

200-105.exam

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




Cisco

 **200-105**
Interconnecting Cisco Networking Devices Part 2
VCE To PDF - Free Practice Exam
Version 1.0

Exam A

QUESTION 1

In the network exhibit, the routers are running OSPF and are set to the default configurations. (Click the Exhibit(s) button to view the network.)



What would be the effect of configuring a loopback interface on RouterA with an address of 192.168.1.50/24?

- A. Router B would become the DR
- B. Router A would become the DR
- C. Router C would become the DR
- D. Router A would become the BDR



Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Configuring a loopback interface on RouterA with an address of 192.168.1.50/24 would cause Router A to become the designated router (DR). The designated router (DR) is determined by the router with the highest interface priority number. If the priority numbers are tied, then the router with the highest router ID (RID) becomes the DR.

The default priority number is 1, and can be configured as high as 255. Changing the priority to 0 would make the router ineligible to become the DR or the backup designated router (BDR). The `ip ospf priority #` command is used to manually configure a priority on a specific interface.

Router IDs are determined first by the highest loopback IP address, followed by the highest IP address on an active physical interface. Thus, in the case of a priority tie, the router with the highest loopback IP address will have the highest RID, and will become the DR for the network segment.

The current Router ID for a router can be determined by executing the show ip interface brief command. In the sample output of the show ip interface brief command below, the RID will be 10.108.200.5.

Router# show ip interface brief

```
Interface IP-Address OK? Method Status Protocol
Ethernet0 10.108.00.5 YES NVRAM up up
Ethernet1 unassigned YES unset administratively down down
Loopback0 10.108.200.5 YES NVRAM up up
Serial0 10.108.100.5 YES NVRAM up up
Serial1 10.108.40.5 YES NVRAM up up
Serial2 10.108.100.5 YES manual up up
Serial3 unassigned YES unset administratively down down
```

Neither Router B nor C will be the DR because the IP addresses on their physical interfaces are lower than 192.168.1.50/24.

Router A will not be the backup designated router. Since it is the DR, it cannot also be the BDR.

Router C will not be the BDR because its IP address is lower than that of Router B. Router B will be the BDR.

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

QUESTION 2

Which Cisco IOS command is used to configure encapsulation for a PPP serial link on a Cisco router?

- A. encapsulation ppp
- B. encapsulation ip ppp
- C. ip encapsulation ppp
- D. encapsulation ppp-synch

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

PPP is a Layer 2 protocol encapsulation type that supports both synchronous and asynchronous circuits and provides built-in security mechanisms. The encapsulation ppp interface configuration mode command is used to configure encapsulation for a PPP (Point to Point Protocol) serial link on a Cisco router. PPP encapsulation provides for router-to-router and host-to-network connections over both synchronous and asynchronous circuits. Serial links are configured to use Cisco High Level Data Link Control (HDLC) encapsulation, by default, on Cisco routers. The Cisco version of HDLC is incompatible with the industry standard version used on other router brands because it contains a type field that identifies the underlying network protocol being encapsulated by HDLC. This is a beneficial feature of Cisco HDLC but makes it incompatible with other router brands.

For this reason, a Cisco router that is going to be connected to a non-Cisco router should be configured to use PPP instead of the default. The encapsulation ppp interface configuration mode command will do this. If you set one of the routers for PPP and leave the other router at the default encapsulation for a serial connection, the connection will fail due to incompatible encapsulation.

You would use the show run command to verify matching encapsulation types. In the partial output of the show run command for two routers shown below, it can be seen that although one of the routers has the encapsulation ppp command in its configuration, the other does not. The absence of the encapsulation ppp command means that the default HDLC is being used. This incompatibility will cause both routers to report a serial interface up, line protocol down condition since the connection is live, but the Layer 2 framing is misconfigured.

router1#show run	router2#show run
<output omitted>	<output omitted>
interface serial 0/0	interface serial 0/1
encapsulation ppp	



If authentication between the routers is also required, the authentication pap, authentication ms-chap, or authentication chap commands could be used to apply Password Authentication Protocol (PAP), Microsoft Challenge Authentication Protocol (MS-CHAP), or Challenge Authentication Protocol (CHAP) authentication to the connection, respectively.

A full configuration of a serial link for using PPP with authentication is as shown below:

```
Router1(config)#interface Serial0
Router1(config-if)#encapsulation ppp
Router1(config-if)#ppp authentication pap
```

Note above that the third line enables PAP authentication, which is not secure. Alternately, you can use CHAP authentication (which is secure) with the ppp authentication chap command. Regardless of which authentication mechanism you choose, these authentication commands will only be accepted on an interface where PPP encapsulation has been enabled, which rules out any non-serial interfaces.

The third type of encapsulation that can be configured on a serial WAN link is Frame Relay, which can be selected with the encapsulation frame relay command under the interface.

In summary, the three encapsulation types available for WAN serial links are PPP, HDLC, and Frame Relay. The command for each is as follows, executed under the interface configuration prompt:

```
encapsulation ppp
encapsulation hdlc
encapsulation frame relay
```

All other options are invalid commands.

References:

http://docwiki.cisco.com/wiki/Point-to-Point_Protocol

QUESTION 3

Which service is denoted by TCP/UDP port number 53?

- A. Domain Name Service (DNS)
- B. File Transfer Protocol (FTP)
- C. Telnet
- D. HTTP

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) port number 53 is assigned to Domain Name Service (DNS), which is used to convert hostnames into Internet Protocol (IP) addresses.

Some common TCP and UDP port number assignments are as follows:

- port 25: Assigned to Simple Mail Transfer Protocol (SMTP), a TCP protocol used to send and receive e-mail messages.
- port 23: Assigned to Telnet to allow remote logins and command execution.
- port 21: Assigned to File Transfer Protocol (FTP). It is used to control FTP transmissions. Port number 20 is also used by FTP for FTP data. - port 80: Assigned to Hypertext Transfer Protocol (HTTP), which is the base for transferring Web pages over the Internet.

References:

http://docwiki.cisco.com/wiki/Internetworking_Basics#Multiplexing_Basics

QUESTION 4

Which of the following is a Point-to-Point Protocol (PPP) authentication protocol that supports sending of hashed values instead of sending passwords in clear text?

- A. LCP
- B. NCP
- C. PAP
- D. CHAP

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

There are two authentication methods available when implementing a PPP connection: Password Authentication Protocol (PAP) and Challenge Handshake Authentication Protocol (CHAP).

Challenge Handshake Authentication Protocol (CHAP) uses a one-way hash function based on the Message Digest 5 (MD5) hashing algorithm to hash the password. This hashed value is then sent across the wire. In this situation, the actual password is never sent. No one tapping the wire will be able to reverse the hash to come up with the original password. This is why MD5 is referred to as a one-way function. It cannot be reverse engineered. CHAP uses a three-way handshake process to perform the authentication. Moreover, CHAP periodically repeats the authentication process after link establishment.

When configuring PPP with CHAP authentication, both routers must be configured with a username that will be presented by the other router with a password. Therefore, the username to configure on Router A will be the username of Router B. The password should be the same on both machines. If these settings are not correct, then authentication will fail. The authentication process can be displayed as it happens with the debug PPP authentication command.

Link Control protocol (LCP) is defined in Request for Comments (RFCs) 1548 and 1570 and has primary responsibility to establish, configure, authenticate, and test a PPP connection. LCP negotiates the following when setting up a PPP connection: - Authentication method used (PAP or CHAP), if any

- Compression algorithm used (Stacker or Predictor), if any
- Callback phone number to use, if defined
- Multilink; other physical connections to use, if configured

Network Control Protocol (NCP) defines the process for how the two PPP peers negotiate which network layer protocols, such as IP and IPX, will be used across the PPP connection. LCP is responsible for negotiating and maintaining a PPP connection whereas NCP is responsible for negotiating upper-layer protocols that will be carried across the PPP connection.

Password authentication Protocol (PAP) is simpler than CHAP, but less secure. During the authentication phase, PAP goes through a two-way handshake process. In this process, the source sends its user name (or hostname) and password in clear text, to the destination. The destination compares this information with a list of locally stored user names and passwords. If it finds a match, the destination returns an accept message. If it does not find a match, it returns a reject message.

References:

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

QUESTION 5

You are configuring an authenticated connection between two routers named Tacoma and Lansing. The connection on the Lansing end is correctly set up with a password of keypass. You are directing an assistant to configure the name and password on Tacoma. Which of the following commands would be correct to complete this authenticated connection?

- A. username Tacoma password keypass
- B. username Lansing keypass password
- C. username Tacoma keypass password
- D. username Lansing password keypass

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

To complete the configuration, you should run the command `username Lansing password keypass`. This command creates a user account for the Lansing router with a password of keypass.

When creating an authenticated connection between the routers, a user account must be created for the other router. The password configured must match on both ends.

When examining the output produced by the `show running-configuration` command for two routers, the output should read as below:

<pre>Tacoma# show running-config <some output text omitted> enable password cisco ! hostname Tacoma username Lansing password keypass !</pre>	<pre>Lansing# show running-config <some output text omitted> enable password cisco1 ! hostname Lansing username Tacoma password keypass !</pre>
--	--

The lines that display enable password cisco and enable password cisco1 represent local passwords to enable privileged mode on the local router. These passwords do not have to match. The lines of output that must display matching passwords are username Lansing password keypass and username Tacoma password keypass.

You should not run the command username Tacoma password keypass. The username Tacoma portion of the command will create an account named Tacoma. You need an account for the other router, Lansing.

You should not run the command username Lansing keypass password. The password portion of the command must follow the syntax password [correct_password].

You should not run the command username Tacoma keypass password. The username Tacoma portion of the command will create an account for the wrong router, and the password portion of the command must follow the syntax password [correct_password].

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

QUESTION 6

Which of the following statements are true when discussing link state and distance vector routing protocols? (Choose all that apply.)

- A. After convergence, routing advertisements are only triggered by changes in the network with distance vector protocols
- B. Packets are routed based upon the shortest path calculated by an algorithm with link state protocols
- C. Only one router in an OSPF area can represent the entire topology of the network
- D. Distance vector protocols send the entire routing table to a neighbor
- E. Distance vector protocols send updates regarding the status of their own links to all routers in the network
- F. Link-state protocols place a high demand on router resources running the link-state algorithm
- G. Distance vector protocols require a hierarchical IP addressing scheme for optimal functionality
- H. Link-state protocols use hello packets and LSAs from other routers to build and maintain the topological database
- I. Link-state protocols require a hierarchical IP addressing scheme for optimal functionality.

Correct Answer: BDFHI

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The following statements are true of link-state and distance vector routing protocols:

- Packets are routed based upon the shortest path calculated by an algorithm with link state protocols.
- Distance vector protocols send the entire routing table to a neighbor.

- Link-state protocols place a high demand on router resources running the link-state algorithm.
- Link-state protocols use hello packets and LSAs from other routers to build and maintain the topological database.
- Link-state protocols require a hierarchical IP addressing scheme for optimal functionality.

Link state protocols like OSPF use the Shortest Path First algorithm to calculate the shortest path based on a metric called cost, while distance vector protocols like RIP consider only hop count when determining the best route. Running the algorithm places a high demand on router resources. Distance vector protocols are required to send the entire routing table with each update, while link state protocols only send updates when required by changes in the network. Therefore, less traffic is created with link state protocols.

Sending routing advertisements after convergence only when changes occur in the network is a characteristic of link state protocol's not distance vector protocols. With distance vector protocols, updates occur regularly and include the entire routing table.

All routers in an OSPF area can represent the entire topology of the network, not just one.

Distance vector protocols do not send updates regarding the status of their own links to all routers in the network. Updating link status is a characteristic of link state protocols. Distance vector protocols send the entire routing table.

Distance vector protocols do NOT require a hierarchical IP addressing scheme for optimal functionality. Link-state protocols do require this for optimal functionality, as it supports more efficient route aggregation or summarization. This reduces the number of routes in the table and the number of calculations required by the SPF algorithm, thereby lowering router resource demand.

References: http://docwiki.cisco.com/wiki/Routing_Basics#Link-State_Versus_Distance_Vector

QUESTION 7

Which Enhanced Interior Gateway Routing Protocol (EIGRP) packet is NOT sent reliably over the network?

- A. Update
- B. Query
- C. Reply
- D. Acknowledgement

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Acknowledgement packets are sent unreliably over the network, and there is no guaranteed delivery of acknowledgement packets between neighboring routers.

Acknowledgement packets are a special type of hello packets that do not contain data and have a non-zero acknowledgement number. These are sent as a unicast.

Update, Query, and Reply packets use Reliable Transport Protocol (RTP), which ensures guaranteed delivery of packets between neighboring devices. The RTP mechanism ensures loop-free synchronized network.

References:

http://docwiki.cisco.com/wiki/Enhanced_Interior_Gateway_Routing_Protocol

QUESTION 8

What will an EIGRP router do if the successor route fails and there is no feasible successor?

- A. EIGRP will mark the route as passive until a new successor route is determined.
- B. EIGRP will redistribute routes into RIP or OSPF.
- C. EIGRP will query neighboring routers until a new successor route is determined.
- D. EIGRP will forward traffic to the neighbor with the lowest administrative distance.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Feasible successors are backup routes for the successor (active) route to a remote network. If a successor route fails, and a feasible successor is available, the feasible successor will immediately become the successor and be installed in the routing table. This provides EIGRP with virtually instantaneous convergence. If no feasible successor is available, then the router must send out query packets to neighboring EIGRP routers to find an alternate path to the remote network.

EIGRP routes are marked as active when the network is converging. Passive routes are stable, converged routes.

EIGRP will not redistribute routes into RIP or OSPF. Redistribution allows information learned from one routing protocol to be converted into routes for injection into the autonomous system of another routing protocol. This allows networks learned via EIGRP, for example, to be visible and reachable from hosts in a RIP routing domain. Redistribution has nothing to do with EIGRP convergence or with the determination of a new successor route.

Administrative distance is used to determine which source of routing information is considered more trustworthy when multiple routing protocols have been implemented. Administrative distance has no effect on EIGRP convergence or the determination of a new successor route.

References: <https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16406-eigrp-toc.html>

QUESTION 9

Which of the following excerpts from the output of the show ip eigrp topology command include EIGRP learned routes or pairs of routes that will be included in the routing table? (For excerpts that include multiple routes, do not include the entry unless BOTH routes will be included in the routing table.)

- A. P 172.16.16.0/24, 1 successors, FD is 284244
via 172.16.250.2 (284244/17669856), Serial0/0
via 172.16.251.2 (12738176/27819002),
Serial0/1 B. P 172.16.250.0/24, 1 successors,
FD is 2248564 via Connected, Serial0/0
- C. P 172.16.10.0/24 2 successors, FD is 284244 via
172.16.50.1 (284244/17669856), Serial1/0 via
172.16.60.1 (284244/17669856), Serial1/1
- D. P 172.16.60.0/24, 1 successors, FD is
2248564 via Connected, Serial1/1

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The following excerpt indicates two successor routes, and they will both be included:

P 172.16.10.0/24 2 successors, FD is 284244
via 172.16.50.1 (284244/17669856), Serial1/0
via 172.16.60.1 (284244/17669856), Serial1/1

Both of these routes will be included because they have identical metrics (284244/17669856). Only the EIGRP successor routes will appear in the routing table, as these are considered the best-path routes to each remote network.

The route for 172.16.16.0/24 via 172.16.251.2 (12738176/27819002) will not be included because only successor routes are included, and this route is a feasible successor. Feasible successor routes are routes that are used only as a backup if the successor route(s) becomes unavailable. If you examine the output of each option, it will indicate how many successor routes are in the entry. The entry shows that there is only one successor to this route:

P 172.16.16.0/24, 1 successors, FD is 284244
via 172.16.250.2 (284244/17669856), Serial0/0
via 172.16.251.2 (12738176/27819002), Serial0/1

The first listed is the successor and the second is the feasible successor. The first has the best or lowest metric (284244/17669856), which is the criterion used for selection.

These entries indicate successor routes, but they also indicate they are via Connected, which means they are networks directly connected to the router.

P 172.16.250.0/24, 1 successors, FD is
2248564 via Connected, Serial0/0 and

P 172.16.60.0/24, 1 successors, FD is 2248564
via Connected, Serial1/1

Therefore, they are not EIGRP learned routes.

References:

https://www.cisco.com/c/en/us/td/docs/ios/iproute_eigrp/command/reference/ire_book/ire_s1.html

QUESTION 10

When the auth keyword is used in the snmp-server host command, which of the following must be configured with an authentication mechanism?

- A. the interface
- B. the host
- C. the user
- D. the group

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The auth keyword specifies that the user should be authenticated using either the HMAC-MD5 or HMAC-SHA algorithms. These algorithms are specified during the creation of the SNMP user.

For example, the following command creates a user named V3User who will be a member of the SNMP group V3Group and will use HMAC-MD5 with a password of Password:

snmp-server user V3User V3Group v3 auth md5 Password

The authentication mechanism is not configured on the interface. All SNMP commands are executed at the global configuration prompt.

The authentication mechanism is not configured at the host level. The version and security model (authentication, authentication and encryption, or neither) are set at the host level.

The authentication mechanism is not configured at the SNMP group level. The group level is where access permissions like read and write are set. This is why a user account must be a member of a group to derive an access level, even if it is a group of one.

References:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/configfun/configuration/guide/ffun_c/fcf014.html

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mcl/allreleasemcl/all-book.html>

QUESTION 11

Which of the following are classless routing protocols? (Choose four.)

- A. Open Shortest Path First (OSPF)
- B. Enhanced Interior Gateway Routing Protocol (EIGRP)
- C. Interior Gateway Routing Protocol (IGRP)
- D. Routing Information Protocol version 1 (RIPv1)
- E. Border Gateway Protocol (BGP)
- F. Routing Information Protocol version 2 (RIPv2)



Correct Answer: ABEF

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), Border Gateway Protocol (BGP), and Routing Information Protocol version 2 (RIPv2) are classless routing protocols.

Intermediate-System-to-Intermediate System (IS-IS) is also a classless routing protocol.


The options IGRP and RIPv1 are incorrect because these are classful routing protocols.

The following are characteristics of classless routing protocols:

- The subnet mask is advertised with each route by using classless routing protocols.

- Flexible route summarization and supernetting (CIDR) are allowed in classless routing protocols.
- Classless routing protocols support variable length subnet masks (VLSM), which allow different subnets of a given IP network to be configured with different subnet masks.

One of the main advantages of using a classless routing protocol is its ability to minimize the effects of discontinuous networks. When subnets of the same classful network are separated by another classful network, the networks are called discontinuous. Examine the diagram below:



The LAN networks extending from Router A and Router B are derived from the same Class C network, 192.168.1.0/24. A classful routing protocol such as RIP v1 would not be able to determine the direction to send the packets, but since classless protocols include the subnet mask in advertisements, they would not suffer the same problem. Whenever networks with non-default subnet masks are used, a classless routing protocol will be required.

Below are some examples of networks that do not have default masks. You can recognize them by the fact that they are not /8, /16, or /24.

192.168.10.0/27
10.5.6.0/22
172.68.0.0/18

All of the classless protocols discussed here are interior routing protocols with the exception of Border Gateway Protocol (BGP), which is an external routing protocol used to connect different autonomous systems. For example, BGP would be used to connect two OSPF autonomous systems (AS).

References: <https://www.cisco.com/c/en/us/tech/ip/ip-routing/index.html>

QUESTION 12

What port types are available for Rapid Spanning Tree Protocol (RSTP) but NOT available in Spanning Tree Protocol (STP)? (Choose two.)

- A. Root port
- B. Backup port
- C. Alternate port

- D. Designated port
- E. Learning port

Correct Answer: BC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

RSTP was developed to reduce the high convergence times required in STP, and introduces the alternate port and backup port roles. RSTP is an Institute of Electrical and Electronics Engineers (IEEE) standard, 802.1w, and is interoperable with 802.1d (STP). It operates on the Data Link layer of the OSI model.

An alternate port is a port that has an alternative path or paths to the root bridge, but is currently in a discarding state. A backup port is a port on a segment that could be used to reach the root port, but there is already an active designated port for the segment. An alternate port can also be described as a secondary, unused root port, and a backup port as a secondary, unused designated port.

A root port is a port on non-root switches used to reach the root switch. There can be only one root port on a switch, and it is determined by the least path cost to the root switch. Root ports are used in STP and RSTP.

A designated port is the port used by a network segment to reach the root switch. Designated ports lead away (downstream) from the root switch, and are determined by the lowest path cost to the root switch. While a switch can only have one root port, every other port could potentially be a designated port. Whenever a network segment could be serviced by more than one switch, STP will elect one switch as designated for the segment, and the other(s) will be blocking. This is a core function of the STP protocol, in that only one active Layer 2 path can exist between any two network segments. This port type is available in STP.

A learning port is not a valid port type in STP or RSTP. Learning is one of the possible port states in STP and RSTP. STP has five port states; blocked, listening, learning, forwarding, and disabled. There are only three port states in RSTP; discarding, learning, and forwarding.

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

QUESTION 13

What Cisco Catalyst switch feature can be used to define ports as trusted for DHCP server connections?

- A. DHCP snooping
- B. port security
- C. 802.1x
- D. private VLANs

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

DHCP snooping is used to define ports as trusted for DHCP server connections. The purpose of DHCP snooping is to mitigate DHCP spoofing attacks. DHCP spoofing is an attack that can be used to force user traffic through an attacking device. This is accomplished by an attacker responding to DHCP queries from users. Eliminating the response from the correct DHCP server would make this more effective, but if the attacker's response gets to the client first, the client will accept it.

The DHCP response from the attacker will include a different gateway or DNS server address. If they define a different gateway, the user traffic will be forced to travel through a device controlled by the attacker. This will allow the attacker to capture traffic and gain company information. If the attacker changes the DNS server in the response, they can use their own DNS server to force traffic to selected hosts to go to a device they control. Again, this would allow the attacker to capture traffic and gain information.

DHCP snooping can be used to determine what ports are able to send DHCP server packets, such as DHCPOFFER, DHCPACK, and DHCPNAK, from the company DHCP server. DHCP snooping can also cache the MAC address to IP address mapping for clients receiving DHCP addresses from a valid DHCP server.

The three required steps to implement DHCP snooping are:

1. Enable DHCP snooping globally with the ip dhcp snooping command:

```
switch(config)# ip dhcp snooping
```

2. Enable DHCP snooping for a VLAN with the vlan parameter:

```
switch(config)# ip dhcp snooping vlan vlan #
```

(for example, ip dhcp snooping 10 12 specifies snooping on VLANs 10 and 12)

3. Define an interface as a trusted DHCP port with the trust parameter:

```
switch(config-if)# ip dhcp snooping trust
```

When specifying trusted ports, access ports on edge switches should be configured as untrusted, with the exception of any ports that may have company DHCP servers connected. Only ports where DHCP traffic is expected should be trusted. Most certainly, ports in any area of the network where attacks have been detected should be configured as untrusted.

Some additional parameters that can be used with the ip dhcp snooping command are:

- switch(config)# ip dhcp snooping verify mac-address - this command enables DHCP MAC address verification.
- switch(config)# ip dhcp snooping information option allow-untrusted - this command enables untrusted ports to accept incoming DHCP packets with option 82 information. DHCP option 82 is used to identify the location of a DHCP relay agent operating on a subnet remote to the DHCP server.

When DHCP snooping is enabled, no other relay agent-related commands are available. The disabled commands include:


```
ip dhcp relay information check global configuration ip
dhcp relay information policy global configuration ip
dhcp relay information trust-all global configuration ip
dhcp relay information option global configuration ip
dhcp relay information trusted interface configuration
```

Private VLANs are a method of protecting or isolating different devices on the same port and VLAN. A VLAN can be divided into private VLANs, where some devices are able to access other devices and some are completely isolated from others. This was designed so service providers could keep customers on the same port isolated from each other, even if the customers had the same Layer 3 networks.

Port security is a method of only permitting specified MAC addresses access to a switch port. This can be used to define what computer or device can be connected to a port, but not to limit which ports can have DHCP servers connected to them.

802.1x is a method of determining authentication before permitting access to a switch port. This is useful in restricting who can connect to the switch, but it cannot control which ports are permitted to have a DHCP server attached to it.

References: <https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4500/12-2/20ew/configuration/guide/config/dhcp.html>

QUESTION 14

You have implemented SNMP v3 in your network. After making the configuration changes, you find that technicians in the TECHS group cannot access the MIB. You execute the show run command and receive the following output that relates to SNMP:

<output omitted>

```
snmp-server group NORMAL v3 priv read NORMAL write NORMAL
snmp-server group TECHS v3 priv read TECHS access 99 snmp-
server group TRAP v3 priv
```

!!

```
snmp-server user NORMAL NORMAL v3 auth sha CISCO priv des56 CISCO snmp-
server user TECHS TECHS v3 auth sha CISCO priv des56 CISCO snmp-server user
TRAP TRAP v3 auth sha CISCO priv des56 CISCO
```

```
snmp-server enable traps snmp linkup linkdown snmp-
server host 155.1.146.100 traps version 3 priv TRAP
```

What is preventing the TECHS group from viewing the MIB?

- A. The presence of the keyword priv in the command creating the RESTRICTED group
- B. A mismatch between the authentication mechanism and the encryption type in the command creating the TECHS user
- C. The absence of an access list defining the stations that can be used by the TECHS group
- D. The presence of the keyword auth in the command creating the TECHS user

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The command that creates the TECHS group ends with the parameter access 99: snmp-server group TECHS v3 priv read TECHS access 99

This indicates that the access list number 99 is specifying the IP addresses of the stations allowed to connect to the MIB for the group. Since the access list is missing from the configuration, no IP addresses will be allowed, and no connections can be made by the group.

The presence of the keyword priv in the command creating the TECHS group is not causing the issue. This keyword indicates that encryption (privacy) and authentication should both be used on all transmissions by the group.

In SNMPv3, there are three combinations of security that can be used:

- noAuthNoPriv- no authentication and no encryption; includes the noauth keyword in the configuration
- AuthNoPriv - messages are authenticated but not encrypted; includes the auth keyword in the configuration
- AuthPriv - messages are authenticated and encrypted; includes the priv keyword in the configuration

There is no mismatch between the authentication mechanism and the encryption type in the command creating the TECHS user.

snmp-server user TECHS TECHS v3 auth sha CISCO priv des56 CISCO

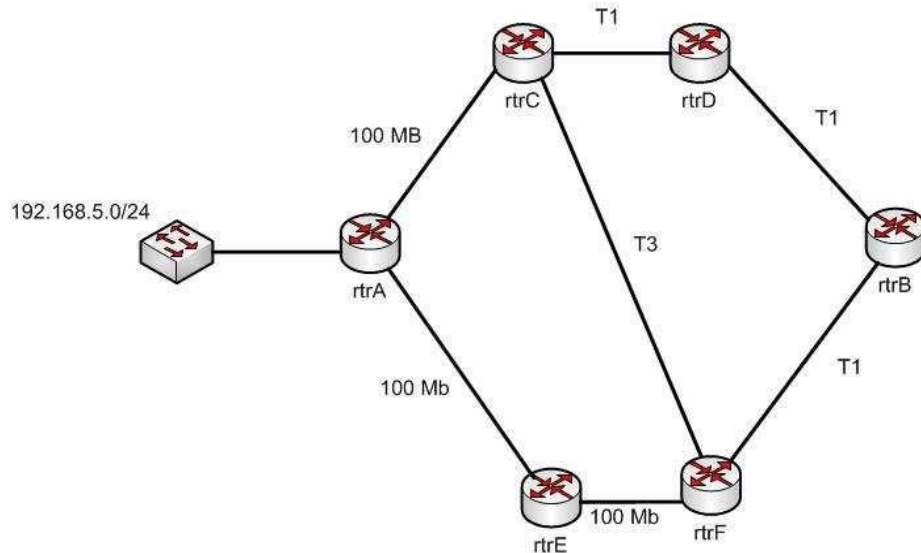
In the preceding command, the section auth sha CISCO specified that messages are authenticated using SHA with a key of CISCO. It does not need to match the section priv des56 CISCO, which indicates that encryption (priv) will be provided using DES56 with a key of CISCO.

The presence of the keyword auth in the command creating the TECHS user is not causing the issue. This line indicates that messages are authenticated using SHA with a key of CISCO.

References: <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/snmp/configuration/xr-3se/3850/snmp-xr-3se-3850-book/nm-snmp-snmpv3.html>

QUESTION 15

Examine the following diagram:



While troubleshooting an OSPF routing problem, you need to determine the cost for Router F to reach the 192.168.5.0 24 network via the best route.

What will that cost be?

- A. 110
- B. 2C. 3
- D. 7

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The best route to the 192.168.5.0/24 network from the perspective of router F will have an OSPF assigned cost of 2. There are three possible loop-free paths to get from router F to the 192.168.5.0/24 network. The default OSPF costs for a 100 MB link, a T1 link, and a T3 link are 1, 64, and 2, respectively.

The three paths and the calculation of their costs are shown:

Router F to Router E to Router A: $1 + 1 = 2$

Router F to Router C to Router A: $2 + 1 = 3$

Router F to Router B to Router D to Router C to Router A: $64 + 64 + 64 + 1 = 193$

Each OSPF route calculates the cost of its path to a network, and passes that value on to the next router, which will then add to it the cost to reach that neighbor. For example, the routing table of Router E would look like this for the route to 192.168.5.0/24:

O 192.168.5.0 [110/1] via <output omitted>

Router F would add its own cost to reach Router E to the cost of reaching 192.168.5.0/24, resulting in the following output:

O 192.168.5.0 [110/2] via <output omitted>

110 is the administrative distance of OSPF.

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

QUESTION 16

Which Cisco IOS command would prompt for input in the following format?

Protocol [ip]:

Target IP address: 10.1.1.1

Repeat count [5]:

Datagram size [100]:

Timeout in seconds [2]:

Extended commands [n]: y

Source address or interface: 192.142.23.10

Type of service [0]:

Set DF bit in IP header? [no]:

Validate reply data? [no]:

Data pattern [0xABCD]:

Loose, Strict, Record, Timestamp, Verbose[none]:

Sweep range of sizes [n]:

Type escape sequence to abort

A. ping 10.1.1.1

B. ping

C. traceroute

D. tracer

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The extended ping command prompts the user for input in the format given in this scenario. The extended ping command is accessed by issuing a ping command without specifying an IP address. This causes the ping command to transit into extended ping command mode, where you can specify and modify various parameters, such as packet size, timeout, and repeat count.

The following code is a sample partial output of the extended ping command:

```
Router A#ping
Protocol [ip]:
Target IP address: 10.1.1.1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 192.142.23.10
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
```



The true value of the extended ping command lies in the ability to ping FROM a different device than the one you are working from. As shown in the above output, you can specify the source address on line 8.

The ping 10.1.1.1 command is incorrect because it sends an ICMP "echo request" to the target host. In turn, the target host replies with the "echo reply" packets. When pinging from one device to another on the network, ICMP and Address Resolution Protocol (ARP) are used. ARP resolves an IP address to its associated MAC addresses.

The tracert command is incorrect because this command is used by Microsoft Windows, not Cisco. It is not a valid utility to run via the Cisco IOS command-line interface. The traceroute command is similar to the traceroute Cisco utility as the traceroute command tests the connectivity or "reachability" of a network device or host. It reports back a reply at each hop, allowing one to determine where the communication link is "broken".

The traceroute command is used to display the path that a packet follows to its destination. This command displays the IP address of each router in the path from the source to the destination address. Unlike the Microsoft tracert command, which uses the ICMP protocol, the Cisco traceroute command is based on User Datagram Protocol (UDP). The following code is the partial output of the traceroute command.

```
RouterA#traceroute 124.10.23.41
```

```
Type escape sequence to abort.
```

```
Tracing the route to 124.10.23.41
```

```
 1 121.10.1.3 6 msec 6 msec 6 msec
```

```
 2 134.10.10.13 30 msec 17 msec 14 msec 3 32.1.2.4 36 msec * 23 msec
```

References:

<https://www.cisco.com/c/en/us/support/docs/ip/routing-information-protocol-rip/13730-ext-ping-trace.html> <https://www.cisco.com/c/en/us/support/docs/ios-nx-os-software/ios-software-releases-121-mainline/12778-ping-traceroute.html>

QUESTION 17

What command would provide the output displayed in the exhibit? (Click on the Exhibit(s) button.)

Interface	Grp	Prio	P	State	Active	Standby	Virtual IP
Vl64	2	100	P	Standby	192.168.64.10	local	192.168.64.1
Vl65	1	110	P	Active	local	192.168.65.20	192.168.65.1
Vl66	2	100	P	Standby	192.168.66.10	local	192.168.66.1
Vl67	1	110	P	Active	local	192.168.67.20	192.168.67.1
Vl68	2	100	P	Standby	192.168.68.10	local	192.168.68.1
Vl69	1	110	P	Active	local	192.168.69.20	192.168.69.1
Vl70	2	100	P	Active	local	192.168.70.20	192.168.70.1

- A. switch# show hsrp
- B. switch# show standby
- C. switch# show interface vlan
- D. switch# show standby brief

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The command `show standby brief` displays the output in the exhibit. It is used to display a summary of the HSRP groups of which the switch is a member. The summary information it provides includes the group number, priority, state, active device address, standby address, and group address. In the exhibit, the interface VLAN 64 is a member of HSRP group 2. Its priority in the group is 100 and it is currently the standby switch. Since preemption is configured (as indicated by the P following the priority), we know that the priority of this switch must be lower than the priority of the active device. The active device has an IP address of 192.168.64.10 and the group IP address is 192.168.64.1.

The command `show standby` can be used to display detailed information about HSRP groups of which a switch is a member. It does not provide the quick summary display of the exhibit. This command displays information about HSRP on all configured interfaces and for all HSRP groups. It also displays hello timer information and the expiration timer for the standby switch. The command syntax is `show standby [type number [group]]`.

Below is an example of this command's output:

RouterA#show standby vlan 5

VLAN 5 - group 1

Local state is Active, priority 105, may preempt

Hellotime 3 sec, holdtime 10 sec

Next hello sent in 1.424

Virtual IP address is 192.12.23.10 configured

Active router is local

Standby router is 192.12.23.3 expires in 9.600

Virtual mac address is 0000.0c07.ac01

2 state changes, last state change 00:01:38

<output omitted>

VLAN 5- group 2

Local state is Standby, priority 100

Hellotime 3 sec, holdtime 10 sec

Next hello sent in 1.424

Virtual IP address is 192.12.23.11 configured

Active router is 192.168.23.3 expires in 9.600

Standby router is local

2 state changes, last state change 00:01:38

<output omitted>

In the above output, Router A is load-sharing traffic for VLAN 5. It is active for group 1 and standby for group 2. The router at address 192.168.23.3 is active for group 2 and standby for group 1. This allows traffic to be sent to both routers while still allowing for redundancy. Router A was also configured with the standby 1 preempt command (results seen in line 1), which allows it to resume its role as active for group 1 if it comes back up from an outage.

The command show interface vlan is not a complete command. A VLAN number must follow the command. When provided with a VLAN number, the output would display the status of the SVI, but no HSRP information.

The command show hsrp is not a valid command due to incorrect syntax.

References:

https://www.cisco.com/c/en/us/td/docs/ios/ipapp/command/reference/iap_s4.html

QUESTION 18

You run the following command:

```
switch# show ip interface brief
```

What information is displayed?

- A. A summary of the IP addresses and subnet mask on the interface
- B. A summary of the IP addresses on the interface and the interface's status
- C. The IP packet statistics for the interfaces
- D. The IP addresses for the interface and the routing protocol advertising the network

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The command show ip interface brief displays a summary of the IP address on the interface and the interface's status. The status shows whether the interface is up. This command is useful when you are connected to a router or switch with which you are not familiar, because it allows you to obtain the state of all interfaces or switch ports.

Sample output of this command is shown below:

```
Switch88# show ip interface brief
```

```
Interface IP-Address OK? Method Status Protocol
```

```
FastEthernet0/1 unassigned YES manual down down
```

```
FastEthernet0/2 unassigned YES manual down down
```


FastEthernet0/3 unassigned YES manual down down
FastEthernet0/4 unassigned YES manual down down
FastEthernet0/5 unassigned YES manual down down
FastEthernet0/6 unassigned YES manual down down
FastEthernet0/7 unassigned YES manual down down
FastEthernet0/8 unassigned YES manual up up
FastEthernet0/9 unassigned YES manual down down
FastEthernet0/10 unassigned YES manual down down

This command does not display subnet mask information. You should use other commands, such as show ip interface or show run interface, to verify the subnet mask.

IP statistics about the interface are displayed with the command show ip interface. Adding the brief keyword tells the switch to leave out everything but the state of the interface and its IP address.

To view the routing protocol advertising an interfaces network, you would use the command show ip protocol.

References: <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr/command/ipaddr-cr-book/ipaddr-r1.html>

QUESTION 19

Which of the following protocols is responsible for negotiating upper-layer protocols that will be carried across a Point-to-Point Protocol (PPP) connection?

- A. LCP
- B. NCP
- C. LMI
- D. ISDN

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Network Control Protocol (NCP) is responsible for negotiating upper-layer protocols that will be carried across the PPP connection. NCP defines how the two PPP peers negotiate with the network layer protocols, such as IP and IPX, which will be used across the PPP connection.

Link Control protocol (LCP) is not responsible for negotiating upper-layer protocols that will be carried across a PPP connection. Link Control protocol (LCP) has the primary responsibility of negotiating and maintaining the PPP connection. LCP, defined in Request for Comments (RFCs) 1548 and 1570, has the primary responsibility to establish, configure, authenticate, and test a PPP connection. LCP negotiates the following when setting up a PPP connection:

- Authentication method used (PAP or CHAP), if any
- Compression algorithm used (Stacker or Predictor), if any
- Callback phone number to use, if defined
- Multilink; other physical connections to use, if configured

Local Management Interface (LMI) is not responsible for negotiating upper-layer protocols that will be carried across the PPP connection. LMI is a characteristic of a frame relay connection. There are three types of LMIs supported by Cisco routers:

- Cisco
- ANSI Annex D
- Q933-A Annex A

LMI has nothing to do with PPP connections.

Integrated Services Digital Network (ISDN) is a type of WAN connection and has nothing to do with PPP connections.

References:

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

QUESTION 20

Which of the following is NOT a true statement regarding Virtual Private Networks (VPNs)?

- A. A VPN is a method of securing private data over public networks
- B. IPsec is a method for providing security over VPN
- C. Frame Relay is a Layer 3 VPN technology
- D. IPsec provides packet-level encryption
- E. A Cisco VPN solution provides increased security, reduced cost, and scalability

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Frame Relay is a Layer 2 VPN technology, providing connectivity over switched carrier Wide Area Networks (WANs). Packets are encapsulated in Frame Relay frames, and assigned Data Link Connection Identifiers (DLCIs) to identify to the local Frame Relay switch the virtual circuit (VC) that the data should follow.

A VPN is a method of securing private data over public networks (such as the Internet), so this is a true statement.

IPsec is a security framework that provides security for data traveling over VPNs, so this is a true statement. It is an open standard protocol framework that is used to secure end-to-end communications.

IPsec allows for encryption at the packet level (Layer 3) when configured in tunnel mode, so this is a true statement.

VPN solutions such as those supported by Cisco ASA firewalls and Cisco integrated routers provide the following benefits:

- Lower desktop support costs
- Threat protection
- Flexible and cost-effective licensing- Reduced cost and management complexity

References:

http://docwiki.cisco.com/wiki/Frame_Relay

http://docwiki.cisco.com/wiki/Virtual_Private_Networks

QUESTION 21

Which of the following commands configures an SNMP host to authenticate a user by username and send clear text notifications, the receipt of which will be acknowledged by the receiver?

- A. Router(config)# snmp-server host 192.168.5.5 informs version 3 noauth public
- B. Router(config)# snmp-server host 192.168.5.5 traps version 3 auth public
- C. Router(config)# snmp-server host 192.168.5.5 informs version 2c public
- D. Router(config)# snmp-server host 192.168.5.5 informs version 3 authpriv public

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The command snmp-server host 192.168.5.5 informs version 3 noauth CISCO will configure the host to authenticate a user by username and send clear text notifications. The receiver will then acknowledge receipt of the notification. The keyword informs indicates that an inform message type will be used. Unlike a trap, an inform message is acknowledged by the receiver.

The version 3 keyword indicates that version 3 is in use, which is the ONLY version that supports authentication and encryption. Finally, the noauth keyword specifies authentication by username only and no encryption.

The command snmp-server host 192.168.5.5 traps version 3 auth public configures the host to send traps rather than informs.

The command snmp-server host 192.168.5.5 informs version 2c public specifies version 2c, which only support community string-based authentication.

The command `snmp-server host 192.168.5.5 informs version 3 authpriv public` specifies the keyword `authpriv`, which indicates encryption will be used and authentication based on HMAC-MD5 or HMAC-SHA algorithms.


Reference:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/configfun/configuration/guide/ffun_c/fcf014.html

<https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mcl/allreleasemcl/all-book.html>

QUESTION 22

Refer to the network diagram in the exhibit. Host A is configured with an incorrect default gateway. All other computers and the Router are known to be configured correctly (Click the Exhibit(s) button.) Which of the following statements is TRUE?



- A. Host C on Network A cannot communicate with Host A on Network A.
- B. Host A on Network A can communicate with all other hosts on Network A.
- C. Host A on Network A can communicate with Router R.
- D. Host C on Network A cannot communicate with Router R.
- E. Host D on Network B cannot communicate with Host B on Network A.

Correct Answer: BC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Host A on Network A can communicate with all other hosts on Network A and with Router R. To communicate with local hosts and the interface of Router R (which are all in the same subnet) only a correct IP address is required. If the default gateway of Host A is incorrect, then it will not be able to communicate with any host on the other side of the router, which includes Network B in the diagram. Packets from hosts on Network B will reach Host A on Network A without any problem, because they possess the correct address of the default gateway or router, but Host A will send the packet to a dead end because Host A has an incorrect default gateway. On the other hand, Host A does not require a default gateway to communicate with other hosts on same network.

Host C on Network A WILL be able to communicate with Host A on Network A, even though Host A has an incorrect default gateway because Host A and C are in the same subnet, which requires no use of the of the gateway or router.

Host C on Network A WILL be able to communicate with Router R because Host C has the correct default gateway address which is the address of Router R.

Host D on Network B WILL be able to communicate with Host B on Network A because both hosts have a correct default gateway address.

References:

http://docwiki.cisco.com/wiki/Routing_Basics

<http://www.microsoft.com/technet/community/columns/cableguy/cg0903.msp>

<http://kb.iu.edu/data/aifx.html>

QUESTION 23

You set up several routers in your lab. Two of them are connected back to back using Data Terminal Equipment (DTE)-to-Data Circuit-terminating Equipment (DCE) cable. You need to configure the clock rate. On which router would you configure the clock rate?

- A. the DCE
- B. the DTE
- C. The clock rate is set by default
- D. The clock rate cannot be configured

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The clock rate is set on the Data Circuit-terminating Equipment (DCE) device. DCE is also known as Data Communications Equipment.

DCE terminates a physical WAN connection and provides clocking and synchronization of a connection between two locations and connects to a DTE. The DCE category includes equipment such as CSU/DSUs and modems. If you were connecting a router to a WAN link, the router would be the DTE end and would be connected to a CSU/DSU or a modem. Either of these devices would provide the clocking.

DTE is an end-user device, such as a router or a PC that connects to the WAN via the DCE device.

Other options are incorrect. By default, no clock rate is configured, but can be set on a DCE device by using the clock rate [bps] command.

QUESTION 24

Which of the following commands will set the line speed of a serial connection that connects to a Channel Service Unit /Digital Service Unit (CSU/DSU) at 56 Kbps?

- A. service-module 56000 clock rate speed
- B. service-module 56k clock rate speed
- C. bandwidth 56k
- D. bandwidth 56000

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The command service-module 56k clock rate speed will configure the network line speed for a 4-wire, 56/64-kbps CSU/DSU module.

The command service-module 56000 clock rate speed is incorrect because the speed must be stated in the form 56k (for Kbps), rather than 56000.

The bandwidth command is used to limit the amount of bandwidth used by an application when utilizing Quality of Service (QoS). It does not set the line speed of a serial connection that connects to a Channel Service Unit /Digital Service Unit CSU/DSU. Therefore, both the bandwidth 56k and the bandwidth 56000 commands are incorrect.

QUESTION 25

How is the designated router (DR) determined by OSPF on a multi-access network segment?

- A. The lowest interface priority, then the highest RID
- B. The highest interface priority, then the highest RID
- C. The lowest interface priority, then the highest OSPF process ID
- D. The highest interface priority, then the highest OSPF process ID

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

OSPF routers elect a designated router (DR) and backup designated router (BDR) on multi-access network segments in order to minimize the amount of update traffic sent between OSPF neighbors. All routers on a multi-access network segment form adjacencies with the DR and BDR, but not with each other. Network events are communicated to the DR, and the DR distributes the event to the rest of the network.

The DR is determined by the router with the highest interface priority number. If the priority numbers tie (which will be the case if they are left to the default of 1), then the router with the highest router ID (RID) becomes the DR. The default priority number is 1, and can be configured as high as 255.

In many cases, it is desirable to intervene in this process and select the router you want to be the DR. If that is the case and the selected router is not becoming the DR for whatever reason, the following options are available to ensure that the selected router wins the election: - Change the priority value of the router to a value higher than the other routers

- Set the priority value of the other routers to 0

- Create a loopback address on the selected router with an IP address higher than the IP addresses used on the other routers

Changing the priority to 0 makes the router ineligible to become the DR or BDR. The `ip ospf priority #` command is used to manually configure a priority on a specific interface.

It is also worth noting that a single OSPF area can have more than one DR. The election is NOT performed per area, but per network segment. So if you had six OSPF routers in area 0 with three in one IP subnet and three in another, there would be two elections, one for each segment.

The lowest interface priority does not determine the DR.

The OSPF process ID has no effect on DR elections.

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

QUESTION 26

Which WAN switching technology is used by Asynchronous Transfer Mode (ATM)?

- A. cell-switching
- B. virtual switching
- C. circuit-switching
- D. packet switching

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Cell switching is used by Asynchronous Transfer Mode (ATM). ATM is an International Telecommunication Union-Telecommunications (ITU-T) standard for transmission of data, voice, or video traffic using a fixed size frame of 53 bytes, known as cells. Out of these 53 bytes, the initial five bytes are header information and the remaining 48 bytes are the payload.

The term virtual switching is incorrect because it is not a valid WAN switching technology.

Circuit switching dynamically establishes a virtual connection between a source and destination. The virtual connection cannot be used by other callers unless the circuit is released. Circuit switching is the most common technique used with the Public Switched Telephone Network (PSTN) to make phone calls. The dedicated circuit is temporarily established for the duration of the call between caller and receiver. Once the caller or receiver hangs up the phone, the circuit is released and is made available to other users.

Packet switching is also used for data transfer but not in an ATM network. With packet switching, the data is broken into labeled packets and is transmitted using packet-switching networks. The Internet and LAN communications use packet switching.

References:

http://docwiki.cisco.com/wiki/Asynchronous_Transfer_Mode_Switching

QUESTION 27

Which of the following is a classful routing protocol?

- A. RIPv1
- B. EIGRP
- C. BGPv4
- D. RIPv2

Correct Answer: A

Section: (none)


Explanation

Explanation/Reference:

Explanation:

The Routing Information Protocol version 1 (RIPv1) is a classful routing protocol, which exchanges routes without including any subnet masking information. IP addresses in the routing table should have the same subnet mask. Because classful routing protocols may not fully utilize the available IP address range, all router interfaces within the same network must have the same subnet mask.

Open Shortest Path First (OSPF), Routing Information Protocol version 2 (RIPv2), Enhanced Interior Gateway Routing Protocol (EIGRP), and Border Gateway Protocol version 4 (BGPv4) are classless routing protocols. These protocols include the subnet mask in the route advertisement and support variable length subnet masks (VLSM). Intermediate System-to-Intermediate System (IS-IS) is also a classless routing protocol. An example of a network using VLSM is shown below. Note the different masks used, indicated with CIDR notation.



References:

<http://www.ciscopress.com/articles/article.asp?p=330807&seqNum=4&rl=1>
http://docwiki.cisco.com/wiki/Routing_Information_Protocol

QUESTION 28

The following is a partial output of the show interfaces command:

```

Serial 0 is up, line protocol is down
Hardware is MCI Ethernet, address is 0000.0c00.750c (bia 0000.0c00.750c)
Internet address is 134.108.28.8, subnet mask is 255.255.255.0
MTU 1500 bytes, BW 10000 Kbit, DLY 100000 usec, rely 255/255, load 1/255
Encapsulation ARPA, loopback not set, keepalive set (10 sec)
ARP type: ARPA, ARP Timeout 4:00:00
Last input 0:00:00, output 0:00:00, output hang never
Last clearing of "show interface" counters 0:00:00
Output queue 0/40, 0 drops; input queue 0/75, 0 drops
Five minute input rate 0 bits/sec, 0 packets/sec
Five minute output rate 2000 bits/sec, 4 packets/sec
1127576 packets input, 447251251 bytes, 0 no buffer
Received 354125 broadcasts, 0 runs, 0 giants, 57186* throttles
0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
  
```

5332142 packets output, 496316039 bytes, 0 underruns
0 output errors, 432 collisions, 0 interface resets, 0 restarts
<<output omitted>>

What does the Serial 0 is up, line protocol is down statement signify in the output? (Choose all that apply.)

- A. the shutdown interface command is present in the router configuration
- B. a cable is unplugged
- C. the interface is displaying normal operation
- D. there are no problems with physical connectivity
- E. there is a configuration problem in the local or remote router

Correct Answer: DE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Serial 0 is up, line protocol is down statement in the output signifies the following:

- There are no problems with the physical connectivity.
- There is a configuration problem in the local or remote router.
- The remote router might not be sending the keep-alives.
- There may be a problem with the leased lines such as line noise and a malfunctioning switch.
- There is an incorrect configuration of the CSU/DSU, which can cause timing issues on the cable. - The local or remote CSU/DSU might have failed.

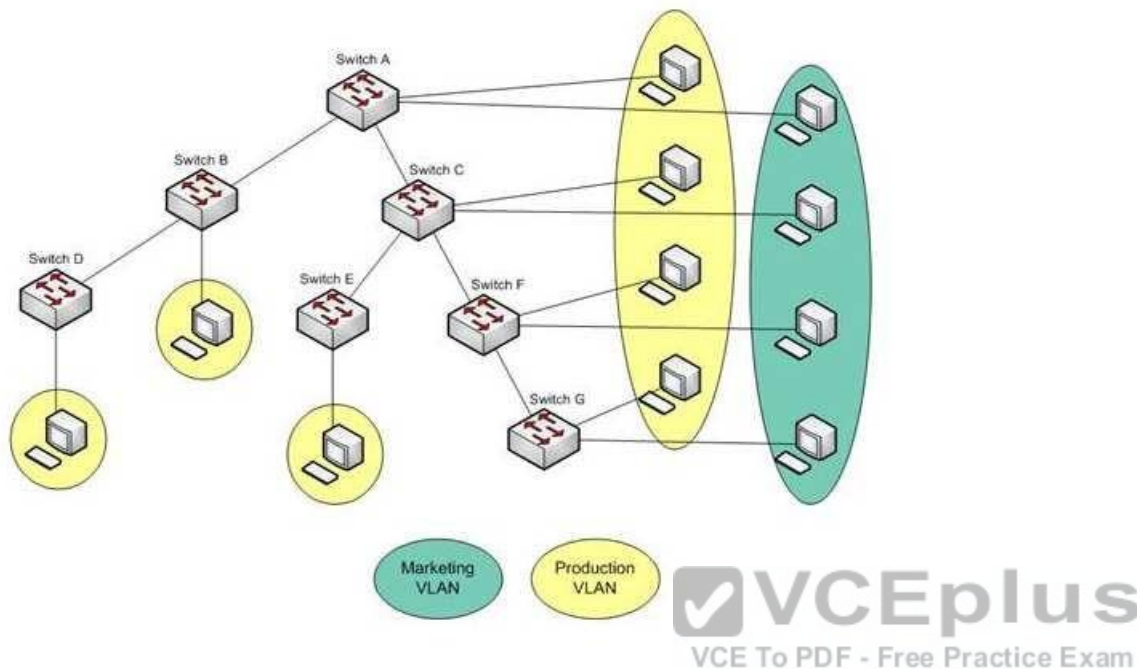
The option stating that the shutdown interface command is present in the router configuration is incorrect because if the shutdown interface command is present in the router configuration, the message displayed would be Serial 0 is administratively down, line protocol is down.

The option stating that a cable is unplugged is incorrect because that would be indicated by Serial 0 is down, line protocol is down. Physical problems such as a bad cable or cable unplugged are addressed in the first part of the output (serial0 is up/down).

The option stating that the message refers to normal operation of the interface is incorrect because the line protocol is shown as down, which indicates a problem.

QUESTION 29

You are a network administrator for your organization. Your organization has two Virtual LANs (VLANs) named Marketing and Production. All switches in the network have both VLANs configured on them. Switches A, C, F, and G have user machines connected for both VLANs, while switches B, D, and E have user machines connected for the Production VLAN only. (Click the Exhibit(s) button to view the network diagram.)



To reduce broadcast traffic on the network, you want to ensure that broadcasts from the Marketing VLAN are flooded only to those switches that have Marketing VLAN users.

Which Cisco switch feature should you use to achieve the objective?

- A. PVST
- B. RSTP
- C. VTP Pruning
- D. Dynamic VLANs

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The VLAN Trunking Protocol (VTP) pruning feature of Cisco VTP allows switches to dynamically delete or add VLANs to a trunk. It restricts unnecessary traffic, such as broadcasts, to only those switches that have user machines connected for a particular VLAN. It is not required to flood a frame to a neighboring switch if that switch does not have any active ports in the source VLAN. A trunk can also be manually configured with its allowed VLANs, as an alternative to VTP pruning.

All other options are incorrect because none of these features can be used to achieve the objective in this scenario.


The Per-VLAN Spanning Tree (PVST) feature allows a separate instance of Spanning Tree Protocol (STP) per VLAN. Each VLAN will have its own root switch and, within each VLAN, STP will run and remove loops for that particular VLAN.

Rapid Spanning Tree Protocol (RSTP) is an Institute of Electrical and Electronics Engineers (IEEE) standard. It reduces high convergence time that was previously required in STP implementations. It is interoperable with STP (802.1d).

With dynamic VLANs, the switch automatically assigns a switch port to a VLAN using information from the user machine, such as its Media Access Control (MAC) address or IP address. The switch then verifies information with a VLAN Membership Policy Server (VMPS) that contains a mapping of user machine information to VLANs.

QUESTION 30

You are the Cisco administrator for Metroil. One of your assistants has submitted the given diagram as a potential addressing plan for two offices. Both offices use EIGRP as the routing protocol. You immediately see a problem with the proposal. Which of the following actions could be a solution? (Choose two. Each correct option is a complete solution.)



- A. Execute the no auto-summary command on both routers.
- B. Change the network on F0/0 of Router A to 192.168.3.0/24 and change the network on F0/0 of Router B to 192.168.2.0/24.
- C. Change the network on F0/0 of Router B to 192.168.1.48/28.
- D. Execute the auto-summary command on both routers.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should either execute the no auto-summary command on both routers, or change the network on F0/0 of Router A to 192.168.3.10/24 and the network on F0/0 of Router B to 192.168.2.0/24. The exhibit is an example of discontinuous subnets, in which two subnets (192.168.1.16 and 192.168.1.32) of the same major network (192.168.1.0) are separated by a completely different network (172.16.0.4/30). The no auto-summary command instructs EIGRP to stop automatically summarizing advertised networks to their classful boundaries. Without the no auto-summary command, EIGRP will automatically summarize these two subnets to 192.168.1.0, and advertise the summary route across the WAN link, losing the subnet-specific information and causing routing problems. The no auto-summary command stops this behavior, and allows EIGRP to advertise specific subnets.

An alternate solution would be to change the network on F0/0 of Router A to 192.168.3.0/24 and the network on F0/0 of Router B to 192.168.2.0/24. If that were done, the two networks would be in separate class C networks and auto summarization would not be a problem.

It would not help to change the network on F0/0 of Router B to 192.168.1.48/28. The two networks would still be in the same class C network and the summarization process would confuse routing.

It would not help to execute the auto-summary command. The command is already in effect by default.

References: <https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16406-eigrp-toc.html>

QUESTION 31

Your network consists of one HSRP group of six routers. All of the routers are functioning properly. The network has been stable for several days. In which HSRP state are most of the routers?

- A. Learn
- B. Listen
- C. Standby
- D. Active

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

If all of the routers in the Hot Standby Routing Protocol (HSRP) group are functioning properly, then most of the routers in the group are in the listen state. Four routers will be in the listen state, one router will be in the standby state, and one router will be in the active state.

HSRP is used by a group of routers to create the appearance of a virtual router with which end stations can communicate in the event that the default gateway becomes unavailable. The active router is responsible for forwarding packets that are sent to the virtual router. The standby router is responsible for assuming the role of active router should the active router fail or become unavailable. All other HSRP routers monitor the hello messages sent by the active and standby routers. Should the active and standby routers both become unavailable, the HSRP router with the highest priority is elected to become the active router by default. For routers with equal priority values, the router with the highest IP address becomes the active router.

HSRP routers can exist in one of the following six states:

- Initial
- Learn
- Listen
- Speak
- Standby
- Active

All HSRP routers start in the initial state. A router in the learn state is waiting for its first hello message from the active router so that it can learn the virtual router's IP address. When the hello message is received and the virtual router's IP address is discovered, the HSRP router is in the listen state. A router in the listen state listens for hello messages from the active and standby routers. If an election for a new active router and a new standby router is required, then an HSRP router will enter the speak state and begin transmitting hello messages. The standby state is reserved for the standby router, and the active state is reserved for the active router. Only routers in speak, standby, and active states will transmit hello packets.

References: <https://www.cisco.com/c/en/us/support/docs/ip/hot-standby-router-protocol-hsrp/9234-hsrpguidetoc.html> <https://www.cisco.com/c/en/us/products/index.html>

QUESTION 32

DRAG DROP

Click and drag the RSTP port state on the left to its matching equivalent STP role, on the right. RSTP port states may be used more than once, and it may not be necessary to use all RSTP port states.

Select and Place:

RSTP Port State

Discarding
Learning
Forwarding

STP Role

	Blocking
	Listening
	Forwarding
	Learning
	Disabled

Correct Answer:

RSTP Port State

Discarding
Learning
Forwarding

STP Role

Discarding	Blocking
Discarding	Listening
Forwarding	Forwarding
Learning	Learning
Discarding	Disabled

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Rapid Spanning Tree Protocol (RSTP) was developed to reduce the high convergence times required in Spanning Tree Protocol (STP), and introduces the alternate port and backup port. RSTP is an Institute of Electrical and Electronics Engineers (IEEE) standard, 802.1w, and is interoperable with 802.1d (STP). There are fewer transitional states used in RSTP than STP. In RSTP, there are only Forwarding, Learning, and Discarding. The three states are defined as follows:

- Forwarding - the state of all root ports and designated ports. The port is passing traffic.
- Learning - the state of a port that was formerly discarding but due to a change in the topology (link down) it has transitioned to learn its new state. The port could return to discarding or move to forwarding depending on the new topology needs
- Discarding - the state of all non-root and non- designated ports. The port is not passing traffic to prevent potential switching loops.

RSTP can reconfigure the spanning tree in less than a second, compared to the 50 seconds that STP may take. This is achieved through having fewer transition states, the use of alternate and backup ports, and faster transitions.

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

QUESTION 33

What is the significance of the 1 in the following configuration?

```
router(config)# router eigrp 1
```

- A. It is the process ID for EIGRP and is locally significant to this router.
- B. It is the process ID for EIGRP and must be the same on all EIGRP routers.
- C. It is the AS number for EIGRP and is locally significant to this router.
- D. It is the AS number for EIGRP and must be the same on all EIGRP routers.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Enhanced Interior Gateway Routing Protocol (EIGRP) configuration requires the specification of an Autonomous System (AS) number with the router eigrp command. Any number can be chosen, but it must match on all EIGRP routers in the domain. This value may appear to be similar to one used in enabling OSPF, which demands a process ID number but that value is locally significant to each router and need not match on each router.

The syntax of this command is router eigrp [autonomous-system]. Therefore, the 1 in the example indicates an Autonomous System (AS) number, not a process ID.

The Autonomous System (AS) number is not locally significant to each router, and must match on all EIGRP routers.

QUESTION 34

You are in the process of verifying the operation of your core switches, which are using HSRP. One core switch was left with the default priority; the other was given a lower priority to make it the standby switch. The command show standby brief was executed on one of the switches. Output of the command is shown below:

Interface	Grp	Prio	P State	Active	Standby	Virtual IP
Vl10	1	90	P Active	local	192.168.10.20	192.168.10.1
Vl20	1	90	P Active	local	192.168.20.20	192.168.20.1

What does this output mean? (Choose all that apply.)

- A. this switch is using the default priority
- B. this switch is the active HSRP switch
- C. the HSRP devices are up and functioning correctly
- D. the switch intended to be the active switch has failed and this switch has taken over
- E. preemption is enabled for the group

Correct Answer: BDE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The output in the exhibit indicates that this switch is the active HSRP switch, the switch intended to be the active switch has failed, and that preemption is enabled for the group.

This is the active switch because Active is the State listed for each interface that is a member of HSRP.

The question states that the switch that was intended to be the standby switch was given a priority lower than the default. The default priority is 100, so this is not the switch intended to be the active switch. This information indicates that the switch intended to be the active switch has failed.

Preemption is enabled, as indicated by the P following the priority value in line 2. Since preemption is enabled, the switch with the priority of 100 is still down. When that switch is corrected and joins the group again, it will take over as active.

The HSRP group is still providing access for users, but not all devices are functioning properly.

References:

https://www.cisco.com/c/en/us/td/docs/ios/ipapp/command/reference/iap_s4.html

QUESTION 35

Which command enables HSRP on an interface?

- A. hsrp
- B. standby ip

- C. standby mode hsrp
- D. switchport mode hsrp

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The standby ip interface configuration command enables Hot Standby Router Protocol (HSRP). The syntax for this command is as follows:

```
switch(config-if)# standby group-number ip ip-address
```

The group-number argument specifies the HSRP group number on the interface. You do not need to enter a group number if there is only one HSRP group.

At least one interface on one of the routers in the group must be configured with the virtual IP address of the group. It is optional on all other interfaces on the other routers, which can learn the address through the hellos sent among the group.

A complete HSRP configuration is shown below with an explanation of each command.

```
RouterA (config) #interface Fa0/1
```

```
RouterA (config-if) # ip address 192.168.5.6 255.255.255.0
```

```
RouterA (config-if) # standby 2 ip 192.168.5.10      RouterA(config-if) # standby 2 priority 150
```

```
RouterA (config-if) #standby 2 Preempt
```

```
RouterA(config-if) #standby 2 track interface fa0/2
```

-Line 1 specifies the interface

-Line 2 addresses the interface

-Line 3 specifies the HSRP group number and the virtual IP address -

Line 4 sets the HSRP priority

-Line 5 allows the router to take the active role if its priority becomes higher than that of the active router

In the above, the router is tracking its own Fa0/2 interface. If that interface goes down it will reduce its priority by 10 (this is the default decrement when not specified). The new value would be 140 if that happened. To specify a decrement value, add it to the track command, as in this example: track interface Fa0/2 20.

When you configure routers to be part of an HSRP group, they listen for the HSRP MAC address for that group as well as their own burned-in MAC addresses.

HSRP uses the following MAC address:

0000.0c07.ac** (where ** is the HSRP group number)

The switchport mode interface configuration command will configure the VLAN membership mode of a port. It is not used to enable HSRP.

The options standby mode hsrp and hsrp are not valid commands.

References:

<https://www.cisco.com/c/en/us/support/docs/ip/hot-standby-router-protocol-hsrp/9234-hsrpguidetoc.html>

<https://www.cisco.com/c/en/us/products/index.html>

QUESTION 36

Which of the following technologies allows a switch port to immediately transition to a forwarding state?

- A. Rapid STP
- B. PortFast
- C. VTP
- D. CDP

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

PortFast is a technology that allows a switch port connected to an end node such as a workstation, server, or printer to bypass the normal Spanning Tree Protocol (STP) convergence process. When a new device is powered up on a switch port, it will immediately transition to a forwarding state.

NOTE: PortFast should only be used on access ports. It should not be used on trunk ports or on ports that connect to hubs, routers and other switches.

Rapid STP (RSTP) is a new STP standard that provides faster convergence than the original 802.1d STP. RSTP supports PortFast, but it must be configured explicitly.

The VLAN Trunking Protocol (VTP) does not allow for immediate transition to a forwarding state. VTP is used to synchronize VLAN databases between switches, and has no effect on STP.

The Cisco Discovery Protocol (CDP) does not allow for immediate transition to a forwarding state. CDP is used to verify connectivity and document directly connected Cisco devices. CDP is not related to STP.

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

QUESTION 37

Which switch will be selected as the root bridge by Spanning Tree Protocol (STP)?

- A. switch with lowest bridge ID
- B. switch with lowest IP address
- C. switch with lowest Media Access Control (MAC) address
- D. switch with lowest number of root ports

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:


STP will use elections to arrive at a fully converged state that will ensure a switching loop free network. It will select:

- The root bridge
- The root port on each non-root bridge
- Designated ports on any shared segments with no direct connection to the root bridge.

The switch with the lowest bridge ID will be selected as the root bridge by STP. A bridge ID has two components: the priority number and the MAC address. On Cisco devices, the priority number may range from 0 to 65535. The priority number constitutes the most significant bits of the bridge ID. If you want to ensure that a particular switch in a topology always becomes a root bridge, regardless of the MAC address, you can set the priority number of that switch to the lowest value among all switches in the topology.

Since the selection of the root bridge influences all other decisions and thus the single loop free path for each VLAN, the selection and location of the root bridge is important and best not left to chance. Once you have determined the best switch for the role of root bridge, you can ensure its election by lowering its bridge priority.

It is best for the root bridge to be centrally located with respect to the clients and the servers that generate the most traffic on the VLAN. For example, in the diagram below, if most of the traffic travels between the clients and the servers on VLAN 20, the best choice for the root bridge for VLAN 20 would be SwitchD. SwitchD is centrally located between the clients on VLAN 20 and the servers on VLAN 20.



To illustrate the type of inefficient traffic that could occur when care is not given to the location of the root bridge, consider the diagram above and assume that Switch B was chosen the root bridge. Next, assume that traffic needs to go from VLAN 10 connected to Switch C to VLAN 10 connected to Switch A. The shortest path would be from Switch C to Switch A. However, because the only port that is forwarding on Switch C is the port that leads to the root bridge (Switch B), then the actual path would be from Switch C, to Switch B, to Switch E, and then to Switch A.

By default, the priority number of all Cisco switches is configured to a value of 32768. For example, consider three switches in network topology with the following MAC addresses and the same default priority number:

0000.0B02.AAAA
0000.0B02.BBBB
0000.0B02.CCCC

The switch with the lowest MAC address, 0000.0B02.AAAA, will become the root bridge.

The switch with the lowest IP address will not be selected as the root bridge by STP because the IP address of the switch does not influence the selection of the root bridge.

The switch with the lowest MAC address will not be selected as the root bridge by STP. A combination of priority number and MAC address determines the selection of the root bridge. The MAC address will determine the root bridge only if there is a tie for the switch with the lowest priority number.

The switch with the lowest number of root ports will not be selected as the root bridge by STP. Root ports are the interfaces on non-root bridges. On a non-root bridge, the least-root-cost interface is known as a root port. Therefore, the switch having the fewest root ports is not the root bridge.


References:

https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst6500/ios/12-2SX/configuration/guide/book/port_sec.html

<https://www.cisco.com/c/en/us/td/docs/routers/7600/ios/12-2SXF/configuration/guide/swcg/spantree.html>

QUESTION 38

You are the Cisco administrator for Verigon Incorporated. The given exhibit displays some of the devices in the network. (Click the Exhibit(s) button.) Workstation A can communicate with Workstation C but cannot communicate with Workstation B.



What is the problem?

- A. Workstation B has an incorrect default gateway
- B. Workstation A has an incorrect subnet mask
- C. Workstation A has an incorrect default gateway
- D. Workstation B has an incorrect subnet mask

Correct Answer: C

Section: (none)

Explanation**Explanation/Reference:**

Explanation:

Workstation A has an incorrect default gateway. To communicate with remote computers or those computers outside of its own subnet, a computer must have the address of the nearest router interface as its default gateway. In this case, the default gateway of Workstation A should be 192.168.10.5/24, which is the Serial0 address of Router A. The diagram shows that it is instead configured as 192.168.10.4/24. This will not cause a problem for Workstation A to communicate with Workstation C, but it will make communication with remote subnets impossible.

Workstation B does not have an incorrect default gateway. Its nearest router interface is 10.0.0.1/8, which is the configuration of its default gateway.

Workstation A does not have an incorrect subnet mask. The mask used by Workstation C and the router interface of Router A, which are in the same subnet, is /24, or 255.255.255.0, which is also the subnet mask used by Workstation A.

Workstation B does not have an incorrect subnet mask. Since the subnet mask of the router interface that is nearest to Workstation B is /8, or 255.0.0.0, then Workstation B also should have an 8 bit mask.

References:

<https://www.cisco.com/c/en/us/support/docs/ip/routing-information-protocol-rip/13788-3.html>

<https://www.cisco.com/c/en/us/support/docs/ip/dynamic-address-allocation-resolution/13711-40.html>

CCNA ICND2 Official Exam Certification Guide (Cisco Press, ISBN 1-58720-181-X), Appendices D, E and H: Subnetting.

QUESTION 39

Which of the following statements is NOT true of Cisco ACI?

- A. It is a comprehensive SDN architecture.
- B. It uses Cisco APIC as the central management system.
- C. It provides policy driven automation support.
- D. It decreases network visibility.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Cisco ACI does not decrease network visibility. On the contrary, the Cisco Application Centric Infrastructure (ACI) increases network visibility. It is a policydriven automaton solution that can keep the network inventory up-to-date automatically whenever a new device is added and provide a graphic representation at all times.

ACI is a comprehensive SDN architecture that integrates physical and virtual environments under one policy model. It uses the Cisco Application Policy Infrastructure Controller (APIC) as the central management system. It provides policy driven automation support through a business-relevant application policy language.

References: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/1-x/aci-fundamentals/b_ACI-Fundamentals/b_ACI-Fundamentals_chapter_010000.html

QUESTION 40

Which technique is used to stop routing loops by preventing route update information from being sent back over the interface on which it arrived?

- A. Holddown timer
- B. Triggered updates
- C. Route poisoning
- D. Split horizon
- E. Maximum hop count

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Split horizon stops routing loops by preventing route update information from being sent back over the interface on which it arrived. Routing loops can occur due to slow convergence and inconsistent routing tables, and can cause excessive use of bandwidth or even complete network failure. Split horizon can prevent routing loops between adjacent routers.

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.

Triggered updates are sent as soon as a change in network topology is discovered, as opposed to waiting until the next regular update interval (every 30 seconds in RIP networks). This speeds convergence and helps prevent problems caused by outdated information.

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 41

You are planning the configuration of an IPsec-protected connection between two routers. You are concerned only with the integrity of the data that passes between the routers. You are less concerned with the confidentiality of the data, and you would like to minimize the effect of IPsec on the data throughput. 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: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should choose Authentication Header (AH) in tunnel mode to meet the scenario requirements. 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 optionally to provide anti-reply service. 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. 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.

You would not choose Authentication Header (AH) in transport mode. Transport mode is used between end stations or between an end station and a VPN gateway.

You would not choose Encapsulating Security Payload (ESP) in tunnel mode or transport mode. Using ESP will slow the connection because of the encryption and decryption process that will occur with each packet.

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> CCNA ICND2 Official Exam Certification Guide (Cisco Press, ISBN 1-58720-181-X), Chapter 15: Virtual Private Networks, pp. 536-537.

QUESTION 42

You are the Cisco administrator for NationalAct Incorporated. One of your assistants is preparing to introduce a new switch to the network. Before doing so, you execute the show vtp status command on OldSwitch and NewSwitch, respectively, and receive the following output:

```
OldSwitch# show vtp status
VTP Version : 2
Configuration Revision : 62
Maximum VLANs supported locally : 1005
Number of existing VLANs : 24
VTP Operating Mode : Server
VTP Domain Name : Corporate
VTP Pruning Mode : Enabled
VTP V2 Mode : Disabled
VTP Traps Generation : Disabled
<output omitted>
```

```
NewSwitch# show vtp status
VTP Version : 2
Configuration Revision : 125
Maximum VLANs supported locally : 1005
Number of existing VLANs : 10
VTP Operating Mode : Server
VTP Domain Name : Corporate
VTP Pruning Mode : Enabled
VTP V2 Mode : Disabled
VTP Traps Generation : Disabled
<output omitted>
```



If NewSwitch is introduced to the network, which of the following will be true?

- A. NewSwitch will delete its current VTP data.
- B. There will be 10 VLANs in the network.
- C. OldSwitch will retain its current VTP data.
- D. There will be 24 VLANs in the network.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

If NewSwitch is introduced to the network, there will be 10 VLANs. The VLAN database of the new switch will overwrite the VLAN databases of the production switches because it is operating in server mode and has a higher VLAN configuration revision number.

VLAN Trunking Protocol (VTP) is used to synchronize VLANs between different switches. The VTP configuration revision number is used to determine which VTP switch has the most current version of the VLAN database, and is incremented whenever a VLAN change is made on a VTP server switch. The Configuration Revision: 125 output indicates that NewSwitch has a configuration revision number of 125, which will be compared to other switches in the same VTP domain, including OldSwitch, which has a revision number of 62. If the production switches have lower configuration revision numbers than the new switch, their VLAN databases will be replaced with the VLAN database of the new switch. Any switch ports that had been assigned to be removed from VLANs in the configuration database of the new switch will be disabled, possibly resulting in catastrophic network failure. All VTP switches in the same VTP domain should have a domain password defined, which will protect against a rogue switch being added to the network and causing VLAN database corruption.

NewSwitch will not delete its current VTP data. If the production switches have lower configuration revision numbers than the new switch, their VLAN databases will be replaced with the VLAN database of the new switch.

The number of VLANs will not remain 24. The 24 VLANs indicated by the Number of existing VLANs: 24 output will be overwritten with the 10 VLANs in the NewSwitch VLAN database.

OldSwitch will not retain its current VTP data. It will be replaced with the VLAN database of the new switch.

References:

<https://www.cisco.com/c/en/us/support/docs/lan-switching/vtp/98154-conf-vlan.html> <https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4500/12-2/25sg/configuration/guide/conf/vlans.html>

QUESTION 43

Which of the following is NOT a benefit of cloud computing to cloud users?

- A. On-demand self-service resources provisioning
- B. Centralized appearance of resources
- C. Highly available, horizontally scaled applications
- D. Cost reduction from standardization and automation

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Cost reduction from standardization and automation is a benefit that accrues to the cloud provider, not the cloud users. Additional benefits to cloud providers are:

- High utilization through virtualization and shared resources
- Easier administration
- Fail-in-place operations model

Benefits that accrue to cloud users include:

- On-demand self-service resources provisioning
- Centralized appearance of resources
- Highly available, horizontally scaled applications - No local backups required

Cloud users can also benefit from new services such as intelligent DNS, which can direct user requests to locations that are using fewer resources.

References: <https://www.cisco.com/c/en/us/products/cloud-systems-management/benefit.html>

QUESTION 44

Which of the following is NOT true of the Cisco APIC-EM?

- A. It can verify the operation of access lists
- B. It provides network topology visualization
- C. It can perform identity tracking
- D. It is appropriate for the datacenter environment



Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

With all of its benefits, the Cisco APIC-EM is not appropriate for the datacenter environment. A more appropriate controller for the datacenter environment is Cisco APIC-DC. Both of these are software-defined network controllers, which can be used to program a network in an automated fashion.

Specific benefits provided by the Cisco APIC-EM include:

- It can verify the operation of access lists with the Path Trace Analysis tool
- It provides network topology visualization
- It can perform identity tracking
- It provides an inventory of devices- It automatically adds new devices

References:

https://www.cisco.com/c/en/us/td/docs/cloud-systems-management/application-policy-infrastructure-controller-enterprise-module/1-2-x/config-guide/b_apicem_config_guide_v_1-2-x/b_apic-em_config_guide_v_1-2-x_chapter_01000.pdf

QUESTION 45

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:

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. For example, if a router had three interfaces with the addresses 192.168.5.1/24, 192.168.6.1/24, and 192.168.7.1/24, and you executed the command `network 192.168.0.0 0.0.0.255.255`, all three of the subnets would be advertised and would be present in the neighboring router's routing table.

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-cfg.html

QUESTION 46

Which of the following statements are TRUE regarding EIGRP operation? (Choose two.)

- A. A successor is a backup route, and is installed in both the routing and topology tables.
- B. A successor is a primary route, and is installed in both the routing and topology tables.
- C. A successor is a primary route, and is installed only in the routing table.
- D. A feasible successor is a backup route, and is installed in both the routing and topology tables.
- E. A feasible successor is a primary route, and is only installed in the routing table.
- F. A feasible successor is a backup route, and is only installed in the topology table.
- G. If the successor route fails and no feasible successor route exists, the router will send an update with the route marked with an unreachable metric of 16.

Correct Answer: BF

Section: (none)

Explanation

Explanation/Reference:

Explanation:

In EIGRP operations, primary or active routes are known as successors. These routes are maintained in both the routing and topology tables. The routing table is the list of network paths that are currently used by the router.

EIGRP also has the ability to maintain backup routes to destination networks. These backup routes are known as feasible successors. If a feasible successor is discovered by EIGRP, it will be maintained only in the topology table, since it is not currently being used to route traffic. In the event of a successor failure, the backup feasible successor will become the successor, and will be installed in the routing table automatically. If the successor route fails and no feasible successor route exists, the router will send queries to all neighbors until a new successor is found.

EIGRP maintains three dynamic tables in RAM:

- Neighbor table, which is a list of all neighboring EIGRP routers on shared subnets
- Topology table, which contains all discovered network paths in the internetwork
- Routing table, which contains the best path (based on lowest metric) to each destination network

A successor is not a backup route. A successor is a primary or active route, and it is stored in both the routing and topology tables.

A feasible successor is not a primary route. It is a backup route, and it is stored only in the topology table.

If the successor route fails and no feasible successor route exists, the router will not send an update with the route marked with an unreachable metric of 16. EIGRP does not send an update with the route marked with an unreachable metric, and even if it did, 16 is not an unreachable metric in EIGRP as it is in RIP. Instead it sends a multicast query packet to all adjacent neighbors requesting available routing paths to the destination network.

References: <https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16406-eigrp-toc.html>

QUESTION 47

Which command is NOT mandatory for inclusion in a plan to implement IP Service Level Agreements (SLAs) to monitor IP connections and traffic?

- A. ip sla
- B. ip sla schedule
- C. ip sla reset
- D. icmp-echo

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The ip sla reset command is not mandatory for an implementation plan to configure IP SLAs for monitoring IP connections and traffic. This command causes the IP SLA engine to either restart or shutdown. As a result, all IP SLAs operations are stopped, IP SLA configuration information is erased, and IP SLAs are restarted. The IP SLAs configuration information will need to be reloaded to the engine.

The following commands are essential to the implementation plan:

```
ip sla ip sla  
schedule  
icmp-echo
```

The ip sla command allows you to configure IP SLAs operations. When you execute this command in the global configuration mode, it enables the IP SLA configuration mode. In the IP SLA configuration mode, you can configure different IP SLA operations. You can configure up to 2000 operations for a given IP SLA ID number.

The icmp-echo command allows you to monitor IP connections and traffic on routers by creating an IP SLA ICMP Echo operation. This operation monitors end-to-end response times between routers.

The ip sla schedule command allows you to schedule the IP SLA operation that has been configured. With this command, you can specify when the operation starts, how long the operation runs, and how long the operation gathers information. For example, if you execute the ip sla schedule 40 start-time now life forever command, the IP SLA operation with the identification number 40 immediately starts running. This is because the now keyword is specified for the start-time parameter. Using the forever keyword with the life parameter indicates that the operation keeps collecting information indefinitely. Note that you cannot re-configure the IP SLA operation after you have executed the ip sla schedule command.

The information gathered by an IP SLA operation is typically stored in RTTMON-MIB. A Management Information Base (MIB) is a database hosting information required for the management of routers or network devices. The RTTMON-MIB is a Cisco-defined MIB intended for Cisco IOS IP SLAs. RTTMON MIB acts as an interface between the Network Management System (NMS) applications and the Cisco IOS IP SLAs operations.

References:

<https://www.cisco.com/c/en/us/support/docs/ip/enhanced-interior-gateway-routing-protocol-eigrp/16406-eigrp-toc.html>
https://www.cisco.com/en/US/technologies/tk648/tk362/tk920/technologies_white_paper09186a00802d5efe.html
https://www.cisco.com/c/en/us/td/docs/ios/ipsla/command/reference/sla_book/sla_02.html
https://www.cisco.com/c/en/us/td/docs/ios/ipsla/command/reference/sla_book/sla_02.html
https://www.cisco.com/c/en/us/td/docs/ios/ipsla/command/reference/sla_book/sla_02.html

QUESTION 48

When a router has been configured with a loopback address, which of the following determines the OSPF router ID?

- A. The highest MAC address assigned to a physical interface on the router
- B. The lowest priority of a physical interface on the router
- C. The lowest IP address assigned to a physical interface on the router
- D. The highest IP address assigned to a loopback interface on the router

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Routers configured with OSPF must be assigned a router ID (RID), which is an IP address unique across the entire OSPF autonomous system. The RID can be assigned manually with the router-id command, or it will be determined automatically by OSPF. If the RID has not been manually assigned, then OSPF will use the highest numerical IP address of a loopback interface on the local router. If there are no configured loopback interfaces, then the RID will be determined by the highest numerical IP address on an active physical interface. The sequence for determining the RID is as follows:

1. Any address manually configured with the router-id command
2. The highest IP address on a loopback interface
3. The highest IP address on an active physical interface

Either of the first two options would be a recommended best practice, since they each offer fault tolerance to the RID. If the RID is determined by a physical interface IP address, then the entire OSPF routing process is bound to an interface that could become unplugged or go down due to network reasons.

Loopback interfaces remain operational unless they are manually shut down. Loopback interfaces are configured as follows:

```
Router(config)# interface loopback0
```

```
Router(config-if)# ip address 192.168.1.254 255.255.255.255
```

The highest media access control (MAC) address assigned to a physical interface on the router is not used. IP addresses are used for the determination of the router ID.

Priorities are not used to determine the OSPF router ID. Priorities are used by OSPF to influence the election of the designated router (DR) and backup designated router (BDR) on a multi-access segment.

Router IDs are determined by the highest IP address on a loopback or physical interface, not the lowest.

References:

CCNA ICND2 Official Exam Certification Guide (Cisco Press, ISBN 1-58720-181-X), Chapter 9: OSPF, pp. 366-367.

QUESTION 49

What is the HSRP virtual router MAC address for the virtual router for HSRP group 31?

- A. 0000.0c07.ac1f
- B. ac1f
- C. 0c07
- D. 07.ac

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Hot Standby Router Protocol (HSRP) virtual MAC address for the virtual router for HSRP group 31 is 0000.0c07.ac1f. A Media Access Control (MAC) address is a 6-byte value that is unique for every networked device. MAC addresses are typically written in hexadecimal notation. The address 0000.0c07.ac1f is a MAC address for an HSRP virtual router; this address can also be written as 00-00-0c-07-ac-1f or 00.00.0c.07.ac.1f. Hexadecimal letters can be written as either lowercase or uppercase letters.

The MAC address for an HSRP virtual router consists of the vendor ID, the HSRP code and the group ID. The vendor ID corresponds to the first three bytes of the MAC address. A vendor ID of 0000.0c indicates that the device was manufactured by Cisco. The HSRP code corresponds to the fourth and fifth bytes of the MAC address. The HSRP code for a virtual router is always equal to 07.ac. Finally, the group ID corresponds to the last byte of the MAC address. For example, a group ID of 1f, when converted to decimal, indicates that the virtual router belongs to HSRP group 31.

QUESTION 50

On Cisco switches, what is the correct order of port transition through the Spanning Tree Protocol (STP) states?

- A. learning, listening, blocking, forwarding
- B. listening, blocking, forwarding, learning
- C. blocking, learning, forwarding, listening
- D. blocking, listening, learning, forwarding

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

There are five states in STP transition:

- Blocking
- Listening
- Learning

- Forwarding
- Disabled

After STP initialization, a port moves from blocking to listening, then to learning, and finally into forwarding state. In case of any errors or exceptions, a port may enter into a disabled state directly from any of the other four states. Once STP has fully converged, all ports on all switches will be in either a forwarding state or a blocking state. All other port states are transitioning states between blocking and forwarding.

When STP is initialized, all ports start in the blocking state to prevent bridge loops. If a switch determines that a blocking port must transition to a forwarding state, the blocked port will first move into a listening state, where it begins sending Bridge Protocol Data Units (BPDUs). Next, the port will transition to a learning state, which allows it to populate its Media Access Control (MAC) address table with addresses learned on the port, but it does not yet forward data frames. Finally, it moves into the forwarding state, where the port is capable of sending and receiving data. The switch only learns MAC addresses during the learning and forwarding states.

QUESTION 51

You are configuring SPAN so you can connect a sniffer to your switch. Which of the following is NOT true with regard to the source port in the configuration?

- A. It can be an EtherChannel
- B. It can also be the destination port
- C. It can be monitored in multiple SPAN sessions
- D. It can monitor both ingress and egress traffic

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The source port in a SPAN configuration cannot also be a destination port. The selected destination port will no longer operate as a normal switch port. It will only pass the traffic redirected from the source port. Therefore, there would be no "source" traffic if it were a destination port as well.

Source ports in a SPAN configuration have the following characteristics:

- It can be any port type, such as EtherChannel, Fast Ethernet, Gigabit Ethernet, and so forth.
- It can be monitored in multiple SPAN sessions.
- It cannot be a destination port.
- Each source port can be configured with a direction (ingress, egress, or both) to monitor. - Source ports can be in the same or different VLANs.
- For VLAN SPAN sources, all active ports in the source VLAN are included as source ports.

References: <https://www.cisco.com/c/en/us/support/docs/switches/catalyst-6500-series-switches/10570-41.html>

QUESTION 52

Which of the following is NOT a packet type used by Enhanced Interior Gateway Routing Protocol (EIGRP)?

- A. Query
- B. Reply
- C. Ack
- D. Response

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Response is not a packet type used by EIGRP. The following are the packet types used by EIGRP:

- Hello/Ack: Establish neighbor relationships. The Ack packet is used to provide acknowledgement of a reliable packet. - Update: Send routing updates.
- Query: Ask neighbors about routing information.
- Reply: Provide response to queries about routing information.
- Requests: Gain specific information from one or more neighbors.

References:

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

QUESTION 53

Refer to the following sample output:

*: interface is up

IHQ: pkts in input hold queue IQD: pkts dropped from input queue

OHQ: pkts in output hold queue OQD: pkts dropped from output queue

RXBS: rx rate (bits/sec) RXPS: rx rate (pkts/sec)

TXBS: tx rate (bits/sec) TXPS: tx rate (pkts/sec)

TRTL: throttle count

Interface IHQ IQD OHQ OQD RXBS RXPS TXBS TXPS TRTL

* FastEthernet0/0 0 0 0 0 0 0 0 0

Serial0/0 0 0 0 0 0 0 0 0

FastEthernet0/1 0 0 0 0 0 0 0 0

Serial0/1 0 0 0 0 0 0 0 0

Which Cisco Internetwork Operating System (IOS) command produces this output?

- A. show interfaces
- B. show interfaces summary
- C. show interfaces serial fast-ethernet
- D. show interfaces fast-ethernet 0/0

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The show interfaces summary command will produce the given output. This command provides a summarized view of all interfaces configured on a device.

The show interfaces command is incorrect because this command does not produce the displayed output. This command is used to view information regarding statistics for specific interfaces. Without specifying an interface, a section for each interface will display, as in the example below for FastEthernet0:

FastEthernet0 is up, line protocol is down
Hardware is Fast Ethernet, address is 0019.e818.a3dd (bia 0019.e818.a3dd)
MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
Auto-duplex, Auto-speed
ARP type: ARPA, ARP Timeout 04:00:00
Last input never, output never, output hang never
Last clearing of "show interface" counters never
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue: 0/40 (size/max)
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
0 packets input, 0 bytes, 0 no buffer
Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
0 input packets with dribble condition detected

0 packets output, 0 bytes, 0 underruns
0 output errors, 0 collisions, 2 interface resets
0 babbles, 0 late collision, 0 deferred
0 lost carrier, 0 no carrier --
More

The show interfaces serial fast-ethernet command is incorrect because this is not a valid Cisco IOS command.

The show interfaces fast-ethernet 0/0 command is incorrect. Although it produces similar output, that output only relates to the FastEthernet 0/0 interface. An example of this output follows:

FastEthernet0 is up, line protocol is up
Hardware is Fast Ethernet, address is 0019.e818.a3dd (bia 0019.e818.a3dd)
MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
Auto-duplex, Auto-speed
ARP type: ARPA, ARP Timeout 04:00:00
Last input 00:00:00, output 00:00:00, output hang never
Last clearing of "show interface" counters never
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops:105
Queueing strategy: fifo
Output queue: 0/40 (size/max)
5 minute input rate 1530000 bits/sec, 201 packets/sec
5 minute output rate 673000 bits/sec, 173 packets/sec
404737363 packets input, 23875417953 bytes, 11 no buffer
Received 1206930011 broadcasts, 0 runts, 0 giants, 0 throttles
0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
0 input packets with dribble condition detected
401877661 packets output, 23875417953 bytes, 0 underruns
0 output errors, 576297 collisions, 0 interface resets
0 babbles, 0 late collision, 2174225 deferred
0 lost carrier, 0 no carrier
0 output buffer failures, 0 output buffers swapped out

Notice that the line of output that says FastEthernet0 is up, line protocol is up indicates that Layers 1 to 3 of the OSI Model are functioning correctly. Also, in the lower portion, there are no values in the error counters such as input errors, output errors, and so on. Finally, make note in line 8 where the interface is set to autosense both the duplex and the speed. Duplex and speed must be in agreement between the NIC on the host and the switch port.

QUESTION 54

Which of the following is NOT true of APIC-EM?

- A. It supports greenfield but not brownfield deployments
- B. It provides a single point for network automation
- C. It saves time and cost
- D. It is open and programmable

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Cisco Application Policy Infrastructure Controller Enterprise Module (APIC-EM) is an SDN controller platform that supports both greenfield implementations, which use no previous code and design from the ground up, and brownfield implementations, which incorporate existing code.

APIC-EM does provide a single point for network automation. This automation leads to both time and cost savings.

APIC-EM uses an open and programmable approach to devices, policies, and analytics.

References: <https://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/application-policy-infrastructure-controller-enterprise-module/datasheet-c78730594.html>

QUESTION 55

You instructed your assistant to add a new router to the network. The routers in your network run OSPF. The existing router, OldRouter, is configured as follows:

```
router ospf 1
network 192.168.5.0 0.0.0.255 area 0
network 192.168.10.0 0.0.0.255 area 0
```

The OldRouter interface that connects to NewRouter is 192.168.5.3/24. Your assistant shows you the configuration that will be implemented:

newrouter(config)# router ospf 1 newrouter(config-router)# network
192.168.5.0 255.255.255.0 area 0 What is wrong with this
configuration?

- A. The area ID is incorrectly configured.
- B. The wildcard mask is incorrectly configured.
- C. The network statement is incorrectly configured.
- D. The process ID number is incorrectly configured.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

When entering network statements for OSPF, a wildcard mask is used instead of a regular mask. Since the network connecting the two routers is a class C network, as shown by the address 192.168.5.0/24, the wildcard mask should be 0.0.0.255 rather than 255.255.255.0. With wildcard masks, the 0s octets must match, and the 255s octets do not have to match.

The area ID is correct. OldRouter is in area 0, so NewRouter should be as well. There must be an area 0 in an OSPF network. There can be multiple areas as well, but they must all connect to area 0. If non-0 areas cannot be directly connected to area 0, they must be configured with a virtual link across an area that does connect to the backbone (area 0).

The network statement is correct. The network between the routers is 192.168.5.0.

The process ID number is correct. The number is stated as OSPF 1 on OldRouter and OSPF 1 on NewRouter. They match in this case but that is not required. Process IDs are only locally significant.

References:

http://docwiki.cisco.com/wiki/Open_Shortest_Path_First

QUESTION 56

Which Enhanced Interior Gateway Routing Protocol (EIGRP) packet type is used for neighbor discovery?

- A. Hello

- B. Update
- C. Queries
- D. Replies

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Hello packets are used for neighbor discovery. These are sent as multicasts and do not require an acknowledgement.

Update packets are sent to communicate the routes used by a router to converge. When a new route is discovered or the convergence process is completed, updates are sent as multicast. During topology table synchronization, updates are sent as unicasts to neighboring peers.

Query packets are sent when a router performs route computation and cannot find a feasible successor. These packets are sent to neighboring peers asking if they have a feasible successor to the destination network.

Reply packets are sent in response of a query packet. These are unicast and sent to the originator of the query.

QUESTION 57

Which Wide Area Network (WAN) switching technology is used by Asynchronous Transfer Mode (ATM)?

- A. packet switching
- B. virtual switching
- C. circuit switching
- D. cell switching

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Cell switching is a WAN switching technology that is used by ATM. ATM is an International Telecommunication Union-Telecommunications (ITU-T) standard for transmission of data, voice, or video traffic using a fixed size frame of 53 bytes, known as cells. Out of these 53 bytes, the initial five bytes are header information and the rest 48 bytes is the payload.

Packet switching is incorrect because packet switching is popularly used for data transfer, as data is not delay sensitive and it does not require real time transfer from a sender to a receiver. With packet switching, the data is broken into labeled packets and transmitted using packet-switching networks.

Virtual switching is incorrect because no such WAN switching technology exists.

Circuit switching is incorrect because circuit switching dynamically establishes a virtual connection between a source and destination. The virtual connection cannot be used by other callers unless the circuit is released. Circuit switching is the most common technique used by the Public Switched Telephone Network (PSTN) to make phone calls. A dedicated circuit is temporarily established for the duration of call between caller and receiver. Once the caller or receiver hangs up the phone, the circuit is released and is available for other users.

References:

http://docwiki.cisco.com/wiki/Introduction_to_WAN_Technologies#Circuit_Switching

QUESTION 58

Which of the following is NOT a characteristic of Open Shortest Path First (OSPF)?

- A. Is a Cisco-proprietary routing protocol
- B. Has a default administrative distance of 110
- C. Supports authentication
- D. Uses cost as the default metric

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

OSPF is not a Cisco-proprietary routing protocol. It is an industry standard protocol supported by a wide range of vendors. The following are characteristics of OSPF:

- Uses Internet Protocol (IP) protocol 89.
- Has a default administrative distance of 110.
- Is an industry standard protocol (non Cisco-proprietary).
- Supports Non-Broadcast Multi-Access (NBMA) networks such as frame relay, X.25, and Asynchronous Transfer Mode (ATM). The default hello interval for NBMA networks is 30 seconds.
- Supports point-to-point and point-to-multipoint connections.
- Supports authentication.
- Uses 224.0.0.6 as multicast address for ALLDRouters.
- Uses 224.0.0.5 as multicast address for ALLSPFRouters.

- Uses link-state updates and SPF calculation that provides fast convergence.
- Recommended for large networks due to good scalability. - Uses cost as the default metric.

References:

<http://www.ciscopress.com/articles/article.asp?p=98156&seqNum=4>

<https://www.cisco.com/c/en/us/obsolete/mixed-technologies/internetworking.html>

CCNA ICND2 Official Exam Certification Guide (Cisco Press, ISBN 1-58720-181-X), Chapter 9: OSPF, pp. 347-361.

QUESTION 59

Which of the following statements describes split horizon?

- A. The router learns from its neighbor that a route has gone down, and the router sends an update back to the neighbor with an infinite metric to that route.
- B. For a period of time, the router will ignore any route advertisements with a lower metric to a downed route.
- C. A router will not send route information back out the same interface over which it was learned.
- D. The moment a router determines a route has gone down, it will immediately send a route update with an infinite metric to that route.
- E. The packets are flooded when a topology change occurs, causing network routers to update their topological databases and recalculate routes.

Correct Answer: C

Section: (none)

Explanation



Explanation/Reference:

Explanation:

Split horizon is used to prevent routing loops in distance vector routing environments. It prevents a router from advertising a network back in the direction of the router from which it was learned. In this sense, route advertisements flow "downstream" (away from the route), but never "upstream" (back towards the advertised route).

Poison reverse describes when a router learns that a network has gone down, and the router sends an update back to the neighbor with an infinite metric.

Holddown describes when a router ignores any route advertisements that have a lower metric to a downed route.

Triggered updates describe when a router immediately sends a route update with an infinite metric, as opposed to waiting for its next regularly scheduled routing update.

Link State Advertisements (LSA) are packets that are flooded when a topology change occurs, causing network routers to update their topological databases and recalculate routes.

References:

<http://www.ciscopress.com/articles/article.asp?p=24090&seqNum=3>

QUESTION 60

You would like for Router25 in your OSPF network to become the DR. You execute the show ip ospf interface command, receiving the output shown below.

```
Router25# show ip ospf interface
Ethernet0 is up, line protocol is up
Internet Address 10.10.10.1/24, Area 0
Process ID 1, Router ID 10.10.10.1, Network Type BROADCAST, Cost: 10
Transmit Delay is 1 sec, State BDR, Priority 1
Designated Router (ID) 172.16.10.1, Interface address 10.10.10.2
Backup Designated router (ID) 10.10.10.1, Interface address 10.10.10.1
<output omitted>
```

You assign an IP address of 192.168.5.6 to the Ethernet1 interface of Router25 and enable the interface. However, Router25 does NOT become the designated router. What additional command must you execute to cause Router25 to become the DR?

- A. Router25(config-router)# network 192.168.5.0 0.0.0.255 area 0
- B. Router25(config-router)# network 192.168.5.0 0.0.0.255 area 1
- C. Router25(config-router)# network 192.168.5.0 255.255.255.0 area 0
- D. Router25(config)# network 192.168.5.0 0.0.0.255 area 0

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The command Router25(config-router)# network 192.168.5.0 0.0.0.255 area 0 must be executed to enable Router25 to become the DR. For an interface to be considered in the DR election, it must be advertised in OSPF. Otherwise, it is not participating in OSPF routing and you may be faced with the situation illustrated by the output of the shown ip ospf interface command below:

```
Router25# show ip ospf interface
Ethernet0 is up, line protocol is up
Internet Address 10.10.10.1/24, Area 0
Process ID 1, Router ID 225.16.33.4, Network Type BROADCAST, Cost: 10
Transmit Delay is 1 sec, State BDR, Priority 1
Designated Router (ID) 172.16.10.1, Interface address 10.10.10.2
```

Backup Designated router(ID)225.16.33.4,Interface address 10.10.10.1
<output omitted>

The RID of Router25, 225.16.33.4, is higher than that of the current DR, which has an RID of 172.16.10.1. Despite that fact, Router 25 did not become the DR because the 225.0.0.0 network has not been advertised. This could be verified by executing the show ip protocols command as shown below:

Router25# show ip protocols

Routing Protocol is "ospf 1"

Outgoing update filter list for all interfaces is not set

Incoming update filter list for all interfaces is not set

Router ID 225.16.33.4

<output omitted>

Routing for Networks:

10.0.0.0 0.0.0.255 area 0

As only the 10.0.0.0 network is being advertised, the 225.16.33.4 IP address will not be a factor in the DR election.

The command Router25(config-router)# network 192.168.5.0 0.0.0.255 area 1 is incorrect because it references area 1 instead of area 0, which is the area in use in this scenario.

The command Router25(config-router)# network 192.168.5.0 255.255.255.0 area 0 is incorrect because it uses a regular mask instead of a wildcard mask. Network commands in OSPF must use a wildcard mask.

The command Router25(config)# network 192.168.5.0 0.0.0.255 area 0 is incorrect because it is executed at the global configuration, router25(config)#, prompt rather than the OSPF configuration prompt, router25(config-router)#.

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

QUESTION 61

You are discovering that there are differences between the configuration of EIGRP for IPv6 and EIGRP for IPv4. Which statement is true with regard to the difference?

- A. A router ID is required for both versions
- B. A router ID must be configured under the routing process for EIGRP for IPv4
- C. AS numbers are not required in EIGRP for IPv6
- D. AS numbers are not required in EIGRP for IPv4

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Both versions of EIGRP require a router ID. The difference is that with EIGRP for IPv6, you must configure a router ID under the routing process if there are no IPv4 addresses on the router. In EIGRP for IPv4, the router can select one of the configured IPv4 addresses as the router ID.

A router ID can be configured under the routing process for EIGRP for IPv4, but it is not required. In EIGRP for IPv4, the router can select one of the configured Pv4 addresses as the router ID.

AS numbers are required in both versions of EIGRP.

Objective:

Routing Technologies Sub-

Objective:


Configure, verify, and troubleshoot EIGRP for IPv6 (excluding authentication, filtering, manual summarization, redistribution, stub)

References:

<http://www.ciscopress.com/articles/article.asp?p=2137516&seqNum=4>

QUESTION 62

With respect to the network shown below, which of the following statements are true when R2 sends a packet to the 192.168.6.0/24 network? (Choose all that apply.)



- A. If RIPv1 is in use, the path taken will be R2 - R4 - R3
- B. If both RIPv2 and EIGRP are in use, the EIGRP route will be placed in the routing table
- C. If EIGRP is in use, the only path taken will be R2 - R4 - R3
- D. If RIPv2 is in use, the path taken will be R2 - R3

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

If both RIPv2 and EIGRP are in use, the EIGRP route will be placed in the routing table. If RIPv2 is in use, the path taken will be R2 - R3.

EIGRP has a default administrative distance (AD) of 90, while RIPv2 has a default administrative distance (AD) of 120. The route learned by the routing protocol with the lowest AD will be placed in the routing table.

If you wanted to force R2 to use the RIPv2 route instead of the EIGRP route, this could be accomplished by changing the administrative distance of RIPv2 to a value less than 90, such as 80. The commands that would accomplish this are:

```

R2(config)# router rip
R2(config-router)# distance 80
  
```


If either of the versions of RIP is in use, hop count is used to determine the route. The path with the least number of hops is R2 - R3.

If RIPv1 is in use, the path taken would be R2 - R3, not R2 - R4 - R3, because R2 - R3 has a lower hop count.

If EIGRP is in use, the path R2 - R4 - R3 will not be the only path taken. EIGRP load-balances two equal cost paths when they exist, and R2 - R4 - R3 and R2 - R1 R3 are of equal cost so would both be used.

References:

<http://www.ciscopress.com/articles/article.asp?p=102174&seqNum=7>

QUESTION 63

What command would be used to verify trusted DHCP ports?

- A. show mls qos
- B. show ip dhcp snooping
- C. show ip trust
- D. show ip arp trust

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The command show ip dhcp snooping is used to verify trusted DHCP ports. This command is used to verify which ports are intended to have DHCP servers connected to them.

DHCP snooping creates an IP address to MAC address database that is used by Dynamic ARP Inspection (DAI) to validate ARP packets. It compares the MAC address and IP address in ARP packets, and only permits the traffic if the addresses match. This eliminates attackers that are spoofing MAC addresses.

DHCP snooping is used to define ports as trusted for DHCP server connections. The purpose of DHCP snooping is to mitigate DHCP spoofing attacks. DHCP snooping can be used to determine what ports are able to send DHCP server packets, such as DHCPOFFER, DHCPACK, and DHCPNAK. DHCP snooping can also cache the MAC address to IP address mapping for clients receiving DHCP addresses from a valid DHCP server.

MLS QOS has no bearing on DHCP services, so show mls qos is not correct.

The other commands are incorrect because they have invalid syntax.

References: <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr/command/ipaddr-cr-book/ipaddr-r1.html>

QUESTION 64

What command should you use to quickly view the HSRP state of the switch for all HSRP groups of which the switch is a member?

- A. switch# show standby brief
- B. switch# show ip interface brief
- C. switch# show hsrp
- D. switch# show standby

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The command show standby brief should be used to quickly view the HSRP state of a switch for all HSRP groups of which it is a member. The summary information it provides includes the group number, priority, state, active device address, standby address, and group address.

The command show standby can be used to display detailed information about HSRP groups of which a switch is a member. This command would not provide a quick view. This command displays information about HSRP on all configured interfaces and for all HSRP groups. It also displays hello timer information and the expiration timer for the standby switch.

The command show ip interface brief is useful in that lists the interfaces and displays the basic IP configuration of each. This output would include the IP address of the interface and the state of the interface, but not HSRP information.

The command show hsrp is not a valid command due to incorrect syntax.

References: <https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ipaddr/command/ipaddr-cr-book/ipaddr-r1.html> <https://www.cisco.com/c/en/us/products/index.html>

QUESTION 65

Which Cisco Internetwork Operating System (IOS) command is used to view the number of Enhanced Interior Gateway Routing Protocol (EIGRP) packets that are sent and received?

- A. show eigrp neighbors
- B. show ip eigrp interfaces
- C. show ip eigrp packets
- D. show ip eigrp traffic
- E. show ip route
- F. show ip eigrp topology

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The show ip eigrp traffic command is used to view the number of EIGRP packets that are sent and received. The syntax of the command is:

```
Router# show ip eigrp traffic [autonomous-system-number]
```

The autonomous-system-number parameter is optional. The output of the command is as follows:

```
Router# show ip eigrp traffic
```

```
IP-EIGRP Traffic Statistics for process 78
```

```
Hello sent/received: 2180/2005
```

```
Updates sent/received: 70/21
```

```
Queries sent/received: 3/1
```

```
Replies sent/received: 0/3
```

```
Acks sent/received: 22/11
```

The show ip eigrp neighbors command is incorrect because it does not show the number of packets sent or received. It does show IP addresses of the devices with which the router has established an adjacency, as well as the retransmit interval and the queue count for each neighbor, as shown below:

```
Router# show ip eigrp neighbors
```

```
IP-EIGRP Neighbors for process 49
```

```
Address Interface Holdtime Uptime Q Seq SRTT RTO
```

```
(secs) (h:m:s) Count Num (ms) (ms)
```

```
146.89.81.28 Ethernet1 13 0:00:41 0 11 4 20
```

```
146.89.80.28 Ethernet0 12 0:02:01 0 10 12 24
```

```
146.89.80.31 Ethernet0 11 0:02:02 0 4 5 20
```

The show ip eigrp interfaces command is incorrect because this command is used to view information about the interfaces configured for EIGRP.

The show ip eigrp packets command is incorrect because it is not a valid Cisco IOS commands.

The show ip route command will not display EIGRP packets that are sent and received. It is used to view the routing table. When connectivity problems occur between subnets, this is the logical first command to execute. Routers must have routes to successfully send packets to remote subnets. Using this command is especially relevant when the underlying physical connection to the remote network has been verified as functional, but routing is still not occurring.

The show ip eigrp topology command is incorrect because it does not show the number of packets sent or received. This command displays all successor and feasible successor routes (if they exist) to each network. If you are interested in that information for only a specific destination network, you can specify that as shown in the output below. When you do, the command output displays all possible routes, including those that are not feasible successors:

```
Router# show ip eigrp topology 25.0.0.5 255.255.255.255
```

```
IP-EIGRP topology entry for 25.0.0.5/32 State is Passive, Query origin flag is 1, 1 Successor(s), FD is 41152000
```

```
<output omitted>
```

```
10.1.0.1 (serial0), from 10.1.0.1 composite  
metric is 46152000/41640000
```

```
<output omitted>
```

```
10.0.0.2 (serial0.1), from 10.0.0.2 composite  
metric is 53973240/120256
```

```
<output omitted>
```

```
10.1.0.3 (serial0), from 10.1.0.3 composite  
metric is 46866176/46354176
```

```
<output omitted>
```

```
10.1.1.1 (serial0.1), from 10.1.1.1 composite  
metric is 46670776/46251776
```

```
<output omitted>
```



In the above output, four routers are providing a route to the network specified in the command. However, only one of the submitted routes satisfies the feasibility test. This test dictates that to be a feasible successor, the advertised distance of the route must be less than the feasible distance of the current successor route. The current successor route has a FD of 41152000, as shown in the first section of the output. In the values listed for each of the four submitted routes, the first number is the feasible distance and the second is the advertised distance. Only the route received from 10.0.0.2 (second section) with FD/AD values of 53973240/120256 satisfies this requirement, and thus this route is the only feasible successor route present in the topology table for the network specified in the command.

References:

https://www.cisco.com/c/en/us/td/docs/ios/12_2/iproute/command/reference/fiprrp_r/1rfeigrp.html

QUESTION 66

What command can be used on a Cisco switch to display the virtual MAC address for the HSRP groups of which the switch is a member?

- A. switch# show standby mac
- B. switch# show hsrp mac

- C. switch# show standby
- D. switch# show standby brief

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The command show standby can be used to display the virtual MAC address for HSRP groups of which a switch is a member. This command displays information about HSRP on all configured interfaces and for all HSRP groups. It also displays hello timer information and the expiration timer for the standby switch. The standby switch will take over as the active switch if the timer expires before it hears a heartbeat from the active switch. Below is an example of the show standby command for the HSRP group 1:

Tacoma# show standby

Fastethernet0/1 - group 1

State is active

3 state changes, last state change 00:22:49

Virtual IP address is 192.168.5.3

Secondary virtual ip address 192.168.5.3

Active virtual MAC address is 0006.6b45.5801

Local virtual MAC address is 0006.6b45.5812(bia)

Hello time is 4 sec, hold time 12 sec

Next hello sent in 1.664 sec

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

Active router is local

Standby router is unknown expired

Priority 95 (configured 120)

Tracking 2 objects, 0 up

Down Interface Fastethernet0/2, pri 15

Down Interface Fastethernet0/3

IP redundancy name is "HSRP1", advertisement interval is 34 sec



In the above output, the following can be determined:

- The router is currently active for the group, as can be seen in line 2. The Active Virtual MAC address is 0006.6b45.5801, which includes the group number (1) in the last two positions, which is why the address is different from the routers actual MAC address shown on the next line. Special Note: Some router models

(Cisco 2500, 4000 and 4500) WILL NOT use this altered MAC address format, but will instead use the real MAC address for the virtual MAC address and will display that MAC address as the virtual MAC address in the output of the show standby command. An example of the output of the show standby command on an older router such as the 2500 would be as follows:

```
Router# show standby
Ethernet0/1 - 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)
```

These routers have Ethernet hardware that only recognize a single MAC address. In either case if for some reason this router becomes the standby router, such as due to loss of interfaces, then when the interfaces come back up it will be able to recover the active role because it is set for preemption, as shown on line 10. - The router is tracking two of its own interfaces. Because both interfaces are down, the router's priority has been reduced by 25 (15 for Fastethernet0/2 and 10 for Fastethernet0/3), from the configured value of 120 to 95. This data is shown on lines 13-16. The default is 10 if not otherwise specified, as is the case for Fastethernet0/3.

- If either of the two interfaces comes back up, the priority will be increased by the amount assigned to the interface. For example, if Fastethernet0/3 comes back up, the priority will become 105 (95 + 10).
- The standby router is unreachable, which can be determined because it is marked unknown expired in line 12. This could be due to either a physical layer issue or an HSRP misconfiguration.

The command show standby brief can be used to view summary information about HSRP groups of which the switch is a member. This information includes the group number, priority, state, active device address, standby address, and group address. It does not include the virtual MAC address.

The commands show standby mac and show hsrp mac are invalid due to incorrect syntax.

References:

https://www.cisco.com/c/en/us/td/docs/ios/ipapp/command/reference/iap_s4.html

<https://www.cisco.com/c/en/us/products/index.html>

QUESTION 67

Which of the following statements is true with regard to SDN?

- A. It combines the control plane and the data plane

- B. It separates the data plane and the forwarding plan
- C. It implements the control plane as software

D. It implements the data plane as software

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

In Software-defined networking (SDN), the control plane is separated from the data (or forwarding) plane and is implemented through software. The data plane remains on each physical device but the control plane is managed centrally for all devices through software.

SDN does not combine the data and control plane. Instead it decouples them.

SDN does not separate the data plane and the forwarding plan. These are both names for the same plane; that is, a data plane is a forwarding plane.

SDN does not implement the data plane as software. The data plane remains on each physical device.

References: <http://www.techrepublic.com/article/software-defined-networking-the-cisco-approach/>

QUESTION 68

Which of the following is NOT managed by the cloud provider in an IaaS deployment?

- A. virtualization
- B. servers
- C. storage
- D. operating system

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Operating systems are not managed by the cloud provider in an Infrastructure as a service (IaaS) deployment. Only storage, virtualization, servers, and networking are the responsibility of the provider. The customer is responsible for the following with IaaS: - Operating systems

- Data
- Applications
- Middleware
- Runtime

In a Platform as a Service (PaaS) deployment, the provider is responsible for all except the following, which is the responsibility of the customer: - Applications - Data

In Software as a Service (SaaS) deployment, the provider is responsible for everything.

References: <https://apprenda.com/library/paas/iaas-paas-saas-explained-compared/>

QUESTION 69

Which redundancy mode for supervisor engine modules exhibits all of the following characteristics?

- Static routes are maintained during a switchover
- The Forwarding Information Base (FIB) is cleared during a switchover
- Dynamic route information is cleared during a switchover
- Route engine is initialized and switch modules are loaded

- A. RPR
- B. RPR+
- C. SSO
- D. NSF

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Redundant supervisor engine modules can be configured in several modes. In route processor redundancy plus (RPR+) mode, the backup module is booted up and the supervisor and route engines initialize. However, no Layer 2 or Layer 3 functions are started, which means it will be necessary to start them after a failover. This also means the routing protocols must re-converge and the FIB table must be rebuilt, since it is derived from the routing table. The static routes are maintained in the running configuration, so they are not lost in the failover.

In route processor redundancy (RPR) mode, the module is booted, but the supervisor and route engines are not initialized.

In stateful switchover (SSO) mode, all functionality provided by RPR+ is available at failover, and the FIB table is not cleared.

Non-stop forwarding (NSF) is not a redundant supervisor engine module mode but an additional redundancy feature designed to reduce the amount of time needed to rebuild the routing information base (RIB) table after a supervisor failure. Instead of waiting for any Layer 3 routing protocols to converge and rebuild the Forwarding Information Base (FIB), the router will use NSF to get assistance from other NSF-aware neighbors, allowing the routing information to be rebuilt quickly.

References: <https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4500/12-2/20ew/configuration/guide/config/RPR.html>

QUESTION 70

DRAG DROP

Click and drag the command line tools used to troubleshoot the network problems on the left to their associated functions on the right. Not all commands may be used.

Select and Place:

Commands:	Function:
ping 127.0.0.1	Displays the local IP address to MAC address mapping table on a Windows PC.
tracert	Verifies Layer 7 connectivity to a remote host.
telnet	Ensures that the TCP/IP protocol stack is running/active.
show ip arp	Used on a Cisco router to determine the routing path to a particular destination.
arp -a	
tracert	

Correct Answer:

Commands:

tracert
show ip arp

Function:

arp -a	Displays the local IP address to MAC address mapping table on a Windows PC.
telnet	Verifies Layer 7 connectivity to a remote host.
ping 127.0.0.1	Ensures that the TCP/IP protocol stack is running/active.
tracert	Used on a Cisco router to determine the routing path to a particular destination.

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The following commands can be used to troubleshoot network connectivity problems:

- ping 127.0.0.1: This command will attempt to contact the local TCP/IP protocol stack. The 127.0.0.1 address is the reserved loopback IP address, which allows applications to communicate with the local system without using an actual IP address assigned to an interface, such as a workstation's Ethernet port. Thus, this command allows you to ping yourself, and if successful, only verifies that TCP/IP is running locally. It does not confirm that the system can communicate with any other host on the network.
- telnet: Telnet is a network application used to establish a remote terminal connection to a host, such as logging in remotely to a Cisco router or switch via TCP/IP. Since network applications reside on the OSI Application Layer (Layer 7), a successful Telnet connection to a remote host confirms that there is network connectivity through Layer 7.
- arp ? a: This command is used to display the local IP address to MAC address mappings on a Windows PC.
- traceroute: This command is used on a Cisco router or switch to verify, or trace, the path that IP packets will take towards a particular destination.
- tracert: This command is used on a Windows PC to verify, or trace, the path that IP packets will take towards a particular destination.
- show ip arp: This command is used to display the local IP address to MAC address mappings on a Cisco router or switch.

References:

<https://www.cisco.com/en/US/docs/internetworking/troubleshooting/guide/tr1902.html>
<http://www.ciscopress.com/articles/article.asp?p=98156&seqNum=4>

QUESTION 71

You have two routers in your OSPF area 0. Router 1 is connected to Router 2 via its Serial 1 interface, and to your ISP via the Serial 0 interface. Router 1 is an ASBR.

After your assistant configures a default route on Router 1, you discover that whenever either router receives packets destined for networks that are not in the routing tables, it causes traffic loops between the two routers.

To troubleshoot, you execute the show run command on Router 1. Part of the output is shown below:

<output omitted>

IP route 0.0.0.0 0.0.0.0 serial 1

Router ospf 1

Network 192.168.5.0 0.0.0.255 area 0

Default-information originate

Which command or set of commands should you execute on Router 1 to stop the looping traffic while maintaining Router 2's ability to send traffic to the Internet?

- A. Execute the no default-information originate command.
- B. Execute the no ip route 0.0.0.0 0.0.0.0 serial 1 command and then execute the ip route 0.0.0.0 0.0.0.0 serial 0 command.
- C. Execute the default-information originate always command.
- D. Execute the no network 192.168.5.0 area 0 command and then execute the network 192.168.5.0 255.255.255.0 area 0 command.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should execute the no ip route 0.0.0.0 0.0.0.0 serial 1 command followed by the ip route 0.0.0.0 0.0.0.0 serial 0 command. The original configuration command was executed on the wrong interface on Router 1 by your assistant. It should be executed on Serial 0, which is the connection to the ISP. The show run command indicates that with the current configuration, if Router 2 receives a packet not in its table, it sends it to Router 1, and then Router 1 sends it back out on Serial 1. This redirects the packet back to Router 2, and the loop begins. By changing the configuration to Serial 0, Router 1 will start forwarding all traffic not in the routing table to the ISP.

You should not execute the `no default-information originate` command. This command instructs Router 1 to NOT inject the default route into area 0, which is the desired behavior. Running this command would stop the loop, but would leave Router2 with no default route to send packets to the Internet.

You should not execute the `default-information originate always` command. It will not change the existing looping behavior. The addition of the `always` parameter instructs Router 1 to inject a default route into area 0, even if one does not exist on Router 1. This is unnecessary, since Router 1 does have a default route configured, and will not change the existing looping behavior. To advertise a default route to other OSPF routers, you should run this command:

```
Router1(config-router)#default information originate
```

You should not execute the `no network 192.168.5.0 area 0` command followed by the `network 192.168.5.0 255.255.255.0 area 0` command. There is nothing wrong with the original network command. Also, the `network 192.168.5.0 255.255.255.0 area 0` command uses an incorrect mask type. The mask must be in the wildcard format. Moreover, since it is incorrect, this will have the effect of disabling OSPF on the network connecting the two routers.

References: <https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/47868-ospfdb9.html>

QUESTION 72

Which Cisco IOS command would you use to troubleshoot IP addressing problems?

- A. `ipconfig /all`
- B. `show config`
- C. `show running-config`
- D. `show config-file`



Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The `show running-config` command will help troubleshoot IP addressing problems, because it shows the details of the router configuration, including the IP address configured on each interface.

The `ipconfig /all` command is a Microsoft command used to verify IP address configuration on a workstation running Windows. This is not a valid Cisco command.

The `show config` command has been replaced by the `show startup-config` command. Both of these commands are used to display the startup configuration of the router stored in NVRAM.

The `show config-file` command is not a valid Cisco command.

References:

https://www.cisco.com/c/en/us/td/docs/ios/fundamentals/command/reference/cf_book/cf_s2.html

QUESTION 73

Which metric does the Open Shortest Path First (OSPF) routing protocol use for optimal path calculation?

- A. MTU
- B. Cost
- C. Delay
- D. Hop count

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

OSPF is a link-state routing protocol which uses cost as a metric for optimal path calculation. It is an open standard protocol based on Dijkstra's Shortest Path First (SPF) algorithm. Metrics are used by routing protocols to determine the lowest cost path to a network number, which is considered the optimal or "fastest" path. Cisco's implementation of OSPF calculates the cost (metric) of a link as inversely proportional to the bandwidth of that interface. Therefore, a higher bandwidth indicates a lower cost, and a more favorable metric.

For this to work properly, the bandwidth of the link must be configured to allow OSPF to arrive at the cost of the link. This is done with the bandwidth command executed in interface configuration mode, and is entered in kbps. For example, if the link were 64 kbps, you would enter the following command:

```
Router(config-if)# bandwidth 64
```

The metric for any OSPF link defaults to $100,000,000/\text{bandwidth}$. The bandwidth used in the formula is in bits per second. So, in this example the calculation would be $100,000,000 / 64000 = 1562.5$. The cost assigned to the link would be 1562. The cost for a network route is the sum of all individual links in the path to that network.

If multiple paths are assigned equal costs, OSPF will load balance across the multiple paths. By default it will limit this load balance to a maximum of four equal-cost paths. When this occurs, all four equal-cost paths will be placed in the routing table. There are two approaches to allow or prevent load balancing when multiple equal cost paths are available:

- Use the bandwidth command to make one or more of the paths either less or more desirable.
- Use the ip ospf cost command to change the cost value assigned to one or more of the paths

Maximum Transmission Unit (MTU), bandwidth, delay, load, and reliability form a composite metric used by Interior Gateway Routing Protocol (IGRP) and Enhanced Interior Gateway Routing Protocol (EIGRP). IGRP is a distance vector routing protocol developed by Cisco Systems. Enhanced IGRP (EIGRP) is a Cisco-proprietary hybrid protocol having features of both distance-vector and link-state protocols.

Hop count is a metric used by Routing Information Protocol (RIP). The fewer hops between the routers, the better the path.

References:

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

http://docwiki.cisco.com/wiki/Open_Shortest_Path_First

QUESTION 74

Which command was used to create the following configuration?

```
Router# show ip protocol
Routing Protocol is "eigrp 1"
Outgoing update filter list for all interfaces is not set
Incoming update filter list for all interfaces is not set
Redistributing: eigrp 1
Automatic network summarization is in effect
Routing for Networks:
 192.168.1.80/28
 192.168.1.128/28
Routing Information Sources:
 Gateway Distance Last Update
 192.168.1.85 90 0:04:01
Distance: internal 90 external 170
```

- A. Router(config-router)# network 192.168.1.0 0.0.0.15
- B. Router(config-router)# network 192.168.1.0 255.255.255.0
- C. Router(config-router)# network 192.168.1.80
- D. Router(config-router)# network 192.168.1.128
- E. Router(config-router)# network 192.168.1.0

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The network 192.168.1.0 command instructs the router to activate EIGRP on every interface that belongs to the class C network 192.168.1.0. The exhibit indicates that the router is running EIGRP on two subnets of 192.168.1.0 (192.168.1.80/28 and 192.168.1.128/28). Since both of these are subnets of the same class C network number, only the class C address needs to be referenced with a network statement.

All interfaces that will participate in EIGRP must be specified with a network command that specifying the network of which the interface is a member. Failure to do so will result in neighbor relationships not forming. In the example below, Router A and Router B are directly connected, but not forming a neighbor relationship. The network they share is the 192.168.5.0/24 network. The output of the show run command for both routers reveals that Router B does not have EIGRP running on the 192.168.5.0 network.

RouterA#show run	Router B#show run
<output omitted>	<output omitted>
router eigrp 36	router eigrp 36
network 192.168.5.0	network 10.0.0.0

The network 192.168.1.0 0.0.0.15 command is incorrect because only the class C network number (192.168.1.0) needs to be referenced to enable EIGRP on all subnets. It is actually valid to include an inverse mask with EIGRP network statements, but it is unnecessary in this case, and the network/mask provided does not match either of the routed networks.

The network 192.168.1.0 255.255.255.0 command is incorrect because the mask is unnecessary in this case, and if masks are included, they must be expressed inversely (0.0.0.255).

It is unnecessary to configure two network commands in this example, as both networks are subnets of the same class C network (192.168.1.0), and a single network command can enable EIGRP on both. Additionally, if specific subnets are referenced in network commands, it is necessary to include an inverse mask after them, or EIGRP will automatically summarize the command to the classful boundary.

References: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_eigrp/configuration/12-4t/ire-12-4t-book.pdf

QUESTION 75

Which VLAN can NOT be filtered through the VLAN Trunking Protocol (VTP) Pruning feature of Cisco switches?

- A. VLAN 1
- B. VLAN 10
- C. VLAN 100
- D. VLAN 1000

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

VLAN 1 traffic cannot be pruned. Cisco recommends that VLAN 1 be used for management of VLANs.

VTP pruning is a Cisco VTP feature that allows switches to dynamically delete or add VLANs to a trunk for traffic transmission. It creates an efficient switching network by optimal use of available trunk bandwidth.

The options 10, 100, and 1000 are incorrect because these VLAN numbers can be pruned. By default, VLANs 2 to 1000 are eligible for pruning.

References:

<http://www.ciscopress.com/articles/article.asp?p=102157&seqNum=6>

QUESTION 76

Which two statements are TRUE of synchronous serial ports? (Choose two.)

- A. These ports can be used to provide leased-line or dial-up communications.
- B. These ports do not support the High-Level Data Link Control (HDLC) encapsulation method.
- C. An AUI connector is used with serial ports.
- D. These ports can be used to configure high-speed lines (E1 or T1).
- E. An RJ-45 connector is used with serial ports.

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Synchronous serial ports can be used to provide leased-line or dial-up communications, and these ports can be used to configure high-speed lines (E1 or T1). The following are also true of synchronous serial ports:

- With the help of synchronous serial lines, dialers can be configured, which are then used to support dial-on-demand routing.
- These ports are found on several serial network interface processors and cards.

The option stating that synchronous serial ports cannot support High-Level Data Link Control (HDLC) encapsulation method is incorrect because HDLC is the default encapsulation method configured on serial interfaces.

The option stating that an AUI connector is used with serial ports is incorrect because AUI is a connector used with Ethernet ports.

The option stating that an RJ-45 connector is used with serial ports is incorrect because RJ-45 and RJ-48 connectors are used with ISDN BRI connections.

QUESTION 77

You have connected two routers in a lab using a Data Terminal Equipment (DTE)-to-Data Circuit-terminating Equipment (DCE) cable. Which command must be issued on the DCE end for the connection to function?

- A. bandwidth
- B. no clock rate
- C. clock rate
- D. no bandwidth

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should issue the clock rate command on the DCE end for the connection to function. The clock rate is set on the Data Circuit-terminating Equipment (DCE) device. DCE is also known as Data Communications Equipment.

The DCE terminates a physical WAN connection, provides clocking and synchronization of a connection between two locations, and connects to a DTE. The DCE category includes equipment such as CSU/DSUs, NT1s, and modems. In the real world, the clock rate is provided by the CSU/DSU end at the telcom provider. In a lab, you must instruct the DCE end to provide a clock rate.

The DTE is an end user device, such as a router or a PC, which connects to the WAN via the DCE device.

You would not issue the bandwidth command. This command is used to inform the router of the bandwidth of the connection for purposes of calculating best routes to locations where multiple routes exist. It is not necessary for the link described to function.

You should not issue the no clock rate command. This command is used to remove any previous settings implemented with the clock rate command.

You would not issue the no bandwidth command. This command is used to remove any previous settings implemented with the bandwidth command.

References:

CCNA ICND2 Official Exam Certification Guide (Cisco Press, ISBN 1-58720-181-X), Chapter 12: Point-to-Point WANs, pp. 446-447.

QUESTION 78

Which of the following TCP port numbers is used by Simple Mail Transfer Protocol (SMTP)?

- A. 23
- B. 21
- C. 53
- D. 80
- E. 57
- F. 25

Correct Answer: F

Section: (none)

Explanation

Explanation/Reference:

Explanation:

TCP port 25 is assigned to SMTP. SMTP is a Transmission Control Protocol (TCP)/ Internet Protocol (IP) protocol used to send and receive e-mail messages.

Important TCP port number assignments are as follows:

- TCP port 23 is used by Telnet to allow remote logins.
- TCP port 21 is assigned to File Transfer Protocol (FTP) for FTP control. FTP also uses port 20 to transmit FTP data.
- TCP and User Datagram Protocol (UDP) port 53 is assigned to Domain Name Service (DNS), which is used to convert hostnames into Internet Protocol (IP) addresses.
- TCP port 80 is used by Hypertext Transfer Protocol (HTTP), which is the base for transferring Web pages over the Internet.
- TCP port 57 is assigned to Mail Transfer Protocol (MTP).
- TCP port 22 is used by Secure Shell (SSH).
- UDP ports 67 and 68 are used by Dynamic Host Configuration Protocol (DHCP).
- UDP port 69 is used by Trivial FTP (TFTP).
- TCP port 110 is used by Post Office Protocol 3 (POP3).
- UDP port 161 is used by Simple Network Management Protocol (SNMP).
- TCP port 443 is used by Secure Sockets Layer (SSL).

TCP port numbers help to direct data to the appropriate application, service, or application window. TCP port numbers ensure that data is displayed in the correct browser window when accessing Web data from multiple sources, and ensures it is directed to the proper application or service when received.

References:

http://docwiki.cisco.com/wiki/Internetworking_Basics#Multiplexing_Basics

QUESTION 79

Which command produced the following output?

<output omitted>

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 0x29BEB

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

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:

Explanation:

The output was produced by the show ip ospf command. The show ip ospf command is used to view information about the OSPF routing processes. 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 show ip ospf database command is incorrect because this command is used to view the OSPF database for a specific router. The following is sample output from the show ip ospf database command when no arguments or keywords are used:

```
Router# show ip ospf database
OSPF Router with id(192.168.239.66) (Process ID 300)
Displaying Router Link States(Area 0.0.0.0)
Link ID ADV Router Age Seq# Checksum Link count
172.16.21.6 172.16.21.6 1724 0x80002CFB 0x69BC 5
172.16.21.5 172.16.21.5 2512 0x800009D2 0xA2B8 3
172.16.1.2 172.16.1.2 1659 0x80000A98 0x4CB6 7
172.16.1.1 172.16.1.1 5115 0x800009B6 0x5F2C 9
172.16.1.5 172.16.1.5 1626 0x80002BC 0x2A1A 4
172.16.65.6 172.16.65.6 1315 0x80001947 0xEE1 9
172.16.241.5 172.16.241.5 1123 0x8000007C 0x7C70 1
172.16.27.6 172.16.27.6 1712 0x80000548 0x8641 4
172.16.70.6 172.16.70.6 1142 0x80000B97 0xEB84 6
Displaying Net Link States(Area 0.0.0.0)
Link ID ADV Router Age Seq# Checksum
172.16.1.3 192.168.239.66 1245 0x800000EC 0x82E
Displaying Summary Net Link States(Area 0.0.0.0)
Link ID ADV Router Age Seq# Checksum
172.16.240.0 172.16.241.5 1152 0x80000077 0x7A05
172.16.241.0 172.16.241.5 1152 0x80000070 0xAEB7
172.16.244.0 172.16.241.5 1152 0x80000071 0x95CB
```



The show ip ospf statistics command is incorrect because this command is used to view the OSPF calculation statistics. The following is sample output from the show ip ospf statistics command that shows a single line of information for each SPF calculation:

```
Router# show ip ospf statistics
OSPF process ID 200
```

Area 0: SPF algorithm executed 10 times
Area 200: SPF algorithm executed 8 times
Summary OSPF SPF statistic
SPF calculation time
Delta T Intra D-Intra Summ D-Summ Ext D-Ext Total Reason
08:17:16 0 0 0 0 0 0 0 R,
08:16:47 0 0 0 0 0 0 0 R, N,
08:16:37 0 0 0 0 0 0 0 R, X
00:04:40 208 40 208 44 220 0 720 R, N, SN, X
00:03:15 0 112 4 108 8 96 328 R, N, SN, X
00:02:55 164 40 176 44 188 0 612 R, N, SN, X
00:01:49 0 4 4 0 4 4 16 R, N, SN, X
00:01:48 0 0 4 0 4 0 12 R, N, SN, SA, X
00:01:43 0 0 4 0 4 0 8 R,
00:00:53 164 40 176 44 188 0 612 R, N, SN, X

The show ip ospf traffic command is incorrect because this is not a valid command.

References:

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

QUESTION 80

Why is it recommended to use Spanning Tree Protocol (STP) in Local Area Networks (LANs) with redundant paths?

- A. To prevent loops
- B. To manage VLANs
- C. To load balance across different paths
- D. To prevent forwarding of unnecessary broadcast traffic on trunk links

Correct Answer: A

Section: (none)


Explanation

Explanation/Reference:

Explanation:

Spanning Tree Protocol (STP) is a Layer 2 protocol used in LANs to maintain a loop-free network topology by recognizing physical redundancy in the network and logically blocking one or more redundant ports.

An example of switch redundancy is shown in the diagram below. The connection from SW4 to SW2, while providing beneficial redundancy, introduces the possibility of a switching loop.



STP probes the network at regular intervals to identify the failure or addition of a link, switch, or bridge. In the case of any topology changes, STP reconfigures switch ports to prevent loops. The end result is one active Layer 2 path through the switch network.

STP is not used for management of Virtual Local Area Networks (VLANs). VLAN Trunking Protocol (VTP) simplifies the management of VLANs by propagating configuration information throughout the switching fabric whenever changes are made. In the absence of VTP, switch VLAN information would have to be configured manually.

STP is not used to load-balance traffic across different redundant paths available in a topology. Load balancing allows a router to use multiple paths to a destination network. Routing protocols, Routing Information Protocol (RIP), RIPv2, Interior Gateway Routing Protocol (IGRP), Enhanced IGRP (EIGRP), and Open Shortest Path First (OSPF) support load balancing. Similarly, multiple links can be combined in a faster single link in switches. This can be achieved with the Fast EtherChannel or Gigabit EtherChannel features of Cisco switches.

STP does not prevent forwarding of unnecessary broadcast traffic on trunk links. This is achieved by manually configuring VLANs allowed on the trunk, or through VTP pruning.

References: <https://www.cisco.com/c/en/us/td/docs/routers/7600/ios/12-2SXF/configuration/guide/swcg/spantree.html>

QUESTION 81

Which of the following items are NOT required to match for two routers to form an OSPF adjacency?

- A. Area IDs
- B. Hello/Dead timers
- C. Passwords (if OSPF authentication has been configured)
- D. Process IDs

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

All of the listed items must match except for the process IDs. The process IDs are locally significant, which keeps multiple instances of OSPF separate on a router, and do not need to match between neighboring routers for the adjacency to form. Process identifiers can be valued from 1 to 65535.

Adjacencies must be formed before routing updates can be exchanged. OSPF routers will form neighbor adjacencies on common subnets if the following three items match:

- Area IDs
- Hello/Dead timers
- Passwords (if OSPF authentication has been configured)

Once an adjacency has been formed it will be maintained by the exchange of Hello messages. On a broadcast medium like Ethernet, they will be sent every 10 seconds. On point-to-point links, they will be sent every 30 seconds.

The show ip ospf interface interface number command can be used to display the state of the DR/BDR election process.

Consider the following output:

```
RouterA# show ip ospf interface fastethernet0/0
```

```
Fastethernet0/0 is up, line protocol is up
Internet Address 192.168.30.2/24, Area 0
Process ID 1, Router ID 192.168.45.1, Network Type BROADCAST, Cost: 10
Transmit Delay is 1 sec, State DR, Priority 1
Designated Router (ID) 192.168.45.1, Interface address 192.168.30.2
No backup designated router on this network
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:06
```

```
RouterB# show ip ospf interface fastethernet0/0
```

```
Fastethernet0/0 is up, line protocol is up
```

Internet Address 192.168.30.1/24, Area 0
Process ID 2, Router ID 192.168.60.1, Network Type BROADCAST, Cost: 10
Transmit Delay is 1 sec, State DR, Priority 2
Designated Router (ID) 192.168.60.1, Interface address 192.168.30.1
No backup designated router on this network
Timer intervals configured, Hello 30, Dead 60, Wait 40, Retransmit 5
Hello due in 00:00:12

The timer intervals' configured output reveals that RouterA is showing a Hello timer of 10 seconds and a Dead timer of 40 seconds. RouterB has a Hello timer of 30 seconds and a Dead timer of 60 seconds. Hello/Dead timers have to match before OSPF routers will form an adjacency. If you executed the debug ip ospf events command on one of the routers, the router at serial /01 will not form a neighbor relationship because of mismatched hello parameters:

```
RouterA# debug ip ospf events
OSPF events debugging is on
RouterA#
*Nov 9 05:41:21.456:OSPF:Rcv hello from 10.16.2.3 area 0 from Serial0/1
192.168.35.1
*Nov 9 05:41:21.698:OSPF:Mismatched hello parameters from
192.168.35.1
```

Hellos are used to establish neighbor adjacencies with other routers. On a point-to-point network, hello packets are sent to the multicast address 224.0.0.5, which is also known as the ALLSPFRouters address.

Area IDs have to match for OSPF routers to form an adjacency. Both of these routers have the interface correctly configured in matching Area 0.

The interface priorities do not have to match for OSPF routers to form an adjacency. Interface priorities can be configured to control which OSPF router becomes the designated router (DR) or backup designated router (BDR) on a multi-access network segment.

References: <https://www.cisco.com/c/en/us/support/docs/ip/open-shortest-path-first-ospf/13699-29.html>

QUESTION 82

In the given exhibit, which combination shows the components of a bridge ID used for Spanning Tree Protocol (STP)?

1	<table><tr><td>VLAN Number</td><td>MAC Address</td></tr></table>	VLAN Number	MAC Address
VLAN Number	MAC Address		
2	<table><tr><td>Priority Number</td><td>Serial Number</td></tr></table>	Priority Number	Serial Number
Priority Number	Serial Number		
3	<table><tr><td>Priority Number</td><td>MAC Address</td></tr></table>	Priority Number	MAC Address
Priority Number	MAC Address		
4	<table><tr><td>VLAN Number</td><td>Serial Number</td></tr></table>	VLAN Number	Serial Number
VLAN Number	Serial Number		

- A. 1
- B. 2
- C. 3
- D. 4

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The bridge ID, also known as the switch ID, is used to elect the root bridge in a redundant network topology. The bridge ID has two components:

- Switch's priority number: Configured as 32768 on Cisco switches by default
- Switch's Media Access Control (MAC) address: The burnt-in hardware address of the network interface card (NIC)

The switch with the lowest bridge ID is elected as the root bridge. If the same priority number is configured on two or more switches in the network, the switch with the lowest MAC address will become the root.

Bridge Protocol Data Units (BPDUs) communicate the details of the switch with the lowest bridge ID in the network. The election process for the root bridge takes place every time there is a topology change in the network. A topology change may occur due to the failure of a root bridge or the addition of a new switch in the network. The root bridge originates BPDUs every two seconds, which are propagated by other switches throughout the network. BPDUs are used as keepalives between switches. If a switch stops receiving BPDUs from a neighboring switch for ten intervals (20 seconds), it will assume a designated role for the network segment.

The combinations of the remaining options are incorrect because Virtual LAN (VLAN) numbers and serial numbers are not components of a bridge ID.

References:

https://www.amazon.com/gp/product/1119288282/ref%3Das_li_tl?ie=UTF8&camp=1789&creative=9325&creativeASIN=1119288282&linkCode=as2&tag=transcender02-20&linkId=cd2bd2412c028f0db900fe3aef249938
<https://www.cisco.com/c/en/us/td/docs/routers/7600/ios/12-2SXF/configuration/guide/swcg/spantree.html>

QUESTION 83

Which of the following are characteristics of Open Shortest Path First (OSPF)? (Choose three.)

- A. Administrative distance of OSPF is 90
- B. Administrative distance of OSPF is 110
- C. OSPF uses the Dijkstra algorithm to calculate the SPF tree
- D. OSPF uses the Diffusing Update algorithm (DUAL) algorithm to calculate the SPF tree
- E. OSPF uses 224.0.0.5 as multicast address for ALLDRouters
- F. OSPF uses 224.0.0.6 as multicast address for ALLDRouters

Correct Answer: BCF

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The following are characteristics of Open Shortest Path First (OSPF) routing protocol:

- The default administrative distance is 110.
- It uses 224.0.0.6 as the multicast address for ALLDRouters.
- It uses the Dijkstra algorithm to calculate the Shortest Path First (SPF) tree. - It uses Internet Protocol (IP) protocol 89.
- OSPF supports Non-Broadcast Multi-Access (NBMA) networks such as Frame Relay, X.25, and Asynchronous Transfer Mode (ATM). The default hello interval for NBMA networks is 30 seconds.
- OSPF supports point-to-point and point-to-multipoint connections.
- It also supports authentication.
- OSPF uses 224.0.0.5 as the multicast address for ALLSPFRouters.
- It uses link-state updates and SPF calculations that provides fast convergence.
- OSPF is recommended for large networks due to good scalability. - It uses cost as the default metric.
- There is no maximum hop count as with distance vector routing protocols. The number of hops to a network can be unlimited.

The option stating that AD of OSPF is 90 is incorrect because 90 is the default administrative distance for an internal Enhanced Interior Gateway Routing Protocol (EIGRP) route.

The option stating that OSPF uses the Diffusing Update algorithm (DUAL) algorithm to calculate the SPF tree is incorrect. The DUAL algorithm is used by EIGRP to calculate the SPF tree.

Keep the following in mind when comparing OSPF and EIGRP:

- EIGRP is vendor specific; OSPF is not
- EIGRP has an AD of 90; OSPF has an AD of 110
- OSPF elects a DR on each multi-access network; EIGRP does not
- OSPF uses cost as its metric, and EIGRP uses bandwidth as its metric

The option stating that OSPF uses 224.0.0.5 as multicast address for ALLDRouters is incorrect because OSPF uses 224.0.0.6 as multicast address for ALLDRouters, and 224.0.0.5 as multicast address for ALLSPFRouters.


References:

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

https://www.cisco.com/en/US/tech/tk828/tech_brief09186a00800a4415.html

QUESTION 84

The partial output displayed in the exhibit is a result of what IOS command? (Click on the Exhibit(s) button.)

 **VCEplus**
VCE To PDF - Free Practice Exam

```
vlan 1 - Group 1
  State is Active
    2 state changes, last state change 00:30:59
  virtual IP address is 172.16.1.20
  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 172.16.1.6, priority 75 (expires in 9.184 sec)
  Priority 95 (configured 120)
  IP redundancy name is "Group1", advertisement interval is 34 sec
```

- A. switch# show running-config
- B. switch# show standby vlan1 active brief
- C. switch# show hsrp 1
- D. switch# show standby

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The command show standby produces the output displayed in the exhibit. This command displays information about HSRP on all configured interfaces and for all HSRP groups. Important information in the exhibit includes that this router is the active router, the virtual IP address for the HSRP group is 172.16.1.20, the address of the standby router is 172.16.1.6, and the router is configured to preempt.

The command show running-config will display the complete configuration of the device, including the configuration of HSRP, but will not display the current status of HSRP on the switch.

The command show standby vlan 1 active brief provides a summary display of all HSRP groups on the switch that are in the active state. This output would provide basic information, not nearly the detail indicated in the exhibit. The following is an example of output for show standby vlan 1 active brief:

Interface	Grp	Prio	P	State	Active	addr	Standby	addr	Group	addr
Vlan1	0	120		Active	172.16.1.5	Unknown	172.16.1.20			

The command show hsrp 1 is not valid due to incorrect syntax.

References:

https://www.cisco.com/c/en/us/td/docs/ios/ipapp/command/reference/iap_s2.html

<https://www.cisco.com/c/en/us/products/index.html>