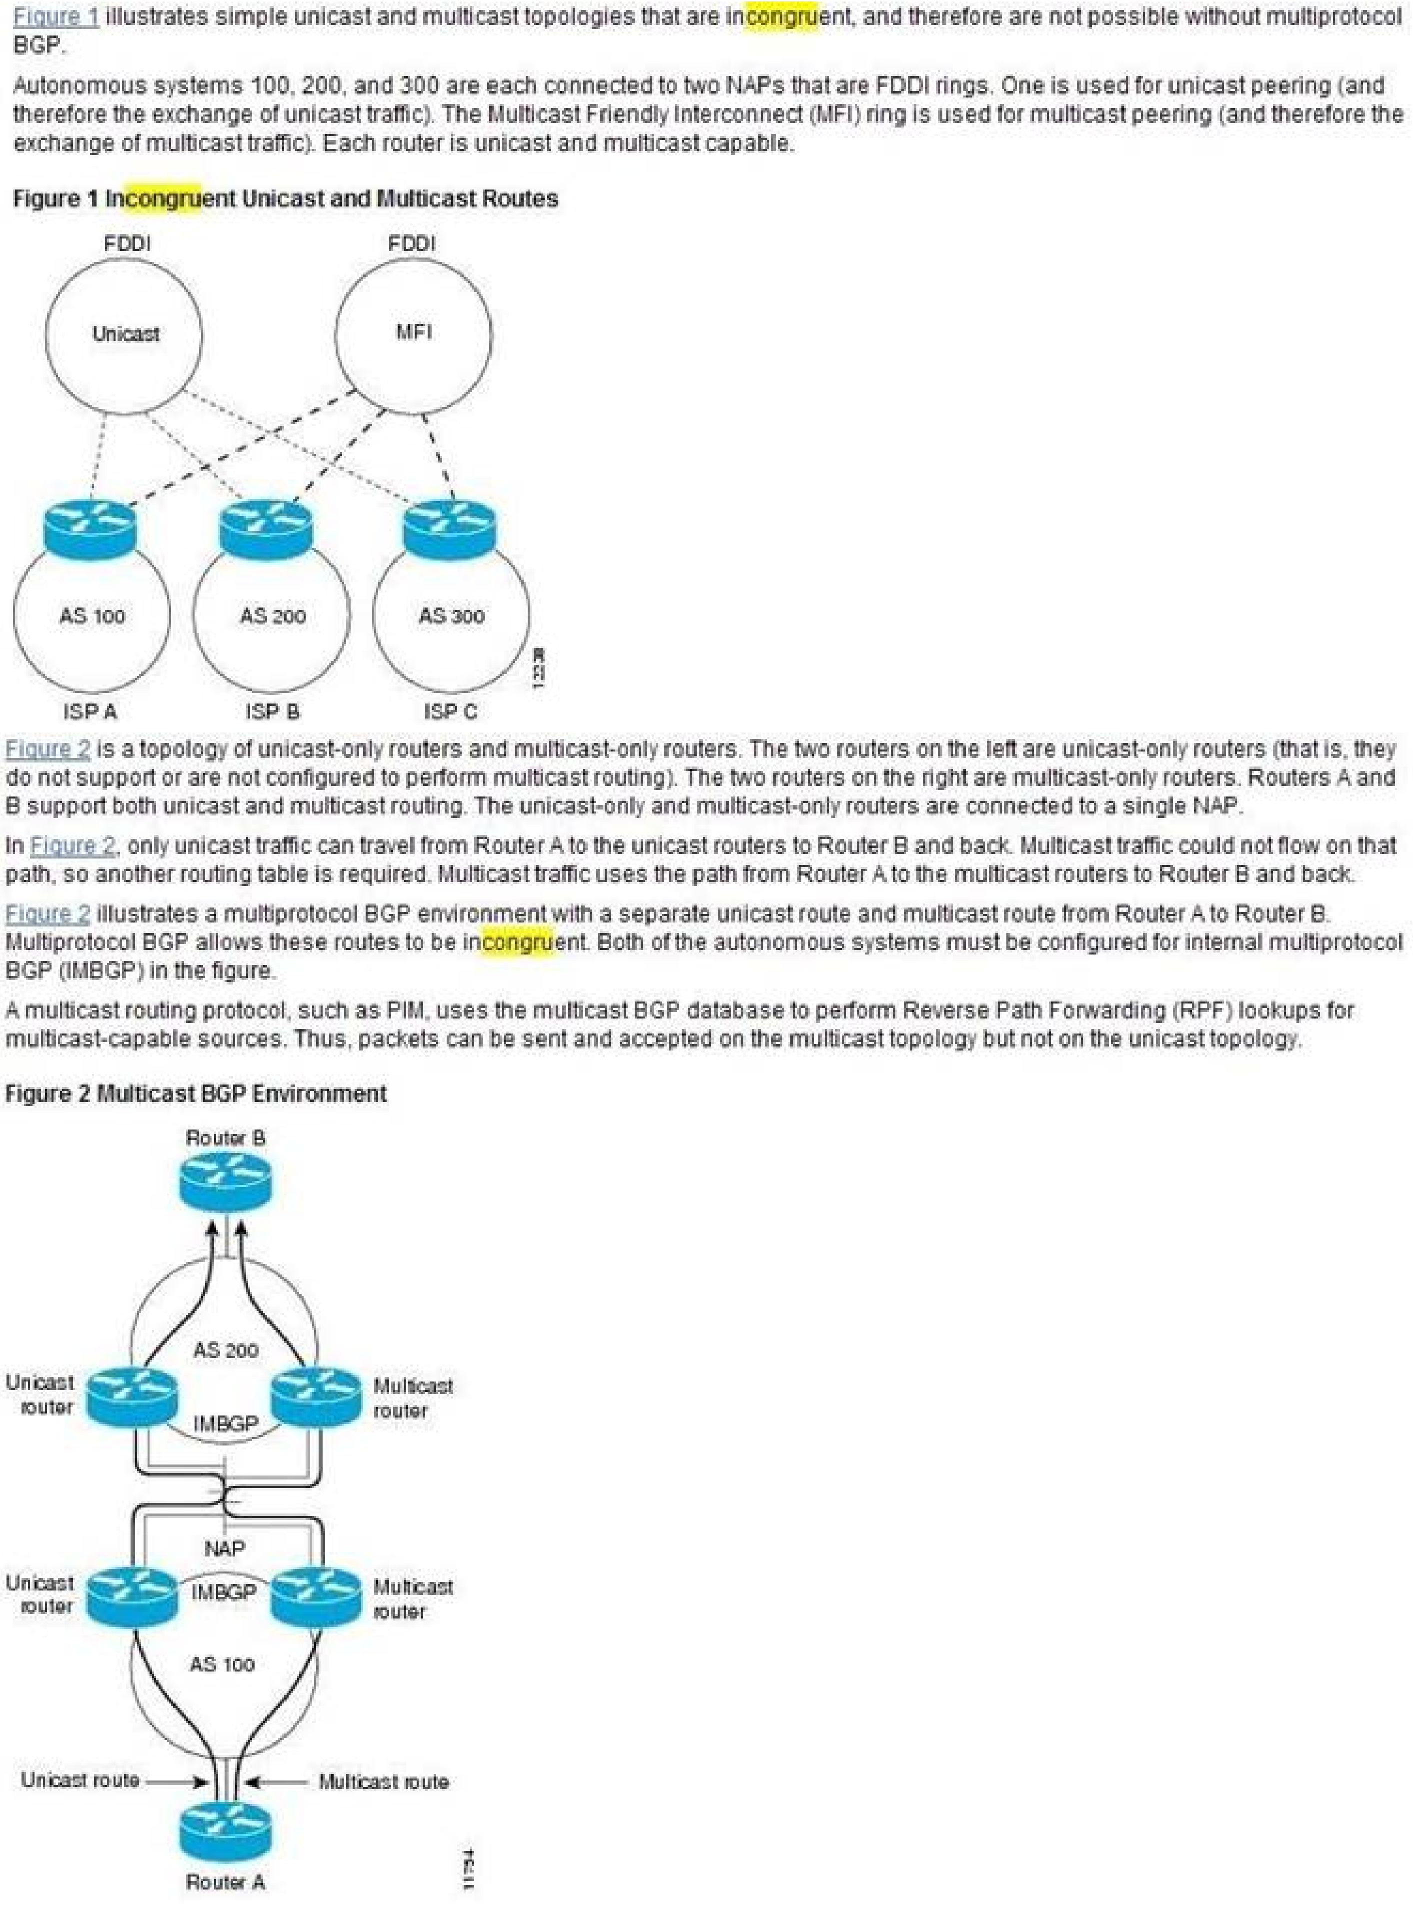

Figure 1 illustrates simple unicast and multicast topologies that are incongruent, and therefore are not possible without multiprotocol BGP.

Autonomous systems 100, 200, and 300 are each connected to two NAPs that are FDDI rings. One is used for unicast peering (and therefore the exchange of unicast traffic). The Multicast Friendly Interconnect (MFI) ring is used for multicast peering (and therefore the exchange of multicast traffic). Each router is unicast and multicast capable.

Figure 1 Incongruent Unicast and Multicast Routes

Figure 2 is a topology of unicast-only routers and multicast-only routers. The two routers on the left are unicast-only routers (that is, they do not support or are not configured to perform multicast routing). The two routers on the right are multicast-only routers. Routers A and B support both unicast and multicast routing. The unicast-only and multicast-only routers are connected to a single NAP.

In Figure 2, only unicast traffic can travel from Router A to the unicast routers to Router B and back. Multicast traffic could not flow on that path, so another routing table is required. Multicast traffic uses the path from Router A to the multicast routers to Router B and back.

Figure 2 illustrates a multiprotocol BGP environment with a separate unicast route and multicast route from Router A to Router B. Multiprotocol BGP allows these routes to be incongruent. Both of the autonomous systems must be configured for internal multiprotocol BGP (IMBGP) in the figure.

A multicast routing protocol, such as PIM, uses the multicast BGP database to perform Reverse Path Forwarding (RPF) lookups for multicast-capable sources. Thus, packets can be sent and accepted on the multicast topology but not on the unicast topology.

Figure 2 Multicast BGP Environment

NEW QUESTION 109
Which keyword is used in the syntax to refer to Cisco IOS XR address-family groups, session groups, or neighbor groups?

A. inherit
B. apply
C. use
D. commit

Answer: C

NEW QUESTION 112
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 114
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 118
Which multicast routing protocol is used to forward multicast data along the optimal path from source to receivers?

A. PIM DM
B. PIM Bi-Dir
C. PIM SM
D. SSM
E. IGMP
F. MSDP

Answer: C

NEW QUESTION 122
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 123
A network engineer must configure a Cisco IOS XR router with BGP dampening. Which configuration meets these parameters?

A. router bgp 60 bgp dampening
B. router bgp 60 neighbor 10.0.0.2 bgp dampening
C. router bgp 60address-family ipv4 unicast bgp dampening
D. route-policy dampening_specific drop!router bgp 60address-family ipv4 unicastbgp dampening route-policy dampening_specific
E. router bgp 60 address-family ipv4 bgp dampening

Answer: C

NEW QUESTION 124
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 126
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 129
Refer to the Cisco IOS DHCPv6 configuration shown in the exhibit.

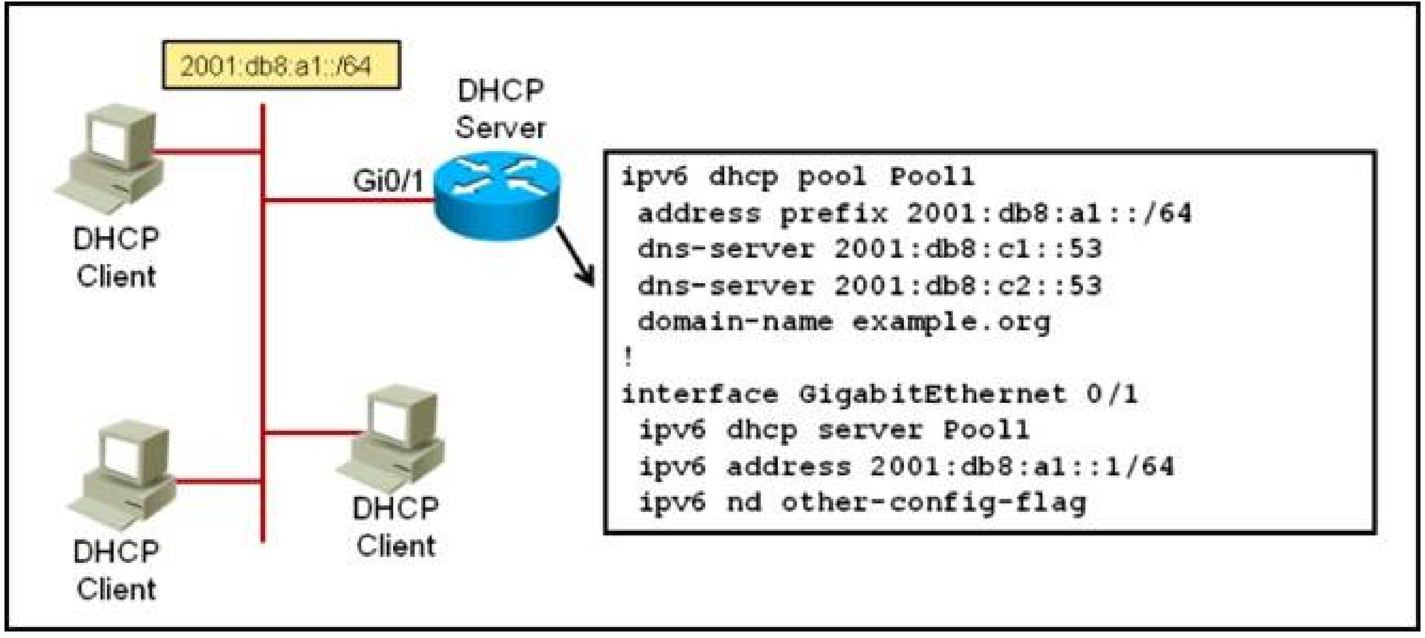

```
ipv6 dhcp pool Pool1
 address prefix 2001:db8:a1::/64
 dns-server 2001:db8:c1::53
 dns-server 2001:db8:c2::53
 domain-name example.org
!
interface GigabitEthernet 0/1
 ipv6 dhcp server Pool1
 ipv6 address 2001:db8:a1::1/64
 ipv6 nd other-config-flag
```

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 131
Which three methods can be used to reduce the full-mesh IBGP requirement in a service provider core network? (Choose three.)

A. Implement route reflectors

B. Enable multi-protocol BGP sessions between all the PE routers
C. Implement confederations
D. Implement MPLS (LDP) in the core network on all the PE and P routers
E. Enable BGP synchronization
F. Disable the IBGP split-horizon rule

Answer: ACD

NEW QUESTION 135
Refer to the exhibit. R2cannot to learn RP information from XR3. Which issue is the source of the problem?

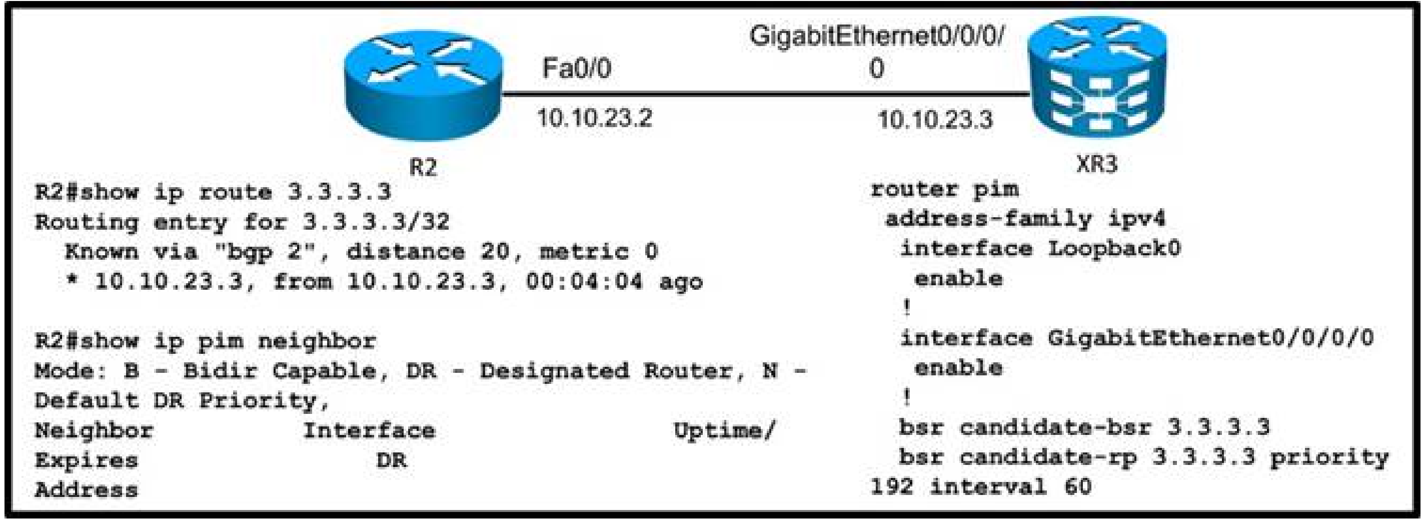

```
R2#show ip route 3.3.3.3
Routing entry for 3.3.3.3/32
  Known via "bgp 2", distance 20, metric 0
  * 10.10.23.3, from 10.10.23.3, 00:04:04 ago

R2#show ip pim neighbor
Mode: B - Bidir Capable, DR - Designated Router, N -
Default DR Priority,
Neighbor            Interface                  Uptime/
Expires                 DR
Address
```

```
router pim
 address-family ipv4
  interface Loopback0
   enable
  !
  interface GigabitEthernet0/0/0/0
   enable
  !
  bsr candidate-bsr 3.3.3.3
  bsr candidate-rp 3.3.3.3 priority
192 interval 60
```

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 139
In which three cases is a dual-stack IPv6/IPv4 router required? (Choose three.)

A. tunnel endpoint routers in the case of IPv6 over GRE
B. transit routers in case of an IPv6 over GRE implementation
C. 6to4 implementation border routers
D. 6to4 implementation border and neighboring routers
E. PE routers in case of an IPv6 over IPv4 tunnel over MPLS implementation
F. PE and P routers in case of an IPv6 over IPv4 tunnel over MPLS implementation

Answer: ACE

NEW QUESTION 140
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 142
Refer to the exhibit.

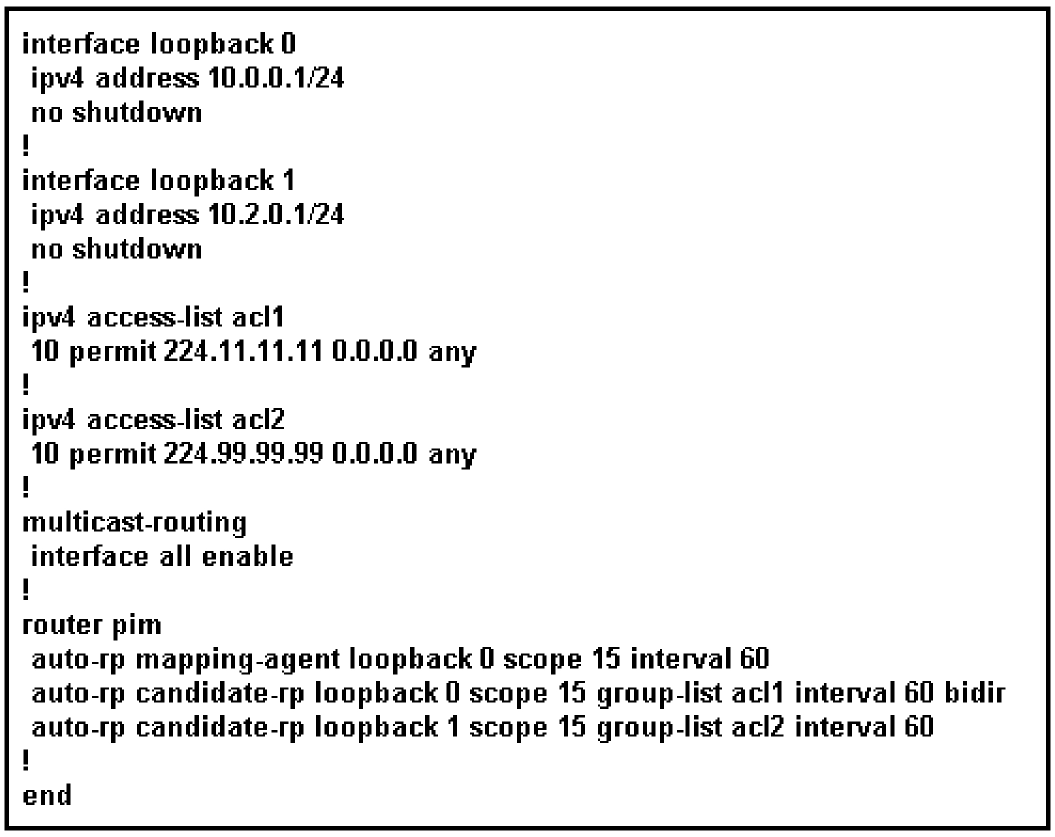

```
interface loopback 0
 ipv4 address 10.0.0.1/24
 no shutdown
!
interface loopback 1
 ipv4 address 10.2.0.1/24
 no shutdown
!
ipv4 access-list acl1
 10 permit 224.11.11.11 0.0.0.0 any
!
ipv4 access-list acl2
 10 permit 224.99.99.99 0.0.0.0 any
!
multicast-routing
 interface all enable
!
router pim
 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 144
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 146
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 151
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 154
Which additional feature is provided using MLDv2 that is not available in MLDv1?

A. Multicast Address Specific Queries
B. Source filtering
C. Done messages
D. Report messages

Answer: B

Explanation: • PIM-SSM is made possible by IGMPv3 and MLDv2. Hosts can now indicate interest in specific sources using IGMPv3 and MLDv2. SSM does not require a rendezvous point (RP) to operate.

NEW QUESTION 157
Refer to the exhibit.

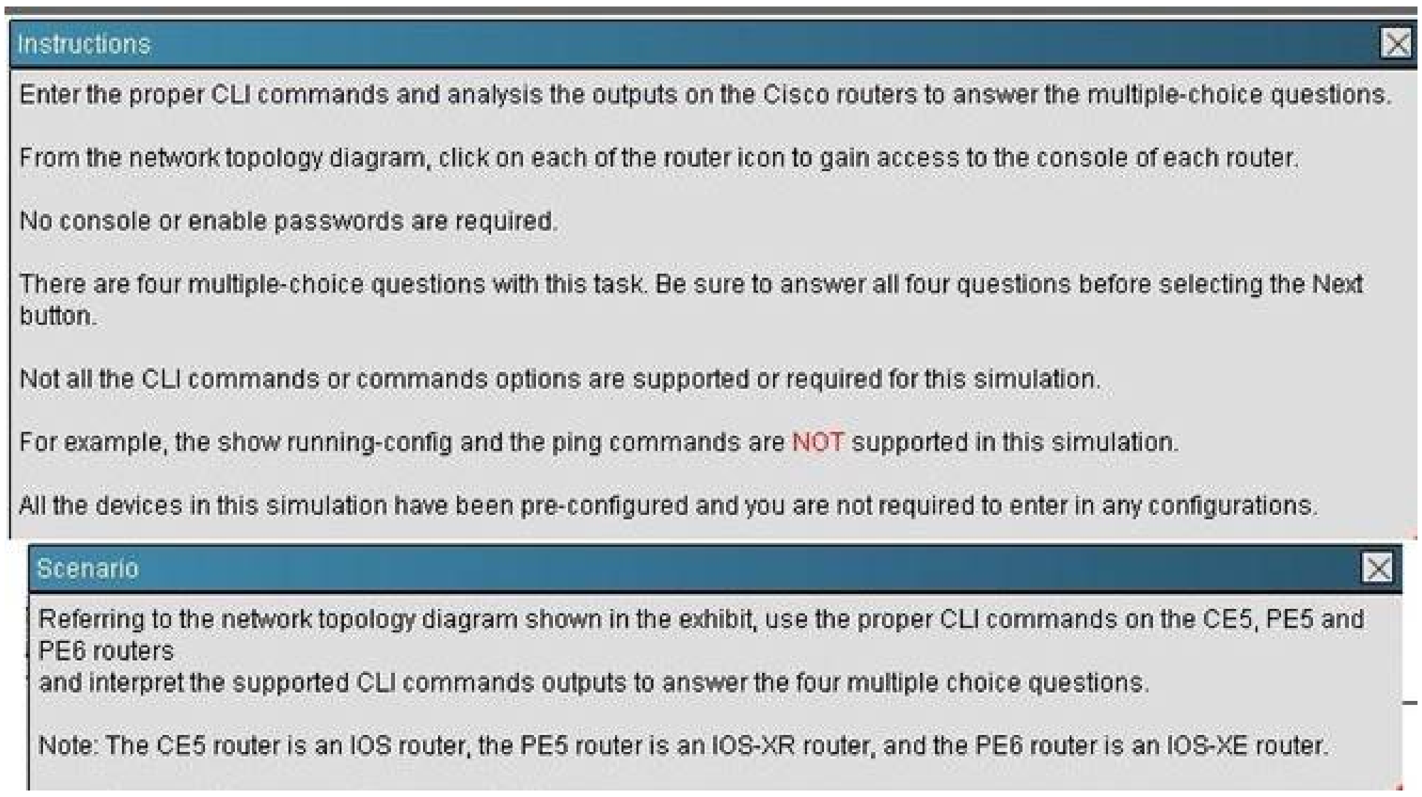

Instructions

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.

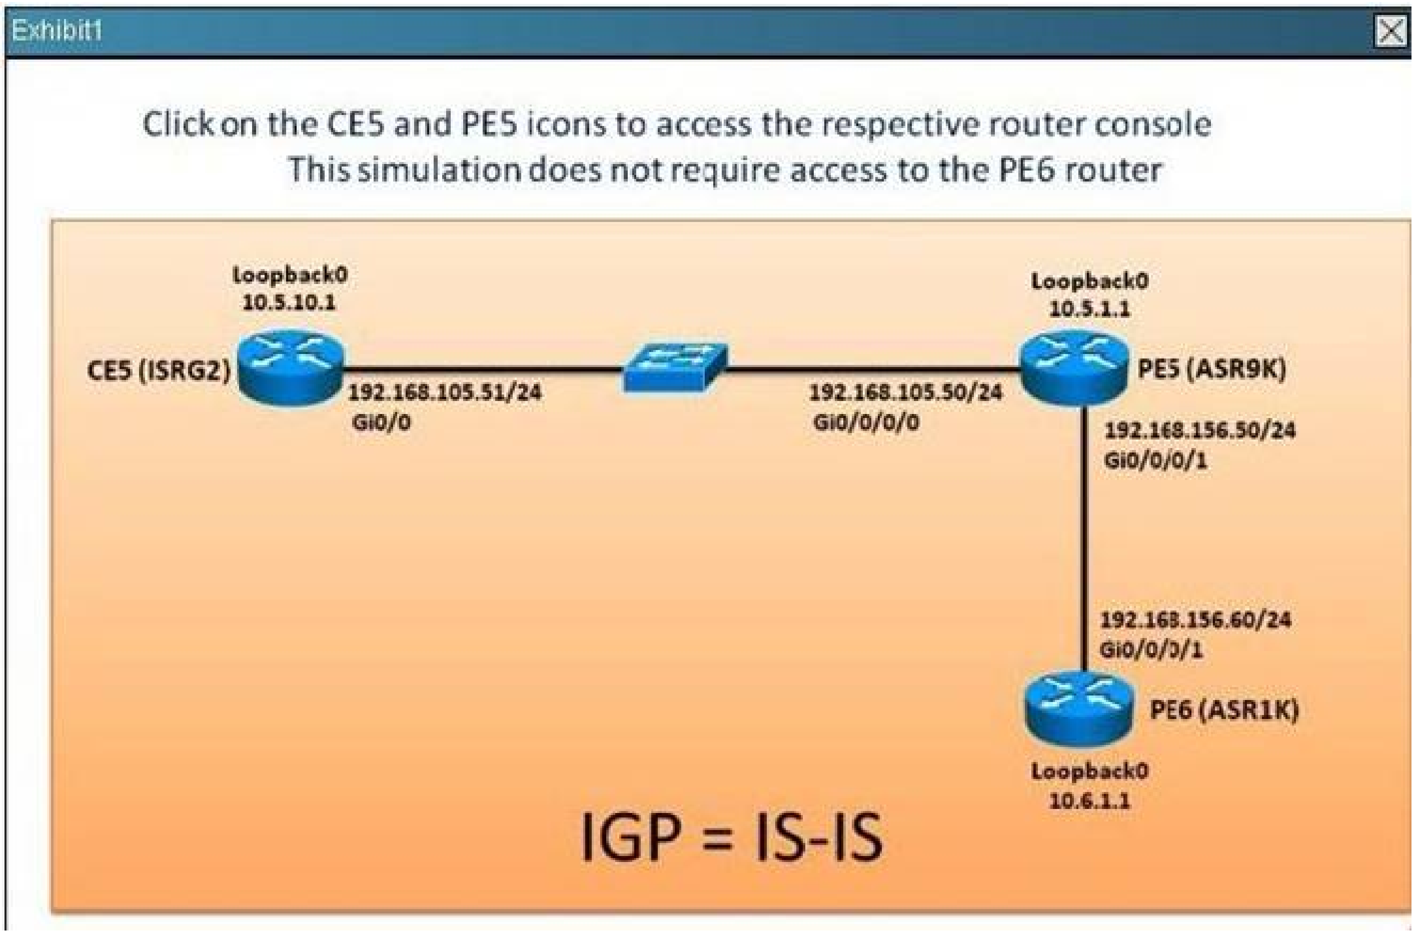

Click on the CE5 and PE5 icons to access the respective router console
This simulation does not require access to the PE6 router

```
CE5#
```

```
PE5#
```

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 159
Refer to the Cisco IOS-XR BGP configuration exhibit.

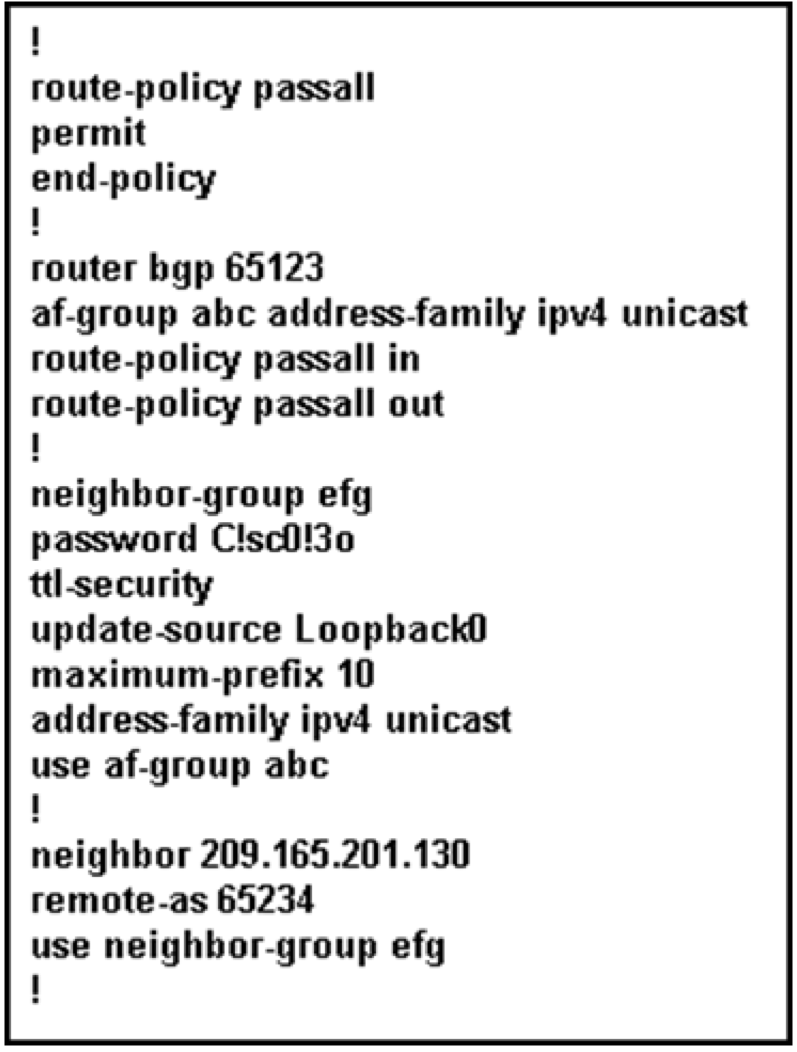

```
!
route-policy passall
permit
end-policy
!
router bgp 65123
af-group abc address-family ipv4 unicast
route-policy passall in
route-policy passall out
!
neighbor-group efg
password C!sc0!3o
ttl-security
update-source Loopback0
maximum-prefix 10
address-family ipv4 unicast
use af-group abc
!
neighbor 209.165.201.130
remote-as 65234
use neighbor-group efg
!
```

Identify two configuration errors. (Choose two.)

A. The neighbor-group efg is missing the ebgp-multihop 2 configuration
B. The ttl-security configuration command is missing the option to set the number of hops
C. The passall route policy is wrong

D. The route-policy passall in and route-policy passall out commands should be configured under the neighbor-group efg instead of the af-group abc
E. The maximum-prefix 10 configuration should be configured under the af-group abc instead of the neighbor-group efg

Answer: CE

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

NEW QUESTION 161
Which two statements regarding Auto RP operations and implementations are correct? (Choose two.)

A. Candidate RPs send RP announcements to the 224.0.1.39 multicast group, and the mapping agents send RP discovery messages to the 224.0.1.40 multicast group
B. Every PIM-SM router must be configured with the RP mapping agent IP address
C. Candidate RPs learn the IP address of the mapping agents via periodic RP discovery messages
D. Administrative scoping can be configured to limit the scope of the RP announcements
E. A Reverse Path Forwarding check is done on the RP discovery messages
F. RP discovery messages are flooded hop by hop throughout the network as multicast to the all PIM routers multicast group with a TTL of 1

Answer: AD

Explanation: Auto-RP
Automatic route processing (Auto-RP) is a feature that automates the distribution of group- to-RP mappings in a PIM network. This feature has these benefits:
It is easy to use multiple RPs within a network to serve different group ranges. It allows load splitting among different RPs.
It facilitates the arrangement of RPs according to the location of group participants.
It avoids inconsistent, manual RP configurations that might cause connectivity problems. Multiple RPs can be used to serve different group ranges or to serve as hot backups for each other. To ensure that Auto-RP functions, configure routers as candidate RPs so that they can announce their interest in operating as an RP for certain group ranges. Additionally, a router must be designated as an RP-mapping agent that receives the RP- announcement messages from the candidate RPs, and arbitrates conflicts. The RPmapping agent sends the consistent group-to-RP mappings to all remaining routers. Thus, all routers automatically determine which RP to use for the groups they support auto- rp candidate-rp
To configure a router as a Protocol Independent Multicast (PIM) rendezvous point (RP) candidate that sends messages to the well-known CISCO-RP-ANNOUNCE multicast group
(224.0.1.39), use the auto-rp candidaterp command in PIM configuration mode. To return to the default behavior, use the no form of this command. auto-rp candidate-rp type interface-path-id scope ttl-value [group-list access-listname] [interval seconds] [bidir] no auto-rp candidate-rp type interface-path-id scope ttl-value [group-list access-listname] [interval seconds] [bidir]

NEW QUESTION 166
What is one of the configuration errors within an AS that can stop a Cisco IOS-XR router from announcing certain prefixes to its EBGP peers?

A. Some prefixes were mistagged with the no-export BGP community
B. Some prefixes were set with an MED of 0
C. The outbound BGP route policy only has set actions defined without any pass actions defined
D. The inbound BGP route policy only has set actions defined without any pass actions defined

Answer: A

NEW QUESTION 168
Refer to the exhibit.

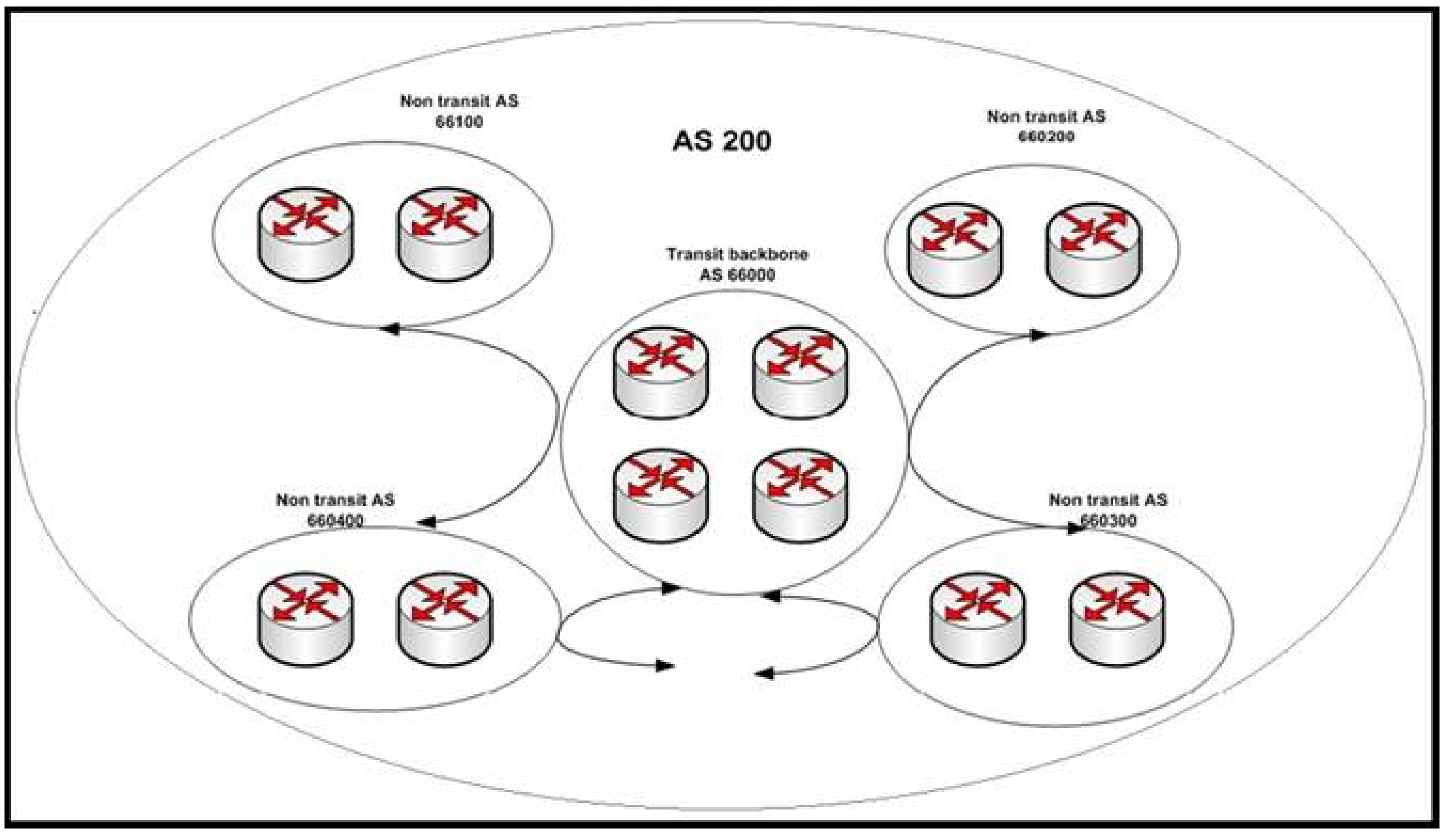

Which option is the function of designing a hub and spoke confederation?

A. allows transit backbone area 66000 to be a blackhole for non-transit ASs
B. reduces the iBGP mesh, iBGP mesh will be in sub non-transit ASs
C. increases eBGP sessions between the confederation sub ASs
D. allows transit backbone area and non-transit ASs to run the same IGP

Answer: B

NEW QUESTION 172
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 177
Which of the following is a feature added in IGMPv3?

A. Support for source filtering
B. Support for Host Membership Report and a Leave Group message
C. Uses a new variation of the Host Membership Query called the Group-Specific Host Membership Query
D. Uses an election process to determine the querying router on the LAN
E. Uses an election process to determine the designated router on the LAN
F. IPv6 support

Answer: A

NEW QUESTION 181
Which two options are characteristics ofconfiguration templates used by Cisco IOS XRto optimize BGP peering implementations? (Choose two.)

A. Session groups are used to inherit address family-specific configurations.
B. Cisco IOS XR provides by default a session group operating with all the supported address families.
C. Session groups are used to inherit address family-independent configurations.
D. Session groups can be included within a neighbor group.
E. Session groups can include neighbor groups.

Answer: CD

NEW QUESTION 182
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 187
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 188
Refer to the exhibit.

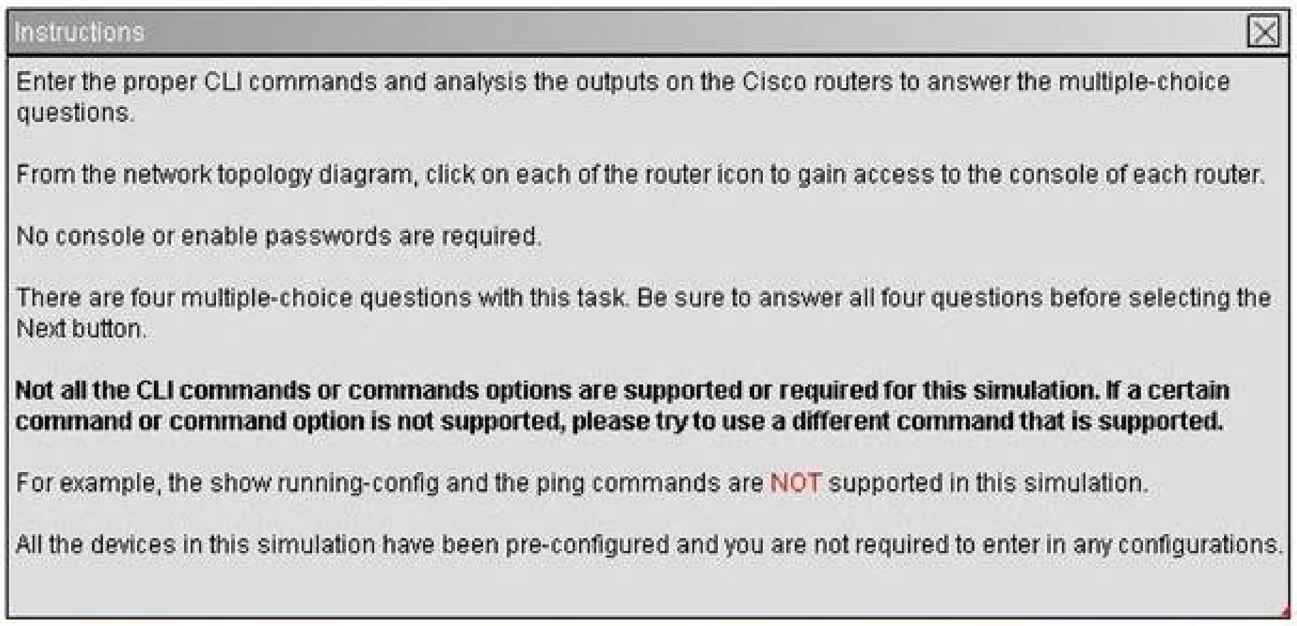

Instructions

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.

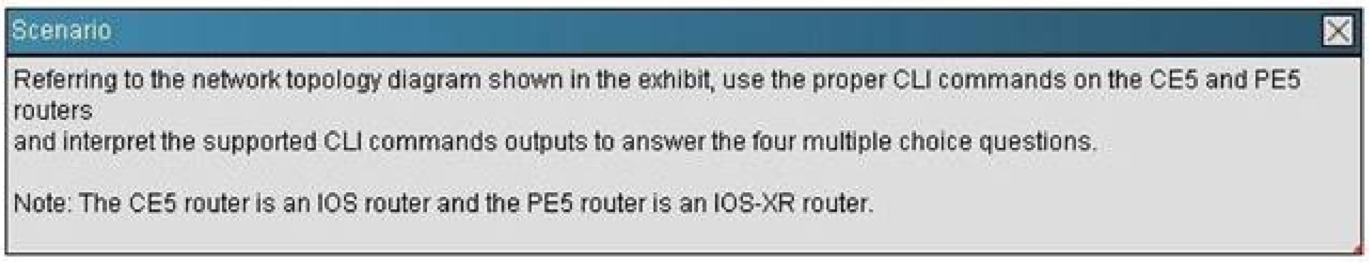

Scenario

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.

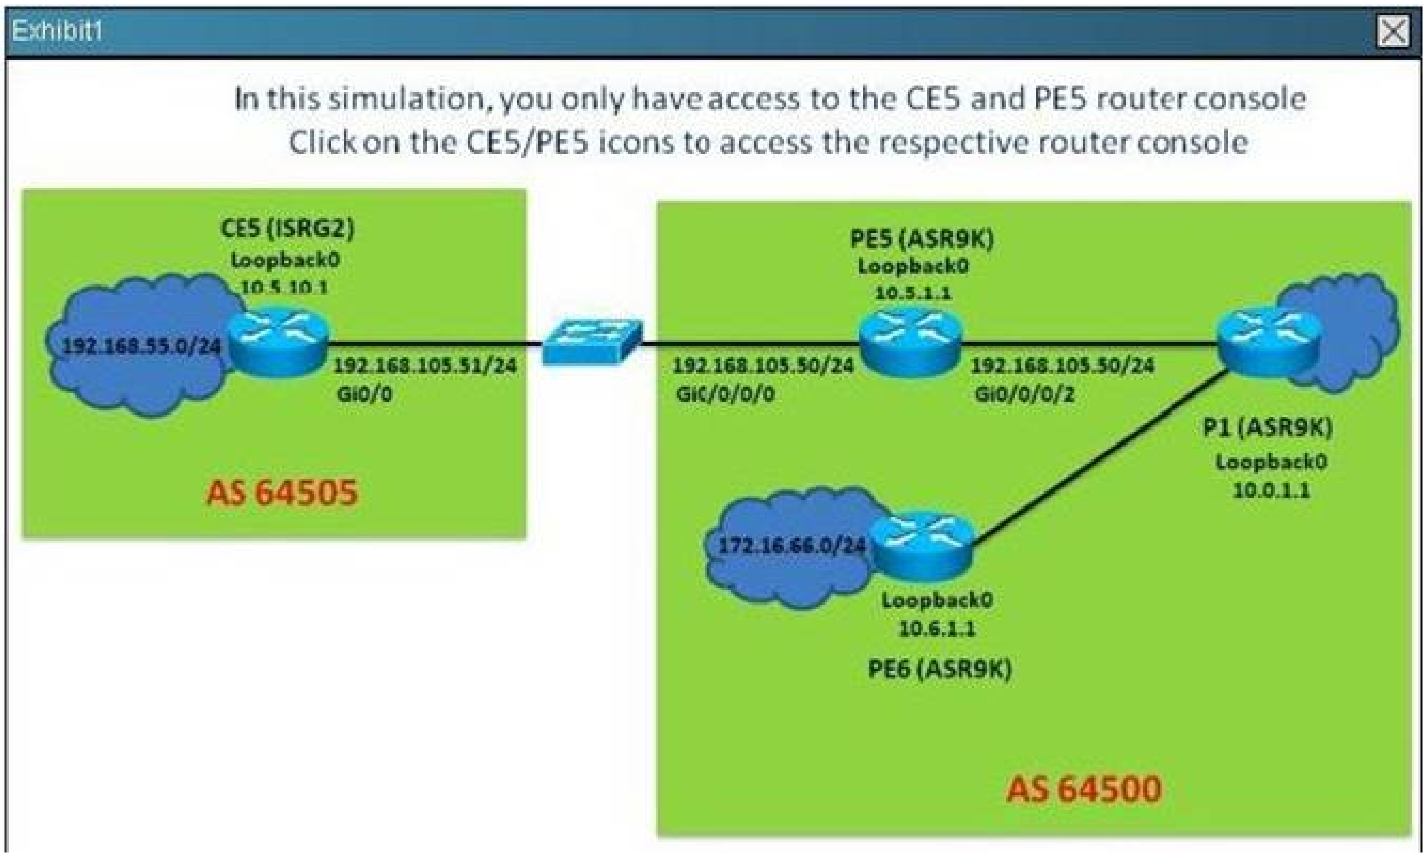

Exhibit1

In this simulation, you only have access to the CE5 and PE5 router console
Click on the CE5/PE5 icons to access the respective router console

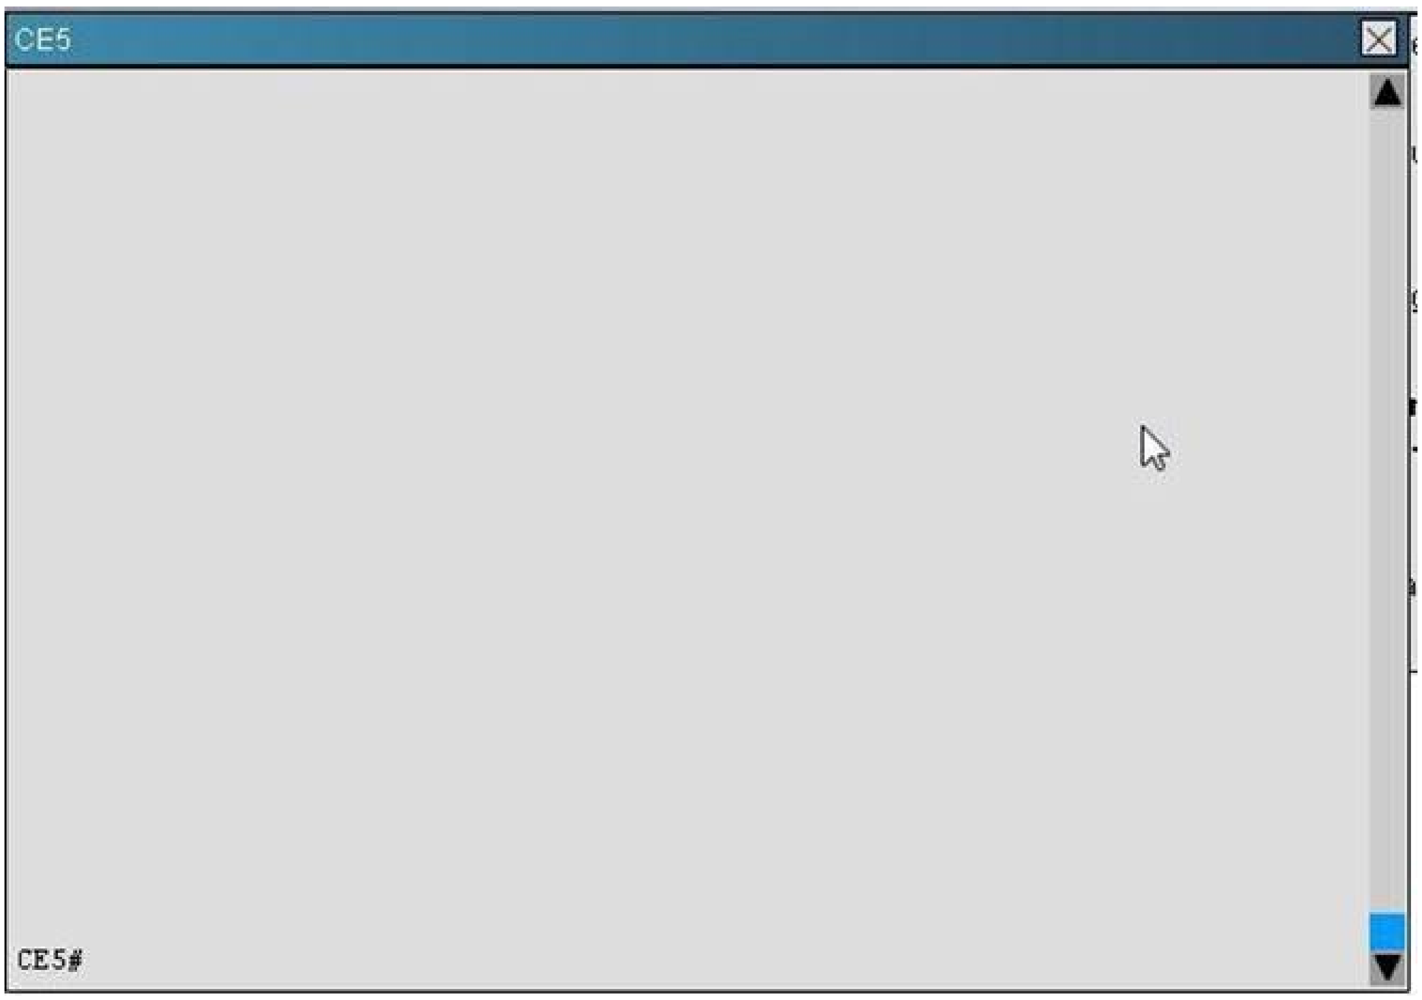

CE5

```
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 193
Which command set is used to configure BFD support for a BGP neighbor that is reachable through GigabitEthernet 0/0/0/0 on Cisco IOS XR?

A. router bgp 300 bfd multiplier 2bfd minimum-interval 20neighbor 10.20.20.2remote-as 200
B. router bgp 300 bfd multiplier 2bfd minimum-interval 20neighbor 10.20.20.2remote-as 200 bfd fast-detect
C. bfdecho disable router bgp 300neighbor 10.20.20.2remote-as 200
D. bfdrouter bgp 300neighbor 10.20.20.2remote-as 200
E. interface Gi0/0/0/0ipv4 verify unicast source reachable-via rx router bgp 300bfd multiplier 2bfd minimum-interval 20neighbor 10.20.20.2remote-as 200 bfd fast-detect
F. interface Gi0/0/0/0ipv4 verify unicast source reachable-via rx bfdinterface Gi0/0/0/0 echo disable router bgp 300bfd multiplier 2bfd minimum-interval 20neighbor 10.20.20.2remote-as 200

Answer: B

NEW QUESTION 198
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 199
Refer to the Cisco IOS configuration exhibit.

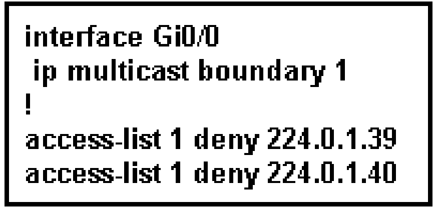

```
interface Gi0/0
 ip multicast boundary 1
!
access-list 1 deny 224.0.1.39
access-list 1 deny 224.0.1.40
```

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 203
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 208
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 210
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 213
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 218
Whichtwo attributes does BGP select before MED? (Choose two.)

A. local preference

B. weight
C. lowest router ID
D. lowest neighbor IP
E. oldest route

Answer: AB

NEW QUESTION 223
Which mechanism is used by an IPv6 multicast receiver to join an IPv6 multicast group?

A. IGMP report
B. IGMP join
C. MLD report
D. General query
E. PIM join

Answer: C

Explanation: MLD Reports
The processing of MLDv1 join messages is essentially the same as with IGMPv2. When no IPv6 multicast routers are detected in a VLAN, reports are not processed or forwarded from the switch.
When IPv6 multicast
routers are detected and an MLDv1 report is received, an IPv6 multicast group address and an IPv6 multicast MAC address are entered in the VLAN MLD database. Then all IPv6 multicast traffic to the group within the VLAN is forwarded using this address. When MLD snooping is disabled, reports are flooded in the ingress VLAN.
When MLD snooping is enabled, MLD report suppression, called listener message suppression, is automatically enabled. With report suppression, the switch forwards the first MLDv1 report received by a group to IPv6 multicast routers; subsequent reports for the group are not sent to the routers. When MLD snooping is disabled, report suppression is disabled, and all MLDv1 reports are flooded to the ingress VLAN.
The switch also supports MLDv1 proxy reporting. When an MLDv1 MASQ is received, the switch responds with MLDv1 reports for the address on which the query arrived if the group exists in the switch on another port and if the port on which the query arrived is not the last member port for the address.

NEW QUESTION 227
Refer to the exhibit.

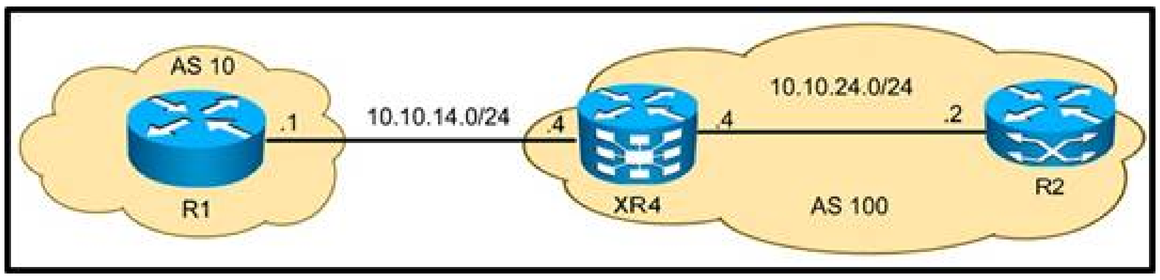

XR4 must protect itself from a DOS attack against its BGP process from R1 by using the TTL security feature. Which configuration achieves this goal?

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 230
......