Interconnecting Cisco Networking Devices Part 2 (ICND2 v3.0)

Cisco 200-105 Dumps Available Here at:

https://www.certification-questions.com/cisco-exam/200-105-dumps.html

Enrolling now you will get access to 170 questions in a unique set of 200-105 dumps

Question 1

You manage the EIGRP subnet in your organization. You have enabled EIGRP for IPv6 on all the routers in the EIGRP AS 260 using the following commands on all the routers:
The ipv6 unicast-routing command in global configuration mode
The interface command in global configuration mode
The ipv6 enable command in interface configuration mode
The ipv6 eigrp command in interface configuration mode
The ipv6 router eigrp command in global configuration mode
The eigrp router-id command in global configuration mode
During verification, you discover that EIGRP for IPv6 is not running on the routers.
Which of the following should be done to fix the issue?

Options:

A. The ipv6 address command should be executed in interface configuration mode.

B. The ipv6 address command should be executed in router configuration mode.

C. The eigrp router-id command should be executed in interface configuration mode.

D. The eigrp router-id command should be executed in router configuration mode

Answer: D

Explanation:
The eigrp router-id command should be executed in router configuration mode to fix the issue. This command specifies a fixed router IPv4 address to the router. If this command is missing or incorrectly configured on the router, EIGRP for IPv6 will not run properly.
Another command that you should perform so that EIGRP for IPv6 runs on the routers is the no shutdown command. You should execute this command in interface configuration mode. The no shutdown command is necessary because all the interfaces with EIGRP for IPv6 enabled on them are in a shutdown state by

default.
A sample configuration to implement EIGRP for IPv6 on a router is as follows:
Rtr63(config)# ipv6 unicast-routing
Rtr63(config) # interface Fa0/1
Rtr63(config-if) # ipv6 enable
Rtr63(config-if) # ipv6 eigrp 260
Rtr63(config-if)# no shutdown
Rtr63(config-if) # exit
Rtr63(config)# ipv6 router eigrp 260
Rtr63(config-rtr)# eigrp router-id 1.1.1.1
The two options stating that the ipv6 address command should be executed on the routers are incorrect. EIGRP for IPv6 can be configured on router interfaces without explicitly specifying a global unicast IPv6 address. If you specify the ipv6 enable command, as in this scenario, then the IPv6 address command is not required.
The option stating that the eigrp router-id command should be executed in interface configuration mode is incorrect. This command should be executed in router configuration mode instead of interface or global configuration modes.
References:
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/configuration/15-2mt/ipv6-15-2mt-book/ip6-eigrp.html#GUID-0A728310-E5CB-4914-A657-BF1C0C656997

Question 2

You just finished configuring VLAN Trunking Protocol (VTP) in a network containing five switches. One of the switches is not receiving VLAN information from the switch that is acting as the server.
Which of the following could NOT be a reason why the switch is not receiving the information?

Options:

A. The VTP domain name on the switch may be misspelled

B. The VTP password may be misspelled on the switch

C. The configuration revision number may be out of sync

D. The VTP version used on the switch may be different

Answer: C

Explanation:
The configuration revision number does not need to match on the switches. The configuration number cannot be directly configured, but is instead synchronized during VTP updates.
For VTP to function correctly, all of the following conditions must be true:
The VTP version must be the same on all switches in a VTP domain.
The VTP password must be the same on all switches in a VTP domain.
The VTP domain name must be the same on all switches in a VTP domain.

References:
CCNA Routing and Switching Complete Study Guide: Exam 100-105, Exam 200-105, Exam 200-125, 2nd Edition, Chapter 2: LAN Switching Technologies - Configure, verify, and troubleshoot STP protocols

Question 3

Which of the following techniques is NOT used by distance vector protocols to stop routing loops in a network?

Options:

A. Split horizon

B. Spanning Tree Protocol (STP)

C. Holddowns

D. Route poisoning

Answer: B

Explanation:
Spanning Tree Protocol (STP) is not used by distance vector protocols to stop routing loops in a network. STP is used to prevent switching loops in a switched network.
Routing loops can occur due to slow convergence and inconsistent routing tables, and can cause excessive use of bandwidth or complete network failure. An example of a routing table problem would be incorrectly configured static default routes. Suppose that Router A is connected to Router B, and the addresses of the interfaces on each end of the link connecting the two routers are as follows:
Router A 192.168.5.1/24
Router B 192.168.5.2/24
A partial output of the routing tables of the two routers is shown below. Router B hosts the connection to the Internet.
routerA# show ip route
Gateway of last resort is 192.168.5.2 to network 0.0.0.0
<Output omitted>
routerB# show ip route
Gateway of last resort is 192.168.5.1 to network 0.0.0.0
<<output omitted>>
From the limited information shown above, you can see that Router A is pointing to Router B for the default route, and Router B is pointing to Router A for the default route. This will cause a routing loop for any traffic that is not in their routing tables. For example, if a ping were initiated to the address 103.5.6.8 and that address was not in the routing tables of Routers A and B, the most likely message received back would NOT be "destination unreachable" but "TTL expired in transit." This would be caused by the packet looping between the two routers until the TTL expired.
The following techniques are used by distance vector protocols to stop routing loops in a network:
Split horizon stops routing loops by preventing route update information from being sent back over the

same interface on which it arrived.
Holddown timers prevent regular update messages from reinstating a route that is unstable. The holddown timer places the route in a suspended, or "possibly down" state in the routing table and regular update messages regarding this route will be ignored until the timer expires.
Route poisoning "poisons" a failed route by increasing its cost to infinity (16 hops, if using RIP). Route poisoning is combined with triggered updates to ensure fast convergence in the event of a network change.
References:
http://www.ciscopress.com/articles/article.asp?p=24090&seqNum=3

Question 4

On which of the following networks will OSPF elect a designated router (DR)? (Choose two.)

Options:

A. Broadcast

B. NBMA

C. Point-to-point

D. Point-to-multipoint

Answer: A, B

Explanation:
OSPF will perform an election for a designated router (DR) and backup designated router (BDR) on every multi-access network segment. Multi-access segments are defined as segments where more than two hosts can reach each other directly, such as a shared Ethernet segment (broadcast multi-access) or Frame Relay (non-broadcast multi-access, or NBMA).
DR and BDR elections do not occur on point-to-point or point-to-multipoint segments. Point-to-point and point-to-multipoint segments are not considered multi-access segments. OSPF routers on these network types will establish an adjacency without a DR/BDR election.
References:
https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html#t21

Question 5

Examine the following output from SwitchD.

```
switch# show interfaces fastethernet0/1 switchport
Name: Fa0/1
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: Off
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
<<output omitted>>
```

Based on this output, what command MUST be executed for an 802.1q trunk to be created on port Fa0/1?

Options:

A. switchport mode trunk

B. switchport mode nonegotiate

C. switchport trunk encapsulation 802.1q

D. switchport trunk native VLAN

Answer: A

Explanation:
The command switchport mode trunk must be executed for a trunk to form. The output indicates that the Administrative Mode of the port is "static access," which means the port has been configured as a static (fixed) access port. Access mode disables trunking on an access port.
Below is a sample of the configuration required to allow a router to provide inter-VLAN routing between two VLANs residing on the switch:
Router(config)#interface fa0/0
Router(config)#no shut down
Router(config)#interface fa0/0.1
Router(config-subif)#encapsulation dot1q
Router(config-subif)#ip address 192.168.10.1 255.255.255.0
Router(config)#interface fa0/0.2
Router(config-subif)#encapsulation dot1q
Router(config-subif)#ip address 192.168.20.1 255.255.255.0
Switch(config)#interface fa0/1
Switch(config-if)#switchport mode trunk
For this example, the following statements are true:
- The trunk link connects to Fa0/0 on the router and Fa0/1 on the switch.
- The physical interface F0/0 on the router has been divided into two subinterfaces, Fa0/0.1 and Fa0/0.2.

- The encapsulation type of 802.1q has been specified on the two subinterfaces of the router.
- The physical interface on the switch has been specified as a trunk link.
- The IP addresses 192.168.10.1 and 192.168.20.1 should be the default gateways of the computers located in VLANs 1 and 2, respectively.

The switchport mode nonegotiate command does not need to be executed because the switch is already configured for non-negotiation, as indicated by the output Negotiation of Trunking: Off. Trunk negotiation using the Dynamic Trunking Protocol (DTP) does not need to be enabled for a trunk to form.
The switchport trunk encapsulation 802.1q command does not need to be executed for a trunk to form. Also, the output Operational Trunking Encapsulation: dot1q indicates that 802.1q encapsulation is already configured.
The switchport trunk native VLAN command does not need to be executed. This command is used to change the native VLAN from its default of 1, but leaving it set to the default of 1 will not prevent the trunk from forming.
References:
https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3560/software/release/12-2_25_see/configuration/guide/scg/swvlan.html#wp1096213

Question 6

You network team is exploring the use of switch staking. Which of the following statements is NOT true of switch stacking?

Options:

A. The master switch is the only switch with full access to the interconnect bandwidth

B. Switches are connected with special cable

C. The stack has a single IP address

D. Up to nine switches can be added to the stack

Answer: A

Explanation:
All switches in the stack have full access to the interconnect bandwidth, not just the master switch. The master switch is elected from one of the stack members. It automatically configures the stack with the currently running IOS image and a single configuration file.
The switches are connected with special cables that form a bidirectional closed loop path.
The stack has a single management IP address and is managed as a unit.
Up to nine switches can be in a stack.
References:
https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-3750-series-switches/prod_white_paper09186a00801b096a.html

Question 7

What two devices can be connected to a router WAN serial interface that can provide clocking? (Choose two.)

Options:

A. CSU/DSU

B. switch

C. modem

D. hub

Answer: A, C

Explanation:
A router DTE interface must receive a clock rate from the DCE end and the rate can be provided by either a CSU/DSU or a modem. Therefore, the connection between the local router and the service provider can be successfully completed by adding either of these devices between the service provider and the local router.
Switches and hubs are neither capable of providing the clock rate nor able to complete the connection between the local router and the service provider.
References:
http://docwiki.cisco.com/wiki/Internetworking_Technology_Handbook#WAN_Technologies

Question 8

You want to encrypt and transmit data between peer routers with high confidentiality. Which protocol option should you choose?

Options:

A. Authentication Header (AH) in tunnel mode

B. Authentication Header (AH) in transport mode

C. Encapsulating Security Payload (ESP) in tunnel mode

D. Encapsulating Security Payload (ESP) in transport mode

Answer: C

Explanation:
You should choose Encapsulating Security Payload (ESP) in tunnel mode to encrypt and transmit data between peer routers with high confidentiality. Two protocols can be used to build tunnels and protect data traveling across the tunnel:
Authentication Header (AH) uses protocol 51.

ESP uses protocol 50.
AH is defined in Request for Comments (RFC) 1826 and 2402. AH does not perform data encryption and therefore, information is passed as clear text. The purpose of AH is to provide data integrity and authentication, and anti-reply service (optional). It ensures that a packet that crosses the tunnel is the same packet that left the peer device and no changes have been made. It uses a keyed hash to accomplish this.
ESP is defined in RFC 2406. ESP can provide data integrity and authentication, but its primary purpose is to encrypt data crossing the tunnel. There are two reasons why ESP is the preferred building block of IPSec tunnels:
The authentication component of ESP does not include any Layer 3 information. Therefore, this component can work in conjunction with a network using Network Address Translation (NAT).
On Cisco devices, ESP supports encryption using Advanced Encryption Standard (AES), Data Encryption Standard (DES), or Triple DES (3DES).
Tunnel mode is used between Virtual Private Network (VPN) gateways such as routers, firewalls, and VPN concentrators.
Transport mode is used between end-stations or between an end-station and a VPN gateway.
The options AH in tunnel mode and AH in transport mode are incorrect because AH does not provide encryption.
The option ESP in transport mode is incorrect because transport mode is used between end-stations or between an end-stations and a VPN gateway.
References:
http://www.ciscopress.com/articles/article.asp?p=25477&rl=1
https://www.cisco.com/c/en/us/about/press/internet-protocol-journal/back-issues/table-contents-4/ipj-archive/article09186a00800c830b.html

Question 9

You execute the ping command from a host, but the router does not have a path to its destination.
Which of the following ICMP message types will a client receive from the router?

Options:

A. ICMP redirect

B. ICMP time exceeded

C. ICMP destination unreachable

D. ICMP echo-reply

Answer: C

Explanation:
When a router receives a ping packet and has no route to the destination in its routing table, it will respond to the client with an ICMP destination-unreachable message. Internet Control Message Protocol (ICMP) is a Layer 3 protocol used to test the connectivity between hosts in a network. There are six types of

unreachable destination message:
1. Network unreachable
2. Host unreachable
3. Protocol unreachable
4. Port unreachable
5. Fragmentation needed and Don't Fragment (DF) bit set
6. Source route failed
An ICMP redirect message would not be received. This type of response is received when the router is configured to direct clients to a different router for better routing.
An ICMP time-exceeded message would not be received. This type of response occurs when the router successfully sent the packet but did not receive an answer within the allotted time; in other words, the time-to-live of the ICMP packet has been exceeded.
An ICMP echo-reply message would not be received. This would be the response received if the destination received the ping command and responded successfully.
References:
http://docwiki.cisco.com/wiki/Internet_Protocols#Internet_Control_Message_Protocol_.28ICMP.29

Question 10

When executed on a HSRP group member named Router 10, what effect does the following command have?
Router10(config-if)# standby group 1 track serial0 25

Options:

A. It will cause the router to increase its HSRP priority by 25 if the Serial0 interface on the standby router

goes down

B. It will cause the router to shut down the Serial0 interface if 25 packets have been dropped

C. It will cause the router to notify Router 25 is serial 0 goes down

D. It will cause the router to decrement its HSRP priority by 25 if Serial 0 goes down

Answer: D

Explanation:
This command will cause the router to decrement its HSRP priority by 25 if Serial 0 goes down. Interface tracking can be configured in Hot Standby Routing Protocol (HSRP) groups to switch traffic to the standby router if an interface goes down on the active router. This is accomplished by having the active router track its interface. If that interface goes down, the router will decrement its HSRP priority by the value configured in the command. When properly configured, this will cause the standby router to have a higher HSRP priority, allowing it to become the active router and to begin serving traffic.
When the standby router in an HSRP group is not taking over the active role when the active router loses

its tracked interface, it is usually a misconfigured decrement value, such that the value does not lower the HSRP priority of the active router far enough for the standby to have a superior priority value.
The command will not cause the router to increase its HSRP priority by 25 if the Serial0 interface on the standby router goes down. HSRP routers track their own interfaces, not those of another router.
The command will not cause the router to shut down the Serial0 interface if 25 packets have been dropped. It will only do this if the link becomes unavailable.
The command will not cause the router to notify Router 25 is serial 0 goes down. The number 25 in the command is the decrement value, not the ID of another router.
References:
https://www.cisco.com/c/en/us/support/docs/ip/hot-standby-router-protocol-hsrp/13780-6.html
https://www.cisco.com/c/en/us/td/docs/ios/ipapp/command/reference/iap_s5.html#wp1156911

Question 11

Examine the output shown below:

```
R1#show ipv6 eigrp neighbors
IPv6-EIGRP neighbors for process 1
H   Address          Interface Hold Uptime    SRTT    RTO  Q  Seq
                               (sec)          (ms)        Cnt Num
0  Link-local add     Se0/0    13 15:17:58    44      264  0  12
    FE80::2
```

```
R2#show ipv6 eigrp neighbors
IPv6-EIGRP neighbors for process 1
H   Address          Interface  Hold Uptime    SRTT    RTO  Q  Seq
                                (sec)          (ms)        Cnt Num
0   Link-local add    Se0/0     14 16:32:05    30    300   0  12
    FE80::1
```

What is true of this configuration?

Options:

A. The link-local address of R1 is FE80::2

B. The link-local address of R1 is FE80::1

C. The area ID is 1

D. No adjacency has formed

Answer: B

Explanation:
The output shows that the link-local address of R1 is FE80::1. R1's link-local address appears in the output of R2 because the show ipv6 eigrp neighbors command displays information about the neighbor, not the local router.
The link-local address of R1 is not FE80::2. That is the link-local address of R2.
Because the area ID is not displayed in the output, we do not know its value. The only 1 in the output is the value representing the process ID of both routers, IPv6-EIGRP neighbors for process 1.
It is not true that no adjacency has formed. There is an adjacency present; if there were not, the two routers would not appear in each other's output of the show ipv6 eigrp neighbors command.
References:
https://www.cisco.com/c/en/us/td/docs/ios/ipv6/command/reference/ipv6_book/ipv6_13.html?bookSearch=true

Question 12

You apply the following commands to a router named R2:

```
R2(config)# interface Tunnel1
R2(config-if)# ip address 172.16.1.2 255.255.255.0
R2(config-if)# ip mtu 1400
R2(config-if)# ip tcp adjust-mss 1360
R2(config-if)# tunnel source 2.2.2.2
R2(config-if)# tunnel destination 1.1.1.1
```

Which statement is NOT true with regard to this configuration?

Options:

A. The physical IP address of R2 is 2.2.2.2

B. The connection will operate in IP mode

C. The configuration will increase packet fragmentation

D. The configuration alters the maximum segment size

Answer: C

Explanation:
The configuration will not increase packet fragmentation. Conversely, it will reduce it by lowering the maximum transmission unit to 1400 and the maximum segment size to 1360 bytes.
Most transport MTUs are 1500 bytes. Simply reducing the MTU will account for the extra overhead added by GRE. Setting the MTU to a value of 1400 is a common practice, and it will ensure unnecessary packet

fragmentation is kept to a minimum.
The other statements are true. The physical address of R2 is 2.2.2.2, while the tunnel interface address is 172.16.1.2.
Because you have not issued any command that changes the connection, it will operate in the default mode of IP.
The configuration does alter the maximum segment size with the ip tcp adjust-mss 1360 command.
References:
https://supportforums.cisco.com/t5/network-infrastructure-documents/how-to-configure-a-gre-tunnel/ta-p/3131970

Question 13

Which of the following is NOT a feature offered by Enhanced Interior Gateway Routing Protocol (EIGRP)?

Options:

A. variable length subnet masks (VLSM)

B. partial updates

C. neighbor discovery mechanism

D. multiple vendor compatibility

Answer: D

Explanation:
EIGRP is a Cisco-proprietary routing protocol, and does not support multiple vendor environments.
EIGRP is a classless routing protocol, and thus supports variable length subnet masks (VLSM).
EIGRP routers build a neighbor table in memory, and use a multicast-based neighbor discovery mechanism.
EIGRP routers send partial updates when there are network events.
The following are features offered by EIGRP:
Fast convergence
Partial updates
Neighbor discovery mechanism
VLSM
Route summarization
Scalability
References:
https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/13669-1.html

Question 14

Which of the following commands would instruct OSPF to advertise ONLY the 192.168.10.0/24 network in

Area 0?

Options:

A. Router(config)# router ospf 1

Router(config-router)# network 192.168.10.0 0.0.0.255 area 0

B. Router(config)# router ospf 1

Router(config-router)# network 192.168.11.0 0.0.0.255 area 0

C. Router(config)# router ospf 1

Router(config-router)# network 192.168.10.0 255.255.255.0 area 0

D. Router(config)# router ospf 1

Router(config-router)# network 192.168.10.0 0.0.255.255 area 0

Answer: A

Explanation:
The command Router(config-router)# network 192.168.10.0 0.0.0.255 area 0 would instruct OSPF to advertise the 192.168.10.0 network in Area 0. It is executed in OSPF process 1 configuration mode, as indicated by the prompt Router(config-router)#. This command correctly states the network as 192.168.10.0 and uses the proper wildcard mask of 0.0.0.255.
The command Router(config-router)# network 192.168.11.0 0.0.0.255 area 0 is incorrect because it advertises the 192.168.11.0/24 network instead of the 192.168.10.0/24 network.
The command Router(config-router)# network 192.168.10.0 255.255.255.0 area 0 is incorrect because it uses a regular mask instead of a wildcard mask.
The wildcard mask in OSPF network statements must be expressed inversely, and not as a regular subnet mask. If the network you are configuring for OSPF operation is 192.168.10.0/24, then the inverse version of a /24 mask (or 255.255.255.0) would be 0.0.0.255. The correct command, Router(config-router)# network 192.168.10.0 0.0.0.255 area 0,will configure OSPF to run over any local interfaces assigned an IP address beginning with 192.168.10, since the inverse mask dictates that the first three octets must be a match.
The command Router(config-router)# network 192.168.10.0 0.0.255.255 area 0 is incorrect because it uses an improper wildcard mask. This mask would instruct OSPF to advertise any network with a prefix longer than the 192.168.0.0/16 network.
When routing does not seem to be working correctly, one of the first things to check is whether OSPF is operating on the proper interfaces. OSPF is enabled by network statements. To verify the network statements that were entered, you should execute the show run command and examine the output. If the network statement is configured so that the interface on the router is not in that network, OSPF will not operate on that interface. For example, suppose that Router A has an interface of 192.168.5.1/30 and the show run command produces the following output:

<output omitted>
router ospf 2 area 0
network 192.168.5.0 0.0.0.4
In this case, OSPF will not operate on the interface because the router interface is not in the network indicated by the network statement. The problem is not the network address but the wildcard mask. For a 30-bit mask, the wildcard should be 0.0.0.3, not 0.0.0.4. The wildcard mask can be determined by subtracting the regular mask value in the last octet (252) from 255, which is 3. The solution would to remove the incorrect statement and enter the correct statement as follows:
routerA(config)# router ospf 2 area 0
no network 192.168.5.0 0.0.0.4 area 0
network 192.168.5.0 0.0.0.3 area 0
References:
https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_ospf/configuration/12-4t/iro-12-4t-book/iro-cfg.html#GUID-51A06D7A-7099-453C-A9FD-34CE45080796