

Cisco.Premium.200-105.by.VCEplus.Update.172q

Number: 200-105 VCEplus
Passing Score: 800
Time Limit: 120 min
File Version: 1.0



Exam Code: 200-105

Exam Name: Interconnecting Cisco Networking Devices Part 2 - v3.0

Certification Provider: Cisco

Corresponding Certification: CCNA Routing and Switching

Website: www.vceplus.com

Free Exam: <https://vceplus.com/exam-200-105-icnd2-v3/>

Questions & Answers Exam Engine is rigorously checked before being put up for sale. We make sure there is nothing irrelevant in 200-105 exam products and you get latest questions. We strive to deliver the best 200-105 exam product for top grades in your first attempt.

VCE to PDF Converter : <https://vceplus.com/vce-to-pdf/>

Facebook: <https://www.facebook.com/VCE.For.All.VN/>

Twitter : https://twitter.com/VCE_Plus

Google+ : <https://plus.google.com/+Vcepluscom>

LinkedIn : <https://www.linkedin.com/company/vceplus>

Exam A**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?

- 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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

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-BFIC0C656997>

QUESTION 2

You are configuring a PPP connection between two routers, R1 and R2. The password for the connection will be poppycock. When you are finished you execute the show run command on R1 to verify the configuration.

Which of the following examples of partial output of the show run command from R1 represents a correct configuration of PPP on R1?

A.

```
enable password griswald
hostname R1
username R1 password poppycock
interface serial 0/0
ip address 192.168.5.5 255.255.255.0
encapsulation ppp
ppp authentication chap
```

B.

```
enable password griswald
hostname R1
username R1 password poppycok
interface serial 0/1
ip address 192.168.5.5 255.255.255.0
encapsulation ppp
ppp authentication chap
```

 **VCEplus**

VCE To PDF - Free Practice Exam

C.

```
enable password griswald
hostname R1
username R2 password poppycock
interface serial 0/0
ip address 192.168.5.5 255.255.255.0
encapsulation ppp
ppp authentication chap
```

D.

```
enable password griswald
hostname R1
username R1 password griswald
interface serial 0/0
ip address 192.168.5.5 255.255.255.0
encapsulation ppp
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

The correct configuration is as follows:

```
enable password griswald
hostname R1
username R2 password poppycock
interface serial 0/0
ip address 192.168.5.5 255.255.255.0
encapsulation ppp ppp authentication chap
```

The key settings that are common problems are as follows:

The username is set to the hostname of the other router (in this case, R2)

The password is set poppycock which must be the same in both routers

The following set is incorrect because the username is set to the local hostname (R1) and not the hostname of the other router (R2):

```
enable password griswald
hostname R1
username R1 password poppycock
interface serial 0/0
ip address 192.168.5.5 255.255.255.0
encapsulation ppp
ppp authentication chap
```

The following set is incorrect because the password is misspelled. It should be poppycock, not poppycok.

```
enable password griswald hostname R1
username R1 password poppycok interface serial 0/0 ip address 192.168.5.5 255.255.255.0 encapsulation ppp ppp authentication chap
```

The following set is incorrect because the password is set to the enable password of the local router (R1) rather than the agreed upon PPP password, which is poppycock.

```
enable password griswald
hostname R1
username R1 password griswald
interface serial 0/0
ip address 192.168.5.5 255.255.255.0
encapsulation ppp
ppp authentication chap
```

References:

<https://www.cisco.com/c/en/us/support/docs/wan/point-to-point-protocol-ppp/25647-understanding-ppp-chap.html>

QUESTION 3

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?

- 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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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 4

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

- A. Split horizon
- B. Spanning Tree Protocol (STP)
- C. Holddowns
- D. Route poisoning

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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.

router A# 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 5

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

- A. Broadcast
- B. NBMA
- C. Point-to-point
- D. Point-to-multipoint

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

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 6

You are configuring Open Shortest Path First (OSPF) protocol for IPv6 on Router5. The router has two interfaces, which have been configured as follows:

S0/0-192.168.5.1/24

S0/1 -10.0.0.6/8

You would like OSPF to route for IPv6 only on the S0/0 network. It should not route for IPv6 on the S0/1 network. The process ID you have chosen to use is 25. You do not want to apply an IPv6 address yet

Which of the following command sets would enable OSPF for IPv6 as required?

A.

```
Router5(config)#ipv6 ospf 25  
Router5(config)# network 192.168.5.0
```

B.

```
Router5(config)#ipv6 ospf 25  
Router5(config)#router-id 192.168.5.1
```

C.

```
Router5(config)#ipv6 unicast-routing  
Router5(config)#ipv6 router ospf 25  
Router5(config-rtr)#router-id 1.1.1.1  
Router5(config)#interface S0/0  
Router5(config-if)#ipv6 ospf 25 area 0
```

D.

```
Router5(config)#ipv6 unicast-routing  
Router5(config)#ipv6 ospf 25  
Router5(config-rtr)#router-id 1.1.1.1
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

The correct command sequence would be as follows:

```
Router5(config)# ipv6 unicast-routing
Router5(config)# ipv6 router ospf 25
Router5(config-rtr)# router-id 1.1.1.1
Router5(config)# interface S0/0
Router5(config-if)# ipv6 ospf 25 area 0
```

The first line enables IPv6 routing with the ipv6 unicast-routing command. The second line enables OSPF routing for IPv6 with the ipv6 router ospf command. The third assigns a necessary router ID (which was chosen at random) with the router-id command. The last two lines enable OSPF for area 0 on the proper interface.

The following command set is incorrect because it does not enable OSPF routing for IPv6, assign a necessary router ID, or enable OSPF for area 0 on the proper interface:

```
Router5(config)# ipv6 ospf 25
Router5(config)# network 192.168.5.0
```

This command set also displays incorrect use of the network command. The network command would be used with OSPF v2.

The following command set fails to enable OSPF routing for IPv6, assign a necessary router ID, or enable OSPF for area 0 on the proper interface:

```
Router5(config)# ipv6 ospf 25
Router5(config)# router-id 192.168.5.1
```

It also assigns the router ID under global configuration mode, rather than under router ospf 25 configuration mode as required.

The following command set fails to enable OSPF for area 0 on the proper interface:

```
Router5(config)# ipv6 unicast-routing
Router5(config)# ipv6 ospf 25
Router5(config-rtr)# router-id 1.1.1.1
```

References:

<https://search.cisco.com/search?querv=Cisco%20IOS%20IPv6%20Confiauration%20QGuide&locale=enUS&tab=Cisco>

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/command/ipv6-cr-book/ipv6-i5.html#wp2095571844>

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipv6/command/ipv6-cr-book/ipv6-i3.html#wp4031648257>

QUESTION 7

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 802.1 q trunk to be created on port Fa0/1 ?

- A. switchport mode trunk
- B. switchport mode nonegotiate
- C. switchport trunk encapsulation 302.1 q
- D. switchport trunk native VLAN

Correct Answer: A

Section: (none)

Explanation



Explanation/Reference:

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 802.1 q 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 802.1 q command does not need to be executed for a trunk to form. Also, the output Operational Trunking Encapsulation: dot1 q indicates that 802.1 q 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 8

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

- 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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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_paper09186a00801_b096a.html

QUESTION 9

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

- A. CSU/DSU
- B. switch
- C. modem
- D. hub

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

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 10

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

- 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

Correct Answer: C

Section: (none)

Explanation

**Explanation/Reference:**

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 cleartext.

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.asD?D=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 11

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?

- A. ICMP redirect
- B. ICMP time exceeded
- C. ICMP destination unreachable
- D. ICMP echo-reply

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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 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 12

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
```

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

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 if 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 13

```
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?

- 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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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 14

You apply the following commands to a router named R2:

```
R2(config)# interface Tunnell
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?

- 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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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 15

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

- A. variable length subnet masks (VLSM)
- B. partial updates
- C. neighbor discovery mechanism
- D. multiple vendor compatibility

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

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 16

Which of the following commands would instruct OSPF to advertise ONLY the 192.168.10.0/24 network in Area 0?

- 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
```

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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 be 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-cfa.html#GUID-51A06D7A-7099-453C-A9FD-34CE45080796

QUESTION 17

Which Cisco Internetwork Operating System (IOS) command is used to view the VLAN Trunking Protocol (VTP) statistics information?

A. show vtp status

- B. show vtp domain
- C. show vtp statistics
- D. show vtp counters

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

The show vtp counters command is used to view VTP statistics information. The syntax of the command is as follows:

show vtp {counters | status}

The parameters used in the command are counters, which specifies VTP statistics information, and status, which specifies VTP domain status information.

The following is the output of the show vtp counters command:

Router#show vtp counters

VTP statistics:

Summary advertisements received: 7

Subset advertisements received: 6

Request advertisements received: 0

Summary advertisements transmitted: 894

Subset advertisements transmitted: 13

Request advertisements transmitted: 3

Number of config revision errors: 0

Number of config digest errors: 0

Number of V1 summary errors: 0

VTP pruning statistics:

Trunk Join Transmitted Join Received Summary advts received from on-pruning-capable device

Fa0/2 43450 42691 6

The show vtp status command option is incorrect because this command is used to view VTP domain status information.

The show vtp domain and show vtp statistics commands are invalid options because they are not valid Cisco IOS commands.



QUESTION 18

You are implementing IP SLA and would like to use it to measure hop-by-hop response time between a Cisco router and any IP device on the network. Which of the following IP SLA operations would you use for this?

- A. ICMP path echo operation
- B. Internet Control Message Protocol Echo Operation
- C. UDP Jitter Operation for VoIP
- D. UDP Jitter Operation

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

The ICMP path echo operation discovers the path using the traceroute command, and then measures response time between the source router and each intermittent hop in the path. IP SLAs allow users to monitor network between the source router and each intermittent hop in the path. IP SLAs allow users to monitor network performance between Cisco routers or from a Cisco router to a remote IP device.

The Internet Control Message Protocol (ICMP) Echo Operation measures end-to-end response time between a Cisco router and any IP-enabled device. Response time is computed by measuring the time taken between sending an ICMP echo request message to the destination and receiving an ICMP echo reply. It does not measure hop-by-hop response time.

The UDP Jitter Operation for VoIP is an extension to the current jitter operations with specific enhancements for VoIP. The enhancements allow this operation to calculate voice quality scores and simulate the codec's directly in CLI and the MIB. It does not measure hop-by-hop response time.

The UDP Jitter Operation is designed to measure the delay, delay variance, and packet loss in IP networks by generating active UDP traffic. It does not measure hop-by-hop response time.

References:

https://www.cisco.com/en/US/technologies/tk648/tk362/tk920/technologies_white_paper09186a00802d5efe.html

QUESTION 19

What will be the effects of executing the following set of commands? (Choose all that apply.)

```
router(config)# router eigrp 44
router (config-router)# network 10.0.0.0
router (config-router)# network 192.168.5.0
```

- A. EIGRP will be enabled in AS 44
- B. EIGRP instance number 44 will be enabled
- C. EIGRP will be activated on the router interface 10.0.0.2/8
- D. EIGRP will be activated on the router interface 192.168.5.9/24
- E. EIGRP will be activated on the router interface 10.0.5.8/16
- F. EIGRP will be activated on the router interface 192.168.6.1/24

Correct Answer: ACDE

Section: (none)

Explanation

Explanation/Reference:

The effects of executing this set of commands will be that Enhanced Interior Gateway Routing Protocol (EIGRP) will be enabled in Autonomous System (AS) 44 and will be active on the router interfaces 10.0.0.2/8, 192.168.5.9/24, and 10.0.5.8/16.

The router eigrp 10 command is used to enable EIGRP on a router. The network 10.0.0.0 and network 192.168.5.0 commands are used to activate EIGRP over any interfaces that fall within the major networks 10.0.0.0 and 192.168.5.0, or within any subnets of these classful networks. The network commands in EIGRP configuration ignore any subnet-specific information by default. Since the IP address 10.0.5.8/24 is in a subnet of the Class A IP network 10.0.0.0, and only the first octet (byte) of a Class A IP address represents the major (classful) network, the remaining bytes are ignored by the network command.

EIGRP instance number 44 will not be enabled. The number 44 in the command does not represent an instance of EIGRP; it represents an autonomous system (AS) number. The autonomous-system parameter of the router eigrp command (router eigrp 44) specifies the autonomous system number. To ensure that all the routers in a network can communicate with each other, you should specify the same autonomous system number on all routers. EIGRP will not be activated on the router interface 192.168.6.1/24. This interface does not exist within the Class C network 192.198.5.0 or Class A network 10.0.0.0, or within any of their subnets.

References:

CCNA ICND2 Official Exam Certification Guide (Cisco Press. ISBN 1-58720-181-X), Chapter 10: EIGRP, pp. 389-390.

<https://search.cisco.com/search?query=Cisco%20IOS%20IP%20Routing%20Configuration%20Guide&locale=enUS&tab=Cisco>

QUESTION 20

You are configuring a WAN connection between two offices. You cannot ping between the routers in a test. The Serial0 interface on RouterA is connected to the Serial1 interface on RouterB.

The commands you have executed are shown below. What is the problem with the configuration?

```
RouterA(config)#username RouterB password lie
RouterA(config)#interface serial0
RouterA(config-if)#encapsulation ppp
RouterA(config-if)#ppp authentication chap
```

```
RouterB(config)#username RouterA password lie
RouterB(config)#interface serial0
RouterB(config-if)#encapsulation ppp
RouterB(config-if)#ppp authentication chap
```



- A. The passwords are incorrectly configured
- B. The usernames are incorrectly configured
- C. The wrong interface has been configured
- D. The encapsulation is incorrect on RouterA
- E. The encapsulation is incorrect on RouterB
- F. The authentication types do not match

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

The two routers are connected using Serial0 on RouterA and Serial1 on RouterB. However, the configuration commands were executed on interface Serial0 on RouterB. So although the configuration itself is completely correct, it is configured on the wrong interface.

The passwords are correct. The passwords should match on both routers. In this case, they are both set to lie. If even one character does not match, including character casing, the authentication and the connection will fail.

The usernames are correct. The username should be set to the host name of the peer router. In this case, RouterA's username is set to RouterB and RouterB's username is set to RouterA, which is correct.

The encapsulations are correct. They are both set to PPP, which is the correct type of encapsulation when using authentication. The authentication types do match. They are both set to CHAP. It is possible to configure two authentication methods, with the second used as a fallback method in cases where the other router does not support the first type. The command below would be used to enable CHAP with PAP as a fallback method:


RouterB(config-if)#ppp authentication chap pap

References:

<https://www.cisco.com/c/en/us/support/docs/wan/point-to-point-protocol-ppp/25647-understanding-ppp-chap.html>

QUESTION 21

Consider the following diagram



Which of the following routing protocols could NOT be used with this design?

- A. RIPv1
- B. RIPv2
- C. EIGRP
- D. OSPF

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

The network design displayed has subnets of a major classful network located in opposite directions from the perspective of some of the individual routers. This configuration can be accommodated by any routing protocol that supports Variable Length Subnet masks (VLSM) or the transfer of subnet mask information in routing advertisements.

RIPv1 supports neither of these. RIPv1 will automatically summarize routing advertisements to their classful network (in this case 192.168.1.0/24). This action will cause some of the routers to have routes to the same network with different next hop addresses, which will NOT work.

EIGRP, RIPv2 and OSPF all support VLSM and can be used in the design shown in the scenario.

References:

<https://www.cisco.com/c/en/us/support/docs/ip/ip-routed-protocols/13722-ripv1-support-vlsm.html>

QUESTION 22

Consider the following output of the show ip interface brief command

```
R1# show ip interface brief
```

Interface	IP-Address	OK?	Method	Status	Protocol
Ethernet0	192.168.12.65	YES	manual	up	up
Ethernet1	192.168.12.129	YES	manual	up	up
Serial0	192.168.12.187	YES	manual	up	up
Serial1	192.168.12.125	YES	manual	up	up
Serial2	192.168.12.121	YES	manual	up	up
Serial3	unassigned	YES	unset	up	up

You have a single area OSPF network. What command should you execute on R1 so that OSPF is operational on the E0, S1, and S2 interfaces ONLY?

- A. R1 (config-router)#network 192.168.12.64 0.0.0.127 area 0
- B. R1 (config-router)#network 192.168.12.64 0.0.0.63 area 0
- C. R1 (config-router)#network 192.168.12.64 0.0.0.66 area 0
- D. R1 (config-router)#network 192.168.12.64 255.255.255.192 area 0
- E. R1 (config-router)#network 192.168.12.64 0 0.0 63 area 1

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

The command R1(config-router)#network 192.168.12.64 0.0.0.63 area 0 would ensure that OSPF is operational on the E0, S1, and S2 interfaces only. When executing the network command in OSPF, a wildcard mask in combination with the network ID used in the command determines which interfaces will participate in OSPF. Any interfaces that are included in the network created by the network ID and the mask will participate in OSPF.

Wildcard masks in OSPF network statements are expressed inversely, and not as a regular subnet masks. For example, 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 network ID is the starting point and the wildcard mask specifies where the network will end or the range of the network. In this case, the network begins at 192.168.12.64. The value in the last octet of the mask indicates the number of values (including 64) that will be included in the network, which means that it will range from 192.168.12.64 -192.168.12.127. 64 to 127 equals 64 values if you include the endpoints 64 and 127.

The network, and therefore the operation of OSPF, includes the interfaces E0 (192.168.12.65), S1 (192.168.12.125), and S2 (192.168.12.121) because these three IP addresses lie within the range 192.1268.12.64-192.168.12.127.

The command R1(config-router)#network 192.168.12.64 0.0.0.127 area 0 is incorrect because the resulting network would range from 192.168.12.64 - 192.168 12.191. This would include all of the required interfaces, but would also include E1 (192.168.12.129) and S0 (192.18.12.187), which is not desired.

The command R1(config-router)#network 192.168.12.64 0.0.0.66 area 0 is incorrect because the resulting network would range from 192.168.12.64 - 192.168 12.129. This would include all of the required interfaces, but would also include E1 (192.168.12.129).

The command R1(config-router)#network 192.168.12.64 255.255.255.192 area 0 is incorrect because the mask, while correct in its breadth and the

exact inverse of the wild card mask 0 0.0.63, is not stated in wildcard mask format.

The command R1 (config-router)#network 192.168.12.64 0.0.0.63 area 1 is incorrect because it specifies area 1. At least one area of an OSPF network must be area 0 and since this is a single area OSPF network, the command must specify area 0.

References:

<https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/7039-1.html#t10>

QUESTION 23

You have executed the following commands on a switch:

```
Switch64(config)# interface range gigabitethernet2/0/1 -2
Switch64(config-if-range)# switchport mode access
Switch64(config-if-range)# switchport access vlan 10
Switch64(config-if-range)# channel-group 5 mode auto
```

In which of the following situations will Switch64 create an Etherchannel?

- A. If the other switch is set for desirable mode
- B. If the other switch is set for auto mode
- C. If the other switch is set for on mode
- D. If the other switch is set for passive mode

Correct Answer: A

Section: (none)

Explanation



Explanation/Reference:

The Etherchannel will be created if the other end is set to desirable mode. The configuration shown in the example is using Port Aggregation protocol (PAGP). This protocol has two settings: desirable and auto. Two ends will negotiate and will only create an Etherchannel under two conditions: if one end is set to auto and the other end is set to desirable, or if both ends are set for desirable.

It will not form an Etherchannel if the other end is set to auto mode. When both ends are set to auto mode, an Etherchannel will not form.

It will not form an Etherchannel if the other end is set to on mode. On mode disables negotiation of any kind, which will prevent an Etherchannel from forming unless the other end is also set for on.


It will not form an Etherchannel if the other end is set to passive mode. Passive is a setting used in Link Aggregation Protocol (LACP). The two protocols are not compatible

References:

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst3750x_3560x/software/release/12-2_55_se/configuration/guide/3750xsca/swethchl.html

QUESTION 24

The four switches in the diagram below have default configurations. Considering the bandwidths indicated on each link and the MAC addresses indicated for each switch, which ports will be forwarding after RSTP has converged? (Choose all that apply.)



- A. SW 1 port that connects to SW 4
- B. SW 1 port that connects to SW 2
- C. SW 1 port that connects to SW 3
- D. SW 2 port that connects to SW 3
- E. SW 2 port that connects to SW 4
- F. SW 3 port that connects to SW 4
- G. SW 3 port that connects to SW 1
- H. SW 3 port that connects to SW 2

Correct Answer: ADFGH

Section: (none)

Explanation

Explanation/Reference:

The ports that will be forwarding after convergence are the SW1 port that connects to SW4, the SW2 port that connects to SW3, and all of the ports on SW3. The process of determining these port states occurs in this order:

Selection of the root bridge. All ports on the root bridge become designated ports and are set to forward. Determination of the root ports on each non-root bridge.

Determination of the designated port on each segment that does not connect directly to the root bridge. Designated and root ports will be set to forwarding, and all other ports will be set to discarding.

For step 1, when all bridge priorities have been left to their default, all switches will have same bridge priority. When that is the case, as in this scenario, the switch with the lowest MAC address will be selected as root bridge. In this case, SW3 has the lowest MAC address and becomes the root bridge.

ALL ports are in a forwarding state on the root bridge, which explains why all of the ports on SW3 will be in a forwarding state.

For step 2, each non-root bridge will select the interface it possesses with the least cost path to the root bridge. Once selected, that port will be placed in a forwarding state. 100 Mbps links will be assigned a cost of 19, and 1 Gbps links will be assigned a cost of 4. Each path cost is the cumulative cost of the links in the path. The root ports for the non-root bridges are determined as follows.

SW1 has four paths to the root bridge, with each path yielding the following costs:

SW1 to SW3 (100 Mbps) cost = 19

SW1 to SW4 to SW3 (1 Gbps + 1 Gbps) cost = 4 + 4 = 8

SW1 to SW2 to SW 4 to SW3 (100 Mbps + 100 Mbps + 1 Gbps) cost=19+ 19 + 4 = 42

SW1 to SW2 to SW3 (100 Mbps + 1 Gbps) cost = 19 + 4 = 23

SW1 will use the lowest cost path (SW1 to SW4 to SW3) as its root path, It will set the SW1 connection to SW4 to

forwarding and the connection from SW1 to SW3 to blocking. The status of its third interface (SW1 to SW2) will be determined in Step 3, since it is a shared segment with SW2 that does not have a direct connection to the root bridge.

Switch 2 (SW2) has three paths to the root bridge, with each path yielding the following costs:

SW2 to SW3 (1 Gbps) cost = 4

SW2 to SW1 to SW3 (100 Mbps + 100 Mbps) cost = 19 + 19 = 38 SW2 to SW4 to SW3 (100 Mbps + 1 Gbps) cost = 19 + 4 = 23

SW2 will use the lowest cost path (SW2 to SW3) as its root path and will set the SW2 connection to SW3 to forwarding. The status of its second and third interfaces (SW2 to SW1 and SW2 to SW4) will be determined in step 3 since both are shared segments with SW2 and SW4 respectively that do not have a direct connection to the root bridge.

Switch 4 (SW4) has four paths to the root bridge, with each path yielding the following costs:

SW4 to SW3 (1 Gbps) cost = 4

SW4 to SW1 to SW3 (1 Gbps + 100 Mbps) cost = 4+19 = 23

SW4 to SW2 to SW1 to SW3 (100 Mbps + 100 Mbps + 100 Mbps) cost =19 + 19 + 19 = 57

SW4 to SW2 to SW3 (1 Gbps + 100 Mbps) cost = 4 +19 = 23

SW4 will use the lowest cost path (SW4 to SW3) as its root path and set the SW4 connection to SW3 to forwarding. The status of SW4's second and third interfaces, SW4 to SW1 and SW4 to SW2, will be determined in step 3. Since these interfaces are shared segments with SW1 and SW2, they do not have a direct connection to the root bridge.

For Step 3, there are two segments in the diagram (SW1 to SW2 and SW2 to SW4) that do not connect directly to the root bridge. The interface on either end of the segment that has the least cost path to the root bridge will be the designated port for that section.

The designated port of each segment is determined in this way.

For the SW1 to SW2 segment, the SW2 end of the segment has a shortest path cost of 1 Gbps (4) to the root, and the SW1 end of the segment has a shortest path cost through SW4 of 2 Gbps (8) to the root. The SW2 port to SW1 will be the designated port and will be forwarding.

For the SW2 to SW4 segment, the SW2 end of the segment has a shortest path cost of 1 Gbps (4) to the root and the SW4 end of the segment has a shortest path cost of 1 Gbps (4) to the root. This is a tie. In the case of a tie, the interface connected to the switch with the lowest MAC address becomes the designated port for the segment. SW4 has the lowest MAC address, so the SW4 port to SW2 will be the designated port and will be forwarding.

Once determined, the designated and root ports will be set to forwarding and all other ports will be set to discarding. The converged state of all ports is shown in the diagram below.



References:

QUESTION 25

A. show running-configuration

- B. show vlans
- C. show standby brief
- D. show standby

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

The show vlans command would NOT be useful for discovering the problem. When troubleshooting a problem with Hot Standby Router Protocol (HSRP), the show vlans command will yield no useful information. The output of the command is shown below, demonstrating that there is no HSRP information provided.

```
router# show vlan trunk
VLAN Name Status IfIndex Mod/Ports, Vlans
-----
1 default active 5 2/1-2
6/4-8
15 VLAN0015 active 18 6/1,6/3
16 VLAN0016 active 19 6/2
23 VLAN0023 active 20
26 VLAN0026 active 21
31 VLAN0031 active 22
39 VLAN0039 active 23
```



All three of the remaining commands will be useful in discovering information. Each is shown below with an example of its application to troubleshooting.

Example A: show running-configuration

Router B is not taking over as the active device for VLAN 101's HSRP group when the Fa0/1 interface on Router A fails. Below is a partial output of show run for both routers with the output focused on the section concerning VLAN 101 's configuration on each.

<pre>routerA interface Vlan101 <output omitted> Standby 5 ip 172.63.51.250 Standby 5 priority 180 standby preempt standby track FastEthernet 0/1 5</pre>	<pre>routerB interface vlan101 Standby 5 ip 172.63.51.250 Standby 5 priority 170 standby preempt Standby track FastEthernet 0/1</pre>
----------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------

The above output displays the source of the problem. Router A has a decrement value of 5 configured for Fa0/1, as shown on the last line of the output after the specification of FastEthernet 0/1. This means that when its Fa0/1 interface goes down, Router A will subtract 5 from its priority for the VLAN 101 group, lowering it to 175. This is still higher than the priority of Router B, which is 170. Therefore, the solution is to change the decrement value for Router A to at least 11. When the interface goes down, Router A's priority will be decremented to 169, allowing Router B to take the role as active for the HSRP group in VLAN 101.

Example B: show standby brief

Router C is not taking over as the active device for VLAN 102's HSRP group when the Fa0/1 interface on Router D fails. Below is a partial output of show standby brief for both routers C and D, with the output focused on the section concerning VLAN 102's configuration on each.

Router C

Interface Grp Prio P State Active addr Standby addr Group addr Fa0/1 102 200 Active local 10.10.10.253 10.10.10.251 Router D

Interface Grp Prio P State Active addr Standby addr Group addr Fa0/1 102 200 P Active local 10.10.10.253 10.10.10.251

The absence of a P in the P (preempt) column in the output for Router C shows that it is not set to preempt. If not configured to preempt, it will never take over for Router D, regardless of its priority with respect to Router D.

Example C: show standby

Router F is supposed to be the active router for VLAN 103's HSRP group. Occasionally both routers are shut down for maintenance over the weekend. After the routers are rebooted, Router F is not taking over as the active device for VLAN 103's HSRP group. Below is a partial output of the show standby command for both routers, with the output focused on the section concerning VLAN 103's configuration on each

Router E

FastEthernet 0/1 - Group 1

State is Active

2 state changes, last state change 00:30:59

Virtual IP address is 10.1.0.20

Secondary virtual IP address 10.1.0.21

Active virtual MAC address is 0004.4d82.7981

Local virtual MAC address is 0004.4d82.7981 (bia)

Hello time 4 sec, hold time 12 sec

Next hello sent in 1.412 secs

Preemption enabled, min delay 50 sec, sync delay 40 sec

Active router is local

Standby router is 10.1.0.7, priority 140(expires in 9.184 sec)

Priority 200 (configured 200)

Tracking interface FastEthernet 0/1 state up decrement 10

Router F

FastEthernet 0/1 - Group 1 State is Active

2 state changes, last state change 00:30:59 Virtual IP address is 10.1.0.20

Secondary virtual IP address 10.1.0.21

Active virtual MAC address is 0004.4d82.7981

Local virtual MAC address is 0004.4d82.7981 (bia)

Hello time 4 sec, hold time 12 sec Next hello sent in 1.412 secs

Preemption enabled, min delay 50 sec, sync delay 40 sec

Active router is 10.1.0.6, priority 190(expires in 9.184 sec)
Standby router is local Priority 190 (configured 190)
Tracking interface FastEthernet 0/1 state up decrement 50

The output shows that Router F is not assuming the active role because of the priority and decrement values configured on the routers. When both routers go down, Router E will decrement its priority (200) by 10, as shown in last two lines of its output, leaving the priority at 190. Router F will decrement its priority (190) by 50 as shown in last two lines of its output, leaving the priority at 140. Therefore, to ensure that Router F maintains its role as active even after the dual shutdowns, the priority of Router F should be increased to at least 241. When both routers decrement their priorities after shutdown, Router F will then have a priority of 191, which will be higher than the priority value of Router E.

References:

<https://www.cisco.com/c/en/us/support/docs/ip/hot-standby-router-protocol-hsrp/10583-62.html>

<https://www.cisco.com/c/en/us/support/docs/ip/hot-standby-router-protocol-hsrp/13780-6.html>

QUESTION 26

You receive the following error message after addressing and enabling an interface

%192.168.16.0 overlaps with FastEthernet0/0

Which two are NOT the causes of the error message? (Choose two.)

- A. incorrect subnet mask in the new interface
- B. incorrect IP address on the new interface
- C. incorrect encapsulation configured
- D. failure to issue the no shutdown command



Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

The error message %192.168.16.0 overlaps with FastEthernet0/0 indicates that the newly configured interface is in the same subnet as an existing interface. This can occur if there is an incorrect subnet mask or an address that inadvertently places the new interface in that subnet. Each router interface must be in a different subnet to function. For example, when the series of commands below is executed on a router, it will elicit the error message because the two IP addresses used are in the same subnet given the subnet mask in use.


```
Router#config t
Router(config)#interface S0
Router(config-if)#ip address 192.168.1.17 255.255.255.0
Router(config-if)#no shutdown
Router(config-if)#interface S1
Router(config-if)#ip address 192.168.1.65 255.255.255.0
Router(config-if)#no shutdown
0%192.168.1.0 overlaps with Serial 0
```

It's also a valuable skill to be able to recognize these problems before the router tells you about them.

An incorrect encapsulation would prevent the interface from working, but would not generate this message.

If the no shutdown command had not been issued, we would not be receiving this error. It is only generated when an attempt is made to enable an incorrectly configured interface.

References:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/ip/configuration/guide/fipr_c/lcfipadr.html

QUESTION 27

You and your team are evaluating the use of OSPFv3 in your IPv6 network. Which of the following statements is true of OSPFv3?

- A. There will be a higher demand on the processor to run the link-state routing algorithm
- B. Router IDs must match for adjacency formation
- C. Area IDs do not need to match for adjacency formation
- D. Area types do not need to match for adjacency formation

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

There will be a higher demand on the processor to run the link-state routing algorithm. As with OSPFv2, OSPFv3 uses the Shortest Path first (SPF) algorithm, which is processor intensive. It is one of the only downsides of using the algorithm.

OSPFv3 also shares a number of other characteristics with its v2 counterpart with respect to adjacency formation. For example:

Router IDs should not match.

Router IDs should reflect the correct router ID for each device.

Area IDs must match.

Area types must match.

References:

<https://supportforums.cisco.com/t5/network-infrastructure-documents/troubleshooting-ospfv3-neighbor-adjacencies/ta-p/3112861>

QUESTION 28

Which of the following technologies should be used to prevent a switching loop if a switch is connected to a port configured for PortFast?

- A. RSTP
- B. BPDU Guard
- C. Root Guard
- D. PVST

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

BPDU Guard prevents switching loops in the case of a switch being connected to a PortFast interface. PortFast is used for ports that connect to host systems, such as workstations and printers, and allows the port to immediately enter a forwarding state. This bypasses the normal 30-second delay that Spanning Tree Protocol would normally use to determine if a switch has been connected to the port. Implementing BPDU Guard will disable the port if a switch is connected and a BPDU is received.

Rapid Spanning Tree Protocol (RSTP) is incorrect because this is an enhanced Spanning Tree standard that operates on the Data Link layer of the OSI model. RSTP was not designed to protect PortFast ports. PortFast and BPDU Guard are supported by RSTP, but they not required or configured by default.

Root Guard is incorrect because it is used to protect the root bridge placement in the Spanning Tree, not to protect PortFast ports.

Per-VLAN Spanning Tree (PVST) is incorrect because this is an implementation of Spanning Tree (the default protocol for Cisco switches), and was not designed to protect PortFast ports. PortFast and BPDU Guard are supported by RSTP, but are not required, and must be configured manually.

References:

<https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/24062-146.html> 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 29

Which Cisco Internetwork Operating System (IOS) command is used to view information about Open Shortest Path First (OSPF) routing processes?

- A. show ip ospf database
- B. show ip ospf statistics
- C. show ip ospf
- D. show ip ospf traffic

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

The show ip ospf command is used to view information about the OSPF routing processes. It does so by displaying the collection of link states present in the database. The syntax of the command is as follows:

Router# show ip ospf [process-id]

The process-id parameter of the command specifies the process ID. The output of the command is as follows:

```
Router# show ip ospf

Routing Process "ospf 203" with ID 21.0.0.1 and Domain ID 21.20.0.1
Supports only single TOS(TOS0) routes
Supports opaque LSA
SPF schedule delay 10 secs, Hold time between two SPFs 20 secs
Minimum LSA interval 10 secs, Minimum LSA arrival 5 secs
LSA group pacing timer 200secs
Interface flood pacing timer 110 msec
Retransmission pacing timer 110 msec
Number of external LSA 1, Checksum Sum 0x0
Number of opaque AS LSA 1, Checksum Sum 0x0
Number of DCbitless external and opaque AS LSA 0
Number of DoNotAge external and opaque AS LSA 0
Number of areas in this router is 3, 1 normal 0 stub 1 nssa
External flood list length 0

Area BACKBONE(0)
Number of interfaces in this area is 4
Area has message digest authentication
SPF algorithm executed 6 times
Area ranges are
Number of LSA 3, Checksum Sum 0x298E5
Number of opaque link LSA 1, Checksum Sum 0x0
Number of DCbitless LSA 3
Number of indication LSA 0
Number of DoNotAge LSA 0
Flood list length 0
```

The show ip ospf database command is incorrect because this command is used to view the OSPF database for a specific router.

The show ip ospf statistics command is incorrect because this command is no longer valid in IOS version 12.4. The show ip ospf traffic command is incorrect because this command is no longer valid in IOS version 12.4.

QUESTION 30

Which commands would be used to enable Enhanced Interior Gateway Routing Protocol (EIGRP) on a router, and configure the IP addresses 10.2.2.2 and 192.168.1.1 as a part of complete EIGRP configuration? (Choose three.)