Cisco

Exam 642-885

Deploying Cisco Service Provider Advanced Network Routing

Version: 7.0

[Total Questions: 131]

Question No : 1

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

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

Answer: D

Question No : 2

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

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

Answer: B

Question No : 3

The following Cisco IOS-XR configuration command will globally enable which multicast process(es) on the router?

RP/0/RP0/CPU0:router(config)# multicast-routing

A. IGMP only
B. PIM only
C. IGMP and MLD only
D. PIM and IGMP only
E. PIM and IGMP and MLD

Answer: E

Explanation:

http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.5/multicast/configuration/guide/mc35mcst.html

Multicast-routing Configuration Submode
When you issue the multicast-routing ipv4 or multicast-routing ipv6 command, all default multicast components (PIM, IGMP, MLD, MFWD, and MRIB) are automatically started, and the CLI prompt changes to "config-mcastipv4" or "config-mcast-ipv6", indicating that you have entered multicast-routing configuration submode

Question No : 4

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

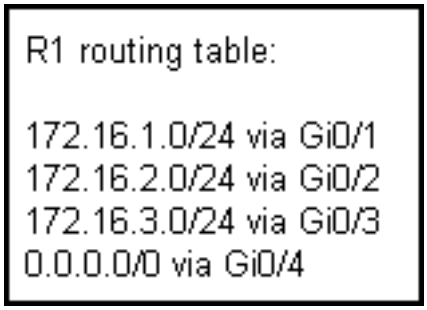

The multicast packet from the 172.16.1.100 source must arrive on which interface on the R1 router for it to be forwarded out the other interfaces?

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

Answer: A

Question No : 5

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: A,C,D

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).

Question No : 6

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 RTD
RTB(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 192.168.1.2
RTB(config-router)#no bgp client-to-client reflection

Answer: A

Question No : 7

Refer to the exhibit.

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

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

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

Answer: D

Question No : 8

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

Question No : 9

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.

Question No : 10

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: B,D

Question No : 11

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

Question No : 12

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: A,B

Question No : 13

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: A,D

Question No : 14

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

Question No : 15

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 the customer 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 updates. This configuration permits VPNv4 communication by accepting the iBGP updates even if no route targets are configured locally.

Answer: D