Cisco

300-410

Implementing Cis Enterprise Advanced Routing and Services (ENARSI)

Question & Answers

(Free – Demo Version)

Buy Full Product by Clicking on the Giving Link:

https://certya.com/product/300-410-exam-questions-answers/

QUESTION: 1
Automatic 6-to-4 tunnels exist between dual-stack routers (A, B, and C). One router has the IPv6 address, 2002:D030:6BC0:173C::26:37D0/48 Which of the following addresses is the IPv4 address of the router with the IPv6 address 2002:D030:6BC0:173C::26:37D0/48?

A. 10.176.15.131
B. 10.200.80.67
C. 208.48.107.192
D. 208.138.16.110

Answer: C

Explanation:
The IPv4 address of the IPv6 router is 208.48.107.192. In an automatic 6-to-4 tunnel, IPv6 addresses have the 2002::/16 prefix. The 32- bit IPv4 address of the IPv6 router is then embedded into the IPv6 address. The 32 bits of the IPv4 address is embedded in the second and third quartet of the IPv6 address. The second and third quarters in the IPv6 address correspond to D030:6BC0. The conversion of these hexadecimal digits into decimal is given as follows:

Hexadecimal Digits (in pairs)	Binary Equivalent	Decimal Equivalent
D0	11010000	208
30	00110000	48
6B	01101011	107
C0	11000000	192

The IPv6 router does not have 10.176.15.131 as its IPv4 address. The 10.176.15.131 address is the IPv4 equivalent of the second and third quarter (05B0:0F81) in the source IPv6 address. The other two IPv4 addresses are incorrect as they pertain to neither of the two IPv6 hosts. Objective:
Network Principles
Sub-Objective:
Recognize proposed changes to the network

References:
Cisco IOS IPv6 Implementation Guide > Implementing Tunneling for IPv6

QUESTION: 2
You have recently joined a company as the network administrator. You have been asked to complete the configuration on the border routers for an automatic 6-to-4 tunnel between several IPv6 network domains. The commands that are currently configured on the routers are as follows:
ipv6 route tunnel
interface tunnel ipv6 address tunnel source
Which of the following additional commands is required to complete the configuration of automatic 6-to-4 tunnel on the border routers?

A. tunnel mode ipv6ip
B. tunnel mode ipv6ip 6to4
C. tunnel mode ipv6ip auto-tunnel
D. tunnel mode ipv6ip isatap

Answer: B

Explanation:
The correct answer is to use the tunnel mode ipv6ip 6to4 command to complete the configuration of an automatic 6-to-4 tunnel. This command requires the use of IPv6 unicast addresses that have the 2002::/16 prefix.

- The types of tunneling mechanisms supported by IPv6 are: Automatic 6-to-4 tunnel
- ISATAP tunnel
-
- Manually configured tunnel GRE tunnel

Apart from using a tunneling mechanism, interoperability between IPv4 and IPv6 can be provided by using a dual-stack infrastructure or Network Address Translation-Protocol Translation (NAT-PT) . A dual-stack infrastructure allows you to use both IPv4 and IPv6 addresses on the same router/host. NAT-PT is used to translate IPv4 addresses to IPv6 and vice versa.
The tunnel mode ipv6ip command should not be used to complete the configuration because this command specifies IPv6 as the passenger protocol and creates a manually configured tunnel.
The tunnel mode ipv6ip auto-tunnel command is not required to enable automatic 6-to-4 tunneling on the border routers. This command creates an automatic IPv4-compatible IPv6 tunnel between the routers.
The tunnel mode ipv6ip isatap command should not be used because this command creates an ISATAP tunnel.
Objective: Network Principles Sub-Objective:
Recognize proposed changes to the network

References:
Cisco IOS IPv6 Configuration Guide; Implementing Tunneling for IPv6 > Configuring Manual IPv6 Tunnels Cisco > Cisco IOS IPv6 Command Reference > tunnel mode ipv6ip

QUESTION: 3
You have implemented IPv6 automatic 6-to-4 tunneling between three IPv6 subnets as shown in the network exhibit. (Click the Exhibit(s) button.)

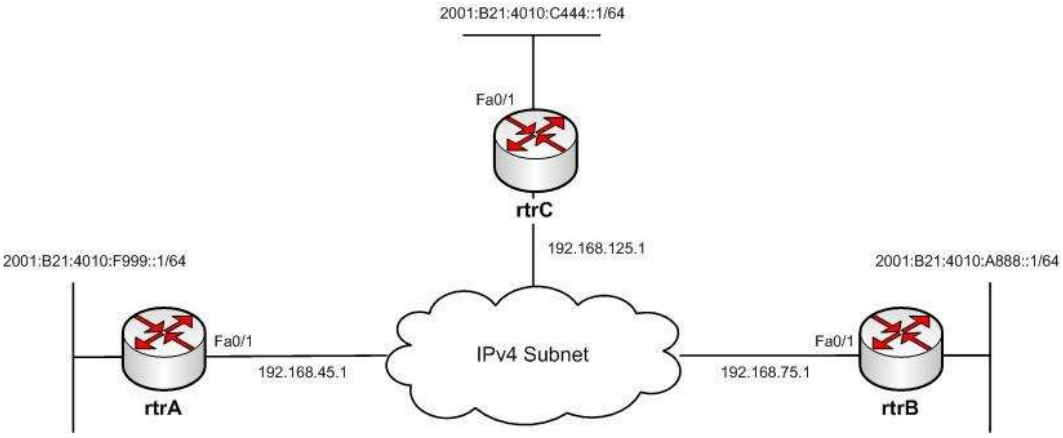

You have used the following commands to implement the automatic 6-to-4 tunnel:

```
rtrA(config)# interface Fa0/1
rtrA(config-if)# ip address 192.168.45.1 255.255.255.0
rtrA(config-if)# exit
rtrA(config)# interface Tunnel0
rtrA(config-if)# no ip address
rtrA(config-if)# tunnel mode ipv6ip 6to4
rtrA(config-if)# tunnel source Fa0/1
rtrA(config-if)# ipv6 address 2002:c0a8:2d01::1/64

rtrB(config)# interface Fa0/1
rtrB(config-if)# ip address 192.168.75.1 255.255.255.0
rtrB(config-if)# exit
rtrB(config)# interface Tunnel0
rtrB(config-if)# no ip address
rtrBconfig-if)# tunnel mode ipv6ip 6to4
rtrBconfig-if)# tunnel source Fa0/1
rtrBconfig-if)# ipv6 address 2002:c0a8:7d01::1/64

rtrC(config)# interface Fa0/1
rtrC(config-if)# ip address 192.168.125.1 255.255.255.0
rtrC(config-if)# exit
rtrC(config)# interface Tunnel0
rtrC(config-if)# no ip address
rtrC(config-if)# tunnel mode ipv6ip 6to4
rtrC(config-if)# tunnel source Fa0/1
rtrC(config-if)# ipv6 address 2002:c0a8:4b01::1/64
```

Your supervisor has assigned the task of verifying the automatic 6-to-4 tunnel to one of

your colleagues. Your colleague runs the show running-config command and finds that incorrect IPv6 addresses have been assigned to the tunnel interfaces of the routers. Which of the following IPv6 addresses should be assigned to rectify the problem? (Choose two.)

A. 2002::c0a8:2d01/64 to the Fa0/1 interface of rtrA
B. 2002:c0a8:4b01::1/64 to the Fa0/1 interface of rtrB
C. 2002:c0a8:7d01::1/64 to the Fa0/1 interface of rtrC
D. 2002:c0a8:4b01::1/64 to the Fa0/1 interface of rtrA

Answer: B, C

Explanation:
The 2002:c0a8:4b01::1/64 and the 2002:c0a8:7d01::1/64 IPv6 addresses should be assigned to the Fa0/1 interfaces of rtrB and rtrC, respectively. Automatic 6-to-4 tunnels embed the IPv4 address of the tunnel interfaces into the second and third quartets of the IPv6 address that has the 2002::/16 prefix.
To assign IPv6 addresses to the tunnel interfaces, perform the following steps:
1. Convert the IPv4 address of the tunnel interface into binary.
2. Convert the binary equivalent of the IPv4 address into hexadecimal (IPv6).
3. Append the hexadecimal equivalent to the 2002::/16 prefix to form the IPv6 prefix of the tunnel interface.
For the Fa0/1 interface of rtrB, its IPv4 address of 192.68.75.1 is equivalent to the IPv6 address c0a8:4b01. This address is then appended to the 2002::/16 prefix, resulting in 2002:c0a8:4b01::/48. The remaining host bits can be filled with zeros. Similarly, the IPv4 address of the Fa0/1 interface of rtrC is converted to the IPv6 address 2002:c0a8:7d01::/48.
The 2002::c0a8:2d01/64 IPv6 address should not be assigned to the Fa0/1 interface of rtrA. The Fa0/1 interface of rtrA has the IPv4 address 192.168.45.1. The IPv6 equivalent of the IPv4 address, which is c0a8:2d01, should be embedded in the second and third quartets of the IPv6 address instead of the seventh and eighth quartets. IPv4 addresses are embedded into the last 32 bits for ISATAP tunnels.
The 2002:c0a8:4b01::1/64 IPv6 addresses should not be assigned to the Fa0/1 interface of rtrA. This IPv6 address is the equivalent of the IPv4 address 192.168.75.1, which is the address of the Fa0/2 interface of rtrB and not rtrA. Therefore, this IPv6 address should be assigned to the Fa0/1 interface of rtrB.
Objective: Network Principles Sub-Objective:
Recognize proposed changes to the network

References:
Cisco Press > Articles > Cisco Certification > CCNP > CCNP Self-Study: Advanced IP Addressing
Cisco Press > Articles > Network Technology > General Networking > Cisco Self-

Study: Implementing Cisco IPv6 Networks (IPV6)
Cisco > Support > Technology Support > IP > IP Version 6 (IPV6) > Configure > Configuration Examples and Technotes > IPv6 Tunnel Through an IPv4 Network
Cisco IOS IPv6 Implementation Guide, Release 15.2M&T > Implementing Tunneling for IPv6

QUESTION: 4
An automatic IPv4-compatible IPv6 tunnel exists between two IPv6 networks. The two IPv6 networks belong to different BGP autonomous systems (AS). The tunnel source has the IPv4 address 172.168.111.65/24 and the tunnel destination has the IPv4 address 172.168.222.80/24. Which of the following statements is TRUE about the tunnel source and tunnel destination IPv6 addresses? (Choose two.)

A. the IPv6 address of the tunnel source is 172.168.111.65::
B. the IPv6 address of the tunnel source is ::172.168.111.65
C. the IPv6 address of the tunnel destination is 172.168.222.80::
D. the IPv6 address of the tunnel destination is ::172.168.222.80

Answer: B, D

Explanation:
The IPv6 address of the tunnel source is ::172.168.111.65 and the IPv6 address of the tunnel destination is
::172.168.222.80. These two addresses are IPv4-compatible IPv6 addresses, which are addresses that contain the IPv4 addresses of the tunnel source and destination.
In automatic IPv4-compatible IPv6 tunnel, the IPv4 addresses of the tunnel source and the tunnel destination are used to determine their IPv6 addresses. The IPv4 addresses of the tunnel source/destination are embedded into the least significant 32 bits of an all-zero unicast IPv6 address. The resultant IPv6 address has zeros in the most significant 96 bits and the IPv4 address of the tunnel source/destination in the remaining 32 bits. In this case, the source of an automatic IPv4-compatible IPv6 tunnel has the IPv6 address 0:0:0:0:0:0:172.168.111.65, abbreviated as ::2.168.111.65. You can also convert this address into pure hexadecimal format, which would be ACA8:6F41.
Any of the following three addresses could be used to identify the BGP neighbor at 172.168.11.65: 0:0:0:0:0:0:172.168.111.65
::172.168.111.65
::ACA8:6F41
Similarly, the tunnel destination has the IPv6 address 0:0:0:0:0:0:172.168.222.80 (abbreviated as ::172.168.222.80). The hexadecimal form of the IPv6 address of the tunnel destination is ::ACA8:DE50. Any of the following three addresses could be used to identify the BGP neighbor at 172.168.222.80:

0:0:0:0:0:0:172.168.222.80
::172.168.222.80
::ACA8:DE50
The other two options state incorrect IPv6 addresses of the tunnel source and the tunnel destination. Both options specify an IPv6 address that has the IPv4 address of the tunnel source/destination in the most significant 32 bits and zeros in the least significant 96 bits.
Objective: Network Principles Sub-Objective:
Recognize proposed changes to the network

References:
Home > Support > Technology Support > IP > IP Version 6 (IPv6) > Configure > Configuration Examples and Technotes > IPv6 Tunnel Through an IPv4 Network > Configure > Configurations (Automatic IPv4-Compatible Mode)
Cisco IOS IPv6 Implementation Guide > Implementing Tunneling for IPv6 Cisco > Support > Technology Support > IP > IP Version 6 (IPv6) > Technology Information > Technology White Paper > IPv6 Deployment Strategies > Selecting a Deployment Strategy > Deploying IPv6 Over IPv4 Tunnels > Automatic IPv4-Compatible Tunnel

QUESTION: 5
Your company has implemented IPv6 addresses and routing on every host, server, and router. Recently, your company acquired another company that has an IPv4 addressing scheme for its entire network. The acquired company's network does not have any support for IPv6. You need to devise a method so that the IPv6 hosts in your company can seamlessly communicate with the IPv4 hosts of the acquired company's network. You do not want to install any additional routers, and you want minimum configuration changes on the networks. Which of the following is the best method to allow communication between the IPv4 and IPv6 hosts?

A. Embedding IPv6 packets into IPv4 packets
B. Translating IPv4 addresses to and from IPv6 addresses
C. Configuring IPv6 on the hosts and routers in the IPv4 network
D. ConfiguringIPv4 on the hosts and routers in the IPv6 network

Answer: B

Explanation:
Translating IPv4 addresses to and from IPv6 addresses is the best method to allow communication between the IPv4 and IPv6 hosts. This translation of IPv4 and IPv6 addresses is known as Network Address Translation-Protocol Translation (NAT-PT) . NAT-PT is a technique available for deploying IPv6 and IPv4 addresses in a unified

network. With NAT- PT, the network requires fewer modifications and software for the IPv4 and IPv6 hosts. Additionally, it provides easy and quick interoperability between the IPv4 and IPv6 hosts.
NAT-PT is configured on one of the routers on the border of the IPv4 and IPv6 networks. Whenever an IPv4 packet intended for a host in the IPv6 network is received by the NAT-PT router, the router applies NAT-PT on the packet and translates all the headers in the IPv4 headers. In addition, it translates the IPv4 source and destination addresses to IPv6 source and destination addresses. The IPv6 packet is then set by the NAT-PT router to the intended IPv6 host. The NAT-PT router performs the reverse translation when an IPv6 host sends a packet to an IPv4 host.
Embedding IPv6 packets into IPv4 packets is not the best method to allow communication between the IPv4 and IPv6 hosts. When IPv6 packets are embedded inside IPv4 packets, the process is referred to as tunneling. Tunneling is appropriate when two IPv6 networks are separated by an IPv4 network. When an IPv6 host of one network sends an IPv6 packet destined for a host on the other IPv6 network, an IPv4 tunnel is created between the two IPv6 networks. The IPv6 packet is then embedded into an IPv4 packet that traverses through the IPv4 tunnel to reach the intended IPv6 host, where the embedded packet is extracted by the recipient. In this scenario, a single IPv6 network is available; hence, a tunnel cannot be formed.
Configuring IPv6 on the hosts and routers in the IPv4 network, or configuring IPv4 on the hosts and routers in the IPv6 network, are not the best methods to allow communication between the IPv4 and IPv6 hosts. Each of these two methods is cumbersome and not the most efficient for providing interoperability between IPv4 and IPv6 in this case. Furthermore, the IPv4 hosts on the acquired company's network do not support IPv6 as stated.
Objective: Network Principles Sub-Objective:
Recognize proposed changes to the network

References:
Cisco NAT Configuration Guide, Release 15M&T > NAT-PT for IPv6

QUESTION: 6
You have implemented an automatic 6-to-4 tunnel between the routers rtrA and rtrB as shown in the following network diagram:

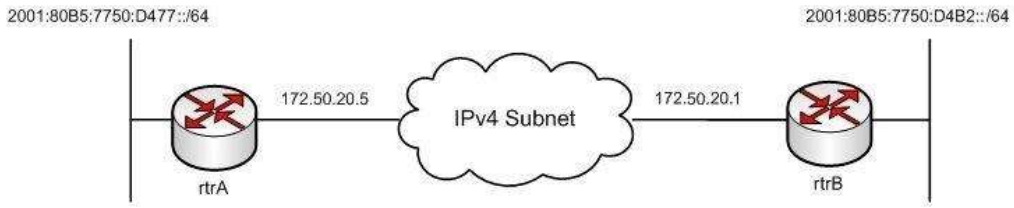

The routers rtrA and rtrB are connected to two IPv6 subnets and are separated by an IPv4 network. You decide to verify whether the tunnel was correctly implemented using

the show running-config command. Which of the following commands should exist in the output of the show running-config command on rtrA and rtrB? (Choose all that apply.)

A. interface tunnel
B. tunnel source
C. tunnel destination
D. tunnel mode ipv6ip
E. tunnel mode ipv6ip 6to4

Answer: A, B, E

Explanation:
The following commands should exist in the output of the show running-config command on rtrA and rtrB:
interface tunnel tunnel source
tunnel mode ipv6ip 6to4
The interface tunnel command is used to define a tunnel interface on the router. The tunnel source command allows you to specify the source of the tunnel, which is the router interface that faces the IPv4 network. The tunnel source must be configured with an IPv4 address. The tunnel mode ipv6ip 6to4 command is used to specify the tunneling mechanism, which in this case is automatic 6-to-4.
The partial output of the show running-config command on rtrA is as follows:
!
interface Tunnel0 no ip address
tunnel mode ipv6ip 6to4 tunnel source 172.50.20.5
ipv6 address 2002:ac32:of06::1/48 !

<output omitted>
The partial output of the show running-config command on rtrB is as follows:
!
interface Tunnel0 no ip address
tunnel mode ipv6ip 6to4 tunnel source 172.50.20.1
ipv6 address 2002:ac32:0f06::2/48 !

<output omitted>
The tunnel destination command and the tunnel mode ipv6ip commands do not appear in the show running- config output when automatic 6-to-4 tunnels are implemented on rtrA and rtrB. Both of these commands are executed for manually configured tunnels. Objective: Network Principles Sub-Objective: Recognize proposed changes to the network

References:

Cisco Press > Articles > Cisco Certification > CCNP > CCNP Self-Study: Advanced IP Addressing Cisco Interface and Hardware Component Configuration Guide > IPv6 Automatic 6to4 Tunnels
Cisco > Support > Technology Support > IP > IP Version 6 (IPV6) > Configure > Configuration Examples and Technotes > IPv6 Tunnel Through an IPv4 Network
Cisco IOS IPv6 Implementation Guide > Implementing Tunneling for IPv6

QUESTION: 7
Which of the following statements are TRUE about manually configured IPV4-to-IP6 tunnels and GRE tunnels? (Choose two.)

A. Manually configured tunnels use the tunnel mode ipv6ip command, while GRE tunnels use the tunnel mode gre ip command.
B. Manually configured tunnels support IPv6 IGPs, while GRE tunnels do not.
C. Manually configured tunnels block IPv6 multicasts, while GRE forwards them.
D. Manually configured tunnels do not support multiple passenger protocols, while GRE tunnels support them.

Answer: A, D

Explanation:
The following statements are TRUE about manually configured tunnels and GRE tunnels:

- Manually configured tunnels use the tunnel mode ipv6ip command, while GRE tunnels use the tunnel mode gre ip command.
- Manually configured tunnels do not support multiple passenger protocols, while GRE tunnels support them.

Manually configured tunnels and Generic Routing Encapsulation (GRE) tunnels are static point -to-point tunneling methods. Both of these tunneling methods provide a permanent link between two IPv6 networks that are separated by an IPv4 backbone. For each link between two IPv6 networks, a separate tunnel needs to be created.
Manually configured tunnels use a particular passenger protocol and do not support multiple passenger protocols at the same time. However, GRE tunnels can simultaneously use various passenger protocols.
It is incorrect to state that manually configured tunnels support IPv6 IGPs, while GRE tunnels do not. GRE tunnels also support IPv6 IGPs, such as OSPF, RIP, and IS-IS. It is incorrect to state that manually configured tunnels block IPv6 multicasts, while GRE forwards them. Manually configured tunnels also forward IPv6 multicasts.
Objective: Network Principles Sub-Objective: Recognize proposed changes to the network

References:
Cisco IOS IPv6 Configuration Guide, Release 12.4 > Implementing Tunneling for IPv6 > Configuration Examples for Implementing Tunneling for IPv6 > Example: Configuring Manual IPv6 Tunnels

QUESTION: 8
Which dialer interface command sets the maximum size of IP packets to 1492?

A. router(config-if)# mtu 1492
B. router(config-if)# ip ppp 1492
C. router(config-if)# ip 1492
D. router(config-if)# ip mtu 1492

Answer: D

Explanation:
The correct interface command to set the maximum size of IP packets (maximum transmission unit or MTU size) to 1492 is router(config-if)# ip mtu 1492. This command is required because RFC 2516 states the maximum receive unit (MRU) must not be negotiated larger than 1492 bytes.
All other answers are invalid commands due to incorrect syntax. Objective:
Network Principles
Sub-Objective:
Explain TCP operations

References:
Cisco > Cisco IOS IP Application Services Command Reference > idle (firewall farm datagram protocol) through ip slb natpool > ip mtu

QUESTION: 9
Examine the following FIB table:

```
Prefix               Next Hop     Interface
0.0.0.0/0            drop         Null0 (default route handler entry)

0.0.0.0/8            drop
0.0.0.0/32           receive
127.0.0.0/8          drop
224.0.0.0/4          drop
224.0.0.0/24         receive
240.0.0.0/4          drop
255.255.255.255/32   receive
```

Which of the following statements is NOT true?

A. These are the default entries in an FIB table
B. No IP addresses have been configured on this router
C. Multicast routing is enabled
D. The gateway of last resort has not been set

Answer: C

Explanation:
The Forwarding Information Base (FIB) table is created when Cisco Express Forwarding (CEF) is enabled on the router. FIB is a mapping of destination networks and IP addresses to next-hop IP addresses and exit interfaces.
In the scenario, multicast routing has NOT enabled in the router. If it were enabled, the next hop for the 224.0.0.0/4 network would not be listed as drop. A drop means any packets sent to multicast IP addresses will be dropped. If multicast routing were enabled, the entry for 224.0.0.0 would appear as follows:

```
Prefix              Next Hop              Interface
224.0.0.0/4         0.0.0.0
```

The next hop of 0.0.0.0 means that this traffic will be process switched, and CEF cannot forward the packets. The table displayed in the scenario contains the default entries in the FIB. These entries will change based on further configuration of the router interfaces and the addition of routes to the routing table through either static routing or through routing protocols.
No IP addresses have been configured on the router. Had they been configured, the addresses of the networks to which they were connected would be in the table. For example, if the IP address of the FastEthernet 0/1 interface were set to 192.168.1.1/24, three entries would have been added to the table as follows:

```
Prefix                  Next Hop        Interface
192.168.1.0/24          attached        FastEthernet0/1
192.168.1.0/32          receive
192.168.1.1/32          receive
192.168.1.255/32        receive
```

While the first IP address represents the directly attached network of which the interface is a member, the second IP address represents the network ID of the network, the third IP address represents the specific IP address assigned to the interface, and the last IP address represents the broadcast address of the network.
The gateway of last resort has not been set on the router. If it were set, it would be listed along with an IP address for the next hop and the exit interface. An entry for a gateway of last resort (or default route) would resemble the following:

```
Prefix          Next Hop        Interface
0.0.0.0/0       192.168.5.5     FastEthernet 0/0
```

Objective: Network Principles Sub-Objective:
Identify Cisco Express Forwarding concepts

References:
Cisco IOS Switching Services Configuration Guide, Release 12.2 > Cisco Express Forwarding Overview Cisco > Home > Support > Product Support > Routers > Cisco 12000 Series Routers > Troubleshoot and Alerts > Troubleshooting Technotes > Understanding Cisco Express Forwarding (CEF) https://www.ccexpert.us/traffic-share/fib-entries.html

QUESTION: 10
Which of the following IPv6/IPv4 interoperability techniques routes both IP versions simultaneously?

A. NAT-PT
B. Dual stack
C. 6to4 tunnels
D. Teredo

Answer: B

Explanation:
When the routers in the network are capable of routing both IPv6 and IPv4 traffic, it is referred to as dual stack. The dual stack routers simply recognize the version a frame is using and react accordingly to each frame.
Network Address Translation- Port Translation (NAT-PT) is a service that runs on a router or server that converts IPv4 traffic to IPv6, and vice versa. This eliminates the need for the routers or clients to be dual stack- capable. When only one router exists between the IPv4 and the IPv6 networks, this will be the only option, since all other methods listed require a dual stack capable device on each end of the tunnel. The IPv6 to IPv4 mapping can be obtained by the host from a DNS server, or the mapping can be statically defined on the NAT device.
6to4 tunnels can be created between dual stack routers or between a dual stack router and a dual stack client. In either case, each tunnel endpoint will have both an IPv6 and an IPv4 address. When traffic needs to cross an area where IPv6 is not supported, the tunnel can be used to transport the IPv6 packet within an IPv4 frame. When the frame reaches the end of the tunnel, the IPv4 header is removed and the IPv6 frame is further routed based on its IPv6 address.
Teredo is an alternate tunneling mechanism that encapsulates the IPv6 frame in an IPv4

UDP packet. It has the added benefit of traversing a NAT device that is converting private IP addresses to public IP addresses. 6to4 tunnels cannot traverse NAT devices by converting private IP addresses to public IP addresses. Objective: Network Principles Sub-Objective:
Recognize proposed changes to the network

References:
Cisco > Home > Products and Services > Cisco IOS and NX-OS Software > Cisco IOS Technologies > IPV6 > Product Literature > White Papers > Federal Agencies and the Transition to IPv6
Cisco > Cisco IOS IPv6 Configuration Guide, Release 15.2MT

QUESTION: 11
An enterprise has implemented an IPv4 addressing scheme on the servers of its core network. To effectively handle the increasing user requests to the server, the enterprise has plans to implement three new subnets with IPv6 addressing in its existing IPv4 network. The network administrator has set up dual-stack routers on the boundary of these subnets, as shown in the network diagram.

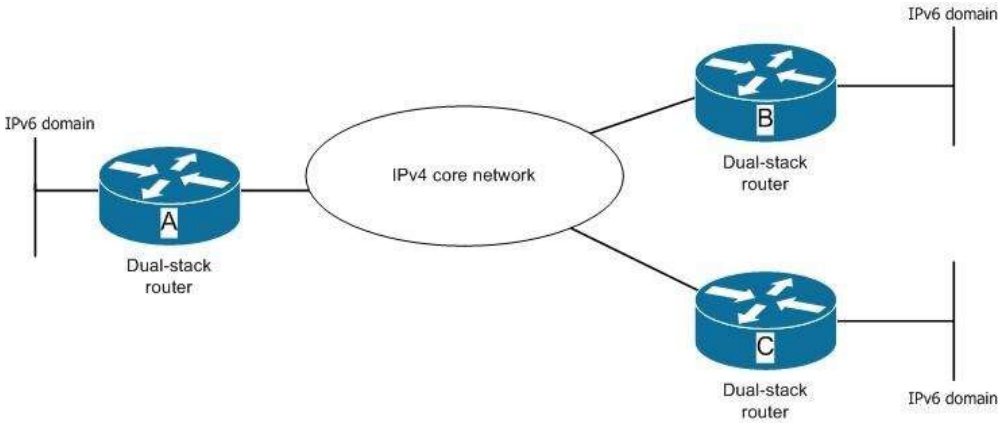

You need to ensure interoperability between IPv4 and IPv6 hosts such that routers A, B, and C can dynamically determine the destination of an IPv6 packet. In addition, global unicast addresses should be supported on all hosts in the three IPv6 subnets. Which of the following tunneling method can be used between the three routers? (Choose two.)

A. GRE tunnel
B. Automatic 6-to-4 tunnel
C. ISATAP tunnel
D. Manually-configured tunnel

Answer: B, C

Explanation:
You can use either automatic 6-to-4 tunnel or an Intra-site Automatic Tunnel Addressing Protocol (ISATAP) tunnel. Both of these tunneling methods are point-to-multipoint tunneling methods. This means that a single router (the point) can send IPv6 packets to different IPv6 routers (multipoints), depending on the destination address. When a router receives an IPv6 packet from an IPv6 host, it encapsulates the IPv6 packet in an IPv4 packet, which is then sent through the IPv4 core network. When the IPv4 packet is received at the destination router, the IPv6 address is extracted from the IPv4 packet and then forwarded to the intended IPv6 host.
Automatic 6-to-4 tunnels are created automatically by two IPv6 routers separated by an IPv4 network. These tunnels consider the IPv4 network as a virtual non-broadcast multi-access (NBMA) link. The tunnel is formed for every IPv6 packet that travels from one IPv6 border router to another IPv6 border router. IPv4 and IPv6 must be supported at both the border routers.
In automatic 6-to-4 tunneling, addresses belonging to the 2002::/16 prefix are used. In such IPv6 addresses, the 32-bit IPv4 address of each edge router is embedded into its IPv6 address increasing the length of the prefix to 48 (16 + 32). In automatic 6-to-4 tunnel, the IPv4 address of the router is embedded into the second the third quartet of the IPv6 address of the router.
ISATAP is also an automatic tunneling mechanism that uses an underlying IPv4 network as a NBMA link for IPv6 networks. However, it is most suitable for exchanging packets within an IPv6 network instead of exchanging packets between two IPv6 networks. With ISATAP tunnels, IPv6 dual-stack routers connected through the same IPv4 network can communicate with one another.
ISATAP works with unicast IPv6 addresses that are identified by a 64-bit prefix. The lowermost 64 bits are used to identify the interface of the router and are in modified EUI-64 format. The 0:5eFe value exists in the first 32 bits of the interface identifier. This value indicates that the IPv6 address is an ISATAP address. The remaining 32 bits contain the hexadecimal value of the IPv4 address; that is, the seventh and the eighth quartets in the IPv6 contain the IPv4 address.
You should not use a GRE tunnel or a manually configured tunnel between the three routers. These two tunneling methods provide static point-to-point tunnel between two IPv6 routers through an IPv4 network. Both these tunneling methods assume a virtual point-to-point link.
Objective: Network Principles Sub-Objective:
Recognize proposed changes to the network

References:
Cisco IPv6 Implementation Guide; Implementing Tunneling for IPv6

QUESTION: 12
Which of the following statements represent characteristics of an automatic 6to4 tunnel through an IPv4 network? (Choose all that apply.)

A. There is a NAT-PT router on either end of the tunnel.
B. There is a dual stack router on either end of the tunnel.
C. Each 6to4 site will have a /48 prefix.
D. Each 6to4 site will have a /16 prefix.
E. The IPv4 addresses of the edge routers are part of the site prefix.
F. The IPv6 addresses of the sending and receiving IPv6 hosts are part of the site prefix.

Answer: B, C, E

Explanation:
When implementing an automatic 6to4 tunnel, each IPv6 site receives a 48-bit prefix. The hexadecimal equivalent of the IPv4 address of the edge router is appended to 0x2002 and followed with the prefix to identify each end of the tunnel. Each end of the tunnel must be a dual stack router, which is one that can route both IPv4 and IPv6 traffic. For example, if the edge router's IPv4 address were 192.168.99.1, the hexadecimal equivalent of the address (c0a8:6301) would be inserted between 0X2002 and the /48 prefix, resulting in a packet with the IPv6 address 2002:c0a8:6301::/48 to arrive at the tunnel endpoint address.
A Network Address Translation - Port Translation (NAT-PT) router performs translation from IPv4 to IPv6. It is not used in a 6to4 tunnel.
Each site does not have a /16 prefix with a 6to4 tunnel. Rather, each site has a /48 prefix.
The IPv6 address of each IPv6 host is not part of the site prefix. These addresses are retained within the IPv6 portion of the header, and will be read after the frame reaches the end of the tunnel for eventual IPv6 routing on the far end. Objective: Network Principles Sub-Objective:
Recognize proposed changes to the network

References:
Cisco > Products > Collateral > Whitepaper > Enterprise IPv6 Transition Strategy > IPv6 Deployment Solution Options

QUESTION: 13
Examine the following output.

```
Router#show adjacency
Protocol     Interface     Address
IP           Serial0       10.10.10.2(2)  (incomplete)
<output omitted>
```

What possible reason(s) can cause the state of the first entry in the adjacency table? (Choose all that apply.)

A. the interface is a multipoint interface
B. the clear ip arp command was executed
C. the Layer 3 information is unknown
D. the clear adjacency command was executed

Answer: B, D

Explanation:
If either the clear ip arp or the clear adjacency commands were issued, the entry would temporarily be listed as incomplete in the adjacency table. The adjacency table is used by Cisco Express Forwarding (CEF) to maintain Layer 2 information about the next hop to remote networks. In CEF, an adjacency refers to a control structure that holds Layer 2 information for an IP address on a particular interface. When that information is not available the entry will be listed as incomplete, as shown in the example. Layer 2 information normally comes from the ARP process. Therefore, if the ARP table is cleared with the clear ip arp command, the Layer 2 information will be temporarily unavailable until the ARP process re-learns it the next time a frame must be sent to that hop. Moreover, if the adjacency table is emptied with the clear adjacency command, the entry must be created again. This will also result in the entry being marked incomplete for a short period of time until the ARP table can be consulted and the Layer 2 information re- added.
The interface in the scenario is not a multipoint interface. A multipoint interface would include entries for multiple next hops, since a multipoint interface connects to multiple Layer 3 destinations. An example of this is shown below in sample output from a Frame Relay interface:

```
Protocol    Interface              Address
IP          Serial0             140.108.1.1(25)
                              0 packets, 0 bytes
                              18410800
                              FR-MAP     never
                              Epoch: 1
IP          Serial0             140.108.1.2(5)
                              0 packets, 0 bytes
                              18510800
                              FR-MAP     never
                              Epoch: 1
```

The layer 3 information of the next hop is present in the entry in the scenario example. It is 10.10.10.2.
Objective: Network Principles Sub-Objective:
Identify Cisco Express Forwarding concepts

References:
Home > Support > Technology support > IP > IP switching > Troubleshoot and alerts > Troubleshooting Technotes > Troubleshooting Incomplete Adjacencies with CEF

QUESTION: 14
You have been alerted that TCP traffic leaving an interface has been reduced to near zero, while UDP traffic is steadily increasing at the same time. What is this behavior called and what causes it?

A. jitter, caused by lack of QoS
B. latency, caused by the MTU
C. starvation, caused improper configuration of QoS queues
D. windowing, caused by network congestion

Answer: C

Explanation:
This behavior is called starvation and is caused by improper configuration of QoS queues. When TCP and UDP flows are assigned to the same QoS queue, they compete with one another. This is not a fair competition because the TCP packets will react to packet drops by throttling back TCP traffic, while UDP packets are oblivious to drops and will take up the slack created by the diminishing TCP traffic. The results from mixing UDP and TCP traffic in the same queue are:

- Starvation Latency
- Lower throughput

While it is true that jitter can be caused by a lack of QoS, jitter is not what is being described in the scenario. Jitter is the variation in latency as measured in the variability over time of the packet latency across a network. This phenomenon seriously impacts time-sensitive traffic, such as VoIP, and can be prevented by placing this traffic in a high-priority QoS queue.
While latency can be caused by the maximum transmission unit (MTU) in the network, this is not a case of latency, although latency may be one of the perceived effects of starvation. Latency is the delay in reception of packets. The MTU is the largest packet size allowed to be transmitted, and an MTU that is set too large can result in latency.
While windowing can be caused by network congestion, this is not a case of windowing. This is a technique used to adjust the number of packets that can acknowledged at once by a receiving computer in a transmission. In times of congestion the window, or number of packets that can be acknowledged at a time, will be small. Later, when congestion goes down, the window size can be increased. Objective: Network Principles
Sub-Objective:
Describe UDP operations

References:

Design Guide > Service Provider Quality of Service > CE Guidelines for Collapsing Enterprise Classes > Mixing TCP with UDP

QUESTION: 15
Refer to the following set of commands:

```
rtrA(config)# ipv6 unicast-routing
rtrA(config)# interface Fa0/0
rtrA(config-if)# ipv6 enable
rtrA(config-if)# ipv6 address 2001:0:1:1:D52::F3C/64
rtrA(config-if)# ip address 130.11.6.1 255.255.255.0
```

Which of the following statements is TRUE about the given set of commands?

A. IPv4 and IPv6 are running simultaneously on rtrA
B. The IPv4 address is translated to an IPv6 address
C. The IPv6 address is an IPv4-compatible address
D. A tunnel is created for the interoperability of the IPv4 and IPv6 addresses

Answer: A

Explanation:
The correct answer is that IPv4 and IPv6 are running simultaneously on rtrA. The set of commands enables IPv6 on the rtrA router and assigns an IPv4 address and an IPv6 address to the Fa0/0 interface. This indicates that the router is a dual-stack router on which both IPv4 and IPv6 are running simultaneously. The IPv4 address is not translated to the IPv6 address by the given set of commands because NAT-PT is not enabled on the router. To enable NAT-PT on a router, you need to use the ipv6 nat command. In addition, the ipv6 nat prefix command should be used to specify an IPv6 prefix.
The IPv6 address is not an IPv4-compatible address. IPv4-compatible IPv6 addresses are used in automatic IPv4-compatible IPv6 tunnels. These addresses refer to those IPv6 unicast addresses that have zeros in the first 96 bits and an IPv4 address in the last 32 bits. For example, 0:0:0:0:0:0:192.156.10.67 is an IPv4- compatible IPv6 address where 192.156.10.67 is an IPv4 address. The IPv6 address (2001:0:1:1:D52::F3C/64), in this case, is not an IPv4-compatible IPv6 address.
A tunnel is not created for the interoperability of the IPv4 and IPv6 addresses because the given set of commands configures the router as a dual-stack router. There are no commands for configuring a tunnel on the router. Objective: Network Principles Sub-Objective:
Recognize proposed changes to the network

References:

Cisco IOS IPv6 Configuration Guide, Release 12.4 > Implementing IPv6 Addressing and Basic Connectivity > Configuration Examples for Implementing IPv6 Addressing and Basic Connectivity > Example: Dual Protocol Stacks Configuration

QUESTION: 16
Which of the following statements is TRUE concerning a 6to4 tunnel?

A. The IPv6 packet is encapsulated in an IPv4 packet using an IPv4 protocol type of 41.
B. The 6to4 tunnel method includes a 20-byte IPv6 header with no options and an IPv4 payload.
C. The maximum transmission unit is increased by 20 octets with the 6to4 tunnel method.
D. The IPv6 packet has its header removed and replaced with an IPv4 header with the 6to4 tunnel method.

Answer: A

Explanation:
When an IPv6 packet is tunneled across a portion of the network that does not support IPv6, the IPv6 packet is encapsulated in an IPv4 packet using an IPv4 protocol type of 41. When it reaches the other end of the tunnel, the IPv4 portion is stripped off and the packet is routed the rest of the way by using the remaining IPv6 header.
This method does not include a 20-byte IPv6 header with no options and an IPv4 payload. On the contrary, it includes a 20-byte IPv4 header with no options and an IPv6 payload.
The maximum transmission unit is not increased by 20 octets with this method. Rather, it is decreased by 20 bytes due to the extra overhead.
The IPv6 packet does not have its header removed and replaced with an IPv4 header. It encapsulates the entire IPv6 packet within an IPv4 header.
Objective: Network Principles Sub-Objective:
Recognize proposed changes to the network

References:
Cisco > Home > Support > Technology Support > IP > IP Version 6 > Configure > Configuration Examples and Technotes > Tunneling IPv6 through an IPv4 Network

QUESTION: 17
Which of the following are valid IPv4 to IPv6 migration strategies? (Choose two.)

A. DHCP
B. Tunnels
C. Dual-stack
D. Encapsulating IPv4 into IPv6

Answer: B, C

Explanation:
Tunnels and dual-stack are valid IPv4 to IPv6 migration strategies.
Tunneling mechanisms can transport IPv6 across an IPv4 infrastructure. Cisco supports the following types of tunneling for this purpose:

- Manual tunnels
- Generic Routing Encapsulation (GRE) tunnels IPv4 compatible tunnels
- 6-to-4 tunnels
- Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) tunnels

For all tunneling types, IPv6 packets are encapsulated in IPv4 packets for delivery across the IPv4 infrastructure. These tunnels require two endpoints, either two routers, or a router and a host. Both endpoints must support IPV4 and IPv6. When implementing an automatic 6-to-4 tunnel each IPv6 site receives a /48-bit prefix. The hexadecimal equivalent of the IPv4 address of the edge router is appended to 0x2002 and followed with the prefix to identify each end of the tunnel. Each end of the tunnel must be a dual stack router, that is, one that can route both IPv4 and IPv6. For example if the edge router's IPv4 address were 192.168.99.1, the hexadecimal equivalent of the address (c0a8:6301) would be inserted between 0X2002 and the /48 prefix, resulting in 2002:c0a8:6301:: /48 to arrive at the tunnel endpoint address.
The following example shows a partial output of the show run command executed on a router hosting one end of a 6-to-4 tunnel:

```
router5# show run
!
interface loopback0
     ip address 64.101.64.1 255.255.255.0
!
interface Tunnel0
     ipv6 unnumbered Ethernet0/1
     tunnel source Loopback0
     tunnel source ipv6ip 6to4
!
interface Ethernet0/1
     ipv6 address 2002:4065:4001:1::/64 eui-64
!
ipv6 route 2002::/16 Tunnel0
```

The least significant 32 bits in the address referenced by the ipv6 route 2002::/16 Tunnel0 command correspond to the IPv4 address (64.101.64.1) assigned to the tunnel source. The hex equivalent is 4065:4001, yielding 2002:4065:4001::/48.
Another example of how IPv4 addresses can be used in the creation of the tunnel endpoint IPv6 identifier is shown in the partial output of the show run command executed on a router that is hosting one end of an automatic IPv4 compatible tunnel:

```
<output omitted>
interface Tunnel0
no ip address
no ip redirects
tunnel source Serial0/0
tunnel mode ipv6ip auto-tunnel
!
router bgp 100
no synchronization
no bgp default ipv4-unicast
bgp log-neighbor-changes
neighbor ::192.168.4.1 remote-as 100
no auto-summary
!
```

In the neighbor statement under the BGP configuration section, the neighbor address is derived from the IPv4 address of the other router (192.168.4.1). This could be implemented in one of three ways:

- ::192.168.4.1
- 0: 0:0:0:0:0:192.168.4.1
- ::c0a8:0401

The IPv6 addresses ::192.168.4.1 and 0:0:0:0:0:0:192.168.4.1 are implemented by inserting the IP address at either the end of :: or 0:0:0:0:0:0. (:: is a IPv6 shortcut for 0:0:0:0:0:0). The IPv6 address::c0a8:0401 is implemented by inserting the hex equivalent of 192.168.4.1 (c0a8:0401) in the same location.
Another potential migration strategy is to run dual stacks. The TCP/IP stack, or stack, is the TCP/IP software that is included in most operating systems. It is possible to run dual TCP/IP stacks on a computer. For example, servers and other infrastructure equipment often run both an IPv4 and IPv6 IP stack for application compatibility. This dual-stack configuration allows applications that require IPv6 to use the IPv6 stack and applications that require IPv4 to use the IPv4 stack. The following partial output of the show run command shows the configuration of a dual stack router: <output omitted>
ipv6 unicast routing
interface fastethernet0/0
ip address 192.168.5.1 255.255.255.0
ipv6 address 3ffe:b00:c19:2::3/127
This configuration allows applications on the same segment to communicate via IPv4 or

IPv6. Dynamic Host Configuration Protocol (DHCP) provides no benefits in migrating from IPv4 to IPv6. IPv4 is not encapsulated in IPv6 in any of the migration strategies. IPv6 is encapsulated into IPv4. Objective:
Network Principles
Sub-Objective:
Recognize proposed changes to the network

References:
Cisco > Cisco IOS IPv6 Implementation Guide, Release 12.4 > Implementing Tunneling for IPv6

QUESTION: 18
You just discovered that a ping packet sent from one of the devices to another took a different path in the return than it did on its way to the destination. What behavior caused this?

A. Windowing
B. Global synchronization
C. MSS
D. Asymmetric routing

Answer: D

Explanation:
This behavior is caused by asymmetric routing. This is quite common in a routed network and usually is not a problem. It can, however, become an issue when firewalls reside in a routed path. Firewalls can cause problems when they maintain state information about connections. State information is used to determine if return connection is allowed. If the return path is routed through a different firewall, it will not have the correct state information for the connection, and the return will be disallowed.
It is not caused by windowing. This is a technique used to adjust the number of packets that can be acknowledged at once by a receiving computer in a transmission. In times of congestion, the window or number of packets that can be acknowledged at a time will be small. Later, when congestion goes down, the window size can be increased.
The behavior is not caused by the maximum segment size (MSS). This value specifies the largest amount of data, in octets, that a computer or communications device can receive in a single TCP segment. This will not cause a packet to take a different path in the return than it did on its way to the destination.
The behavior is not caused by global synchronization. This occurs when congestion on the network causes all devices to reduce their transmission rates at the same time. The result is the network cycling between sharp increases and sharp decreases in traffic.

Objective: Network Principles Sub-Objective:
Explain TCP operations

References:
Home > Services > Technical services newsletter > Tech insights > Chalk talk > Asymmetric Routing and Firewalls

QUESTION: 19
You are configuring a 6to4 tunnel. You want to translate the IPv4 address 192.168.50.4 to the IPv6 address for the tunnel. What would be the correct translation?

A. 2002:c0a8:3204::/16
B. 2002:c0a8:9901::/48
C. 2002:c0a8:3204::/48
D. c0a8:3204:2002::/16

Answer: C

Explanation:
When implementing an automatic 6to4 tunnel, each IPv6 site receives a /48-bit prefix. The hexadecimal equivalent of the IPv4 address of the edge router is appended to 0x2002 and followed with the prefix to identify each end of the tunnel. In this case, if the edge router's IPv4 address were 192.168.50.4, the hexadecimal equivalent of the address (c0a8:3204) would be inserted between 0X2002 and the /48 prefix, resulting in 2002:c0a8:3204::/48 to arrive at the tunnel endpoint address. The correct address would not be 2002:c0a8:3204::/16. The prefix is 48, not 16. The correct address would not be 2002:c0a8:9901::/48. The hexadecimal equivalent of the address 192.168.50.4 is c0a8:3204, not c0a8:9901.
The correct address would not be c0a8:3204:2002::/16. It has an incorrect prefix (/16) and the values in the other sections are out of order. Objective: Network Principles Sub-Objective:
Recognize proposed changes to the network

References:
Cisco IPv6 Implementation Guide, Release 15.2M&T > Implementing Tunneling for IPv6 > Implementing Tunneling for IPv6 > Configuration Examples for Implementing Tunneling for IPv6

QUESTION: 20
In the Active Discovery phase of PPPoE, which of the following is NOT verified by the Broadband Network gateway (BNG) to prevent spoofing?

A. source MAC address
B. arriving access interface
C. PPPoE session ID
D. destination MAC address

Answer: D

Explanation:
The destination MAC address is the address of the BNG, so there is no need for it to be verified. If the traffic arrived on the BNG interface, it is correct. PPPoE is composed of two main phases, the Active Discovery Phase and the PPP Session Phase. The Active Discovery phase consists of the following communications between the PPPoE client and the BNG:
1. The client sends a PPPoE Active Discovery Initiation (PADI) broadcast signal to the remote device (BNG).
2. The remote device sends back a PPPoE Active Discovery Offer (PADO) that contains the MAC address of the BNG and destination MAC address of the subscriber (client).
3. The subscriber (client) send a PPPoE Active Discovery Request (PADR) continuing the destination MAC address of the BNG to which it wishes to establish a session.
4. The BNG responds with a PPPoE Active Discovery Session-Confirmation (PADS) containing the PPPoE session ID.
Once this process is complete, the session moves on to the PPP Session Phase in which Link Control Protocol (LCP) parameters such as maximum transmission unit (MTU) are agreed to, authentication is performed, and Network Control Protocols (NCP) for any Layer 3 protocol that will traverse the link are started. Objective:
Layer 2 Technologies Sub-
Objective: Configure and verify PPP

References:
Cisco Support Community > ASR9000 BNG debugging PPPoE sessions
Cisco > Cisco Security Appliance Command Line Configuration Guide, Version 8.0 > Configuring the PPPoE Client > PPPoE Client Overview