

200-125.exam

Number: 200-125
Passing Score: 800
Time Limit: 120 min
File Version: 1.0

Cisco

200-125

Cisco Certified Network Associate (CCNA)

Version 1.0

Exam A

QUESTION 1

Which of the following is NOT an advantage of static routes over dynamic routing protocols?

- A. Routing protocol overhead is not generated by the router.
- B. Bandwidth is not consumed by route advertisements between network devices.
- C. Static routes are easier to configure and troubleshoot than dynamic routing protocols.
- D. Static route configuration is more fault tolerant than dynamic routing protocols.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Static route configuration is NOT more fault tolerant than dynamic routing protocols. The following lists the true advantages of static routes over dynamic routing protocols:

- Routing protocol overhead is not generated by the router.
- Bandwidth is not consumed by route advertisements between network devices.
- Static routes are easier to configure and troubleshoot than dynamic routing protocols.
- Router resources are more efficiently used.
- Network security is increased by using static routes.

The following are disadvantages of static routes:

- Static routes are not recommended for large networks because static routes are manually configured on the router. Therefore, maintaining routes in a timely manner is nearly impossible.
- Static route configuration is not fault tolerant without configuring multiple static routes to each network with varying administrative distances.

All other options are incorrect because these are the advantages of static routes over dynamic routing protocols.

Objective:

Routing Fundamentals

Sub-Objective:

Compare and contrast static routing and dynamic routing

References:

[Cisco Documentation > Internetworking Technology Handbook > Routing Basics](#)

QUESTION 2

Which settings represent the proper BITS per second, data bits, and parity settings for a HyperTerminal session to the router?

- A. 19200,8,none
- B. 9600,8,none
- C. 9600,8, even
- D. 19200,8,even

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

When setting up a HyperTerminal session to either a router or a switch, the proper settings must be enabled or

the connection will not work. The proper settings are 9600 bits per second, 8 data bits, and none in the parity box. The HyperTerminal application is provided on Windows operating systems up to Windows XP. The path to the tool is Start > Programs > Accessories > Communications > HyperTerminal. For later Windows operating systems, a HyperTerminal program must be downloaded and installed.

After opening the tool, you will name the connection and select a COM port if you want to use a port other than the default, which is the serial port to which the console cable is connected. You will then be presented with the following box:



The settings should reflect 9600 bits per second, 8 data bits, no parity, 1 stop bit, and no flow control, as illustrated in this graphic:



All other options include incorrect settings in at least one category.

Objective:

Infrastructure Management

Sub-Objective:

Configure and verify device management

References:

QUESTION 3

```
Protocol [ip]:
Target IP address: 10.10.10.1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 12.1.10.2
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.
```

Which Cisco IOS command would produce the preceding menu-based prompt for additional information?

- A. `tracert 10.10.10.1`
- B. `traceroute 12.1.10.2`

- C. ping 10.10.10.1
- D. ping

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

This menu-based prompt for additional information shown would be generated by the Cisco IOS ping command when issued without a target IP address. This is also known as issuing an extended ping. This command can be issued on the router to test connectivity between two remote routers. To execute an extended ping, enter the ping command from the privileged EXEC command line without specifying the target IP address. It takes the command into configuration mode, where various parameters, including the destination and target IP addresses, can be defined.

Note: You can only perform an extended ping at the privileged EXEC command line, while the normal ping works in both user EXEC mode and privileged EXEC mode.

The tracert command is incorrect because the tracert command is used by Microsoft Windows operating systems, not Cisco devices. This command cannot be run via the Cisco IOS command line interface. However, Microsoft's tracert utility is similar to Cisco's traceroute utility, which is to test the connectivity or "reachability" of a network device or host. The tracert command uses Internet Control Message Protocol (ICMP) to list all of the "hops" or routers traversed to a destination.

The traceroute command is incorrect because this command uses Internet Control Message Protocol (ICMP) to list all of the "hops" or routers traversed to a destination. It is also used to find routing loops or errors within a network.

The ping 10.10.10.1 command is incorrect because you when you issue this command you will either receive a reply from the destination or a destination unreachable message. It will not prompt for additional information as shown

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot IPv4 and IPv6 static routing

References:

[Cisco > Cisco IOS Command Fundamentals Reference, Release 12.4 > ping](#)

[Cisco > Tech Notes > Using the Extended ping and Extended traceroute Commands > Document ID: 13730 > The Extended ping Command](#)

[Cisco Documentation > Internetwork Troubleshooting Handbook > Troubleshooting TCP/IP](#)

QUESTION 4

On a Cisco 2950 switch, which status LED and color combination indicates a Power On Self-Test (POST) failure?

- A. system LED: no color
- B. system LED: solid red
- C. system LED: solid amber
- D. stat LED: no color
- E. stat LED: green

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A POST failure is indicated by a solid amber color on the system LED. The switch automatically runs POST which is a series of self-tests to verify proper functioning, after the power is connected. The system LED is off (no color) at the time that POST begins. The LED will turn green if POST is successful, or it will turn amber if POST fails.

The system LED will not be colorless. The system LED will show no color at the beginning of the POST cycle, not after a POST failure.

The system LED will not be solid red after a POST failure. Cisco LEDs do not have a red color mode.

The Stat LED indicates the status of each port. If it is amber there is a signal but the port is not forwarding, either because of an address violation or it has been disabled. If it is colorless, there is no signal. In this case:

- Ensure the switch has power
- Ensure the proper cable type is in use (for a switch to switch connection use a crossover cable; for a switch to host and or switch to router connection use a straight through)
- Ensure a good connection by reseating all cables

If it is green, the port has a signal and is functional. Green means:

- Layer 1 media is functioning between the switch and the device on the other end of the cable
- Layer 2 communication has been established between the switch and the device on the other end of the cable

LED color	Status
Off	RPS is either shut down or not installed.
Solid Green	RPS is installed and operational.
Blinking Green	Another switch in the stack is being backed up by RPS.
Solid Amber	Standby mode. It should turn green after pressing the active/standby button on the RPS. If it does not turn green, the RPS power supply or FAN might have failed.
Blinking Amber	Switch internal power supply is down and the switch is functioning on RPS.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot interswitch connectivity

References:

QUESTION 5

Which of the following loop avoidance mechanisms drives the requirement to create subinterfaces for each point-to-point connection in a partially meshed frame relay network?

- A. split horizon
- B. poison reverse
- C. maximum hop count
- D. feasible successor

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Split horizon is the loop avoidance mechanism that drives the requirement to create sub interfaces for each point-to-point connection in a partially meshed frame relay network. Frame relay is a non-broadcast multi-access (NBMA) network and obeys the rules of split horizon. This mechanism prohibits a routing protocol from sending updates out the same physical interface on which it was received. When the same physical interface is used to host multiple frame relay connections, this will prevent an update arriving from remote network A on the physical interface from being sent out the same interface to remote network B.

By creating a subinterface for each frame relay connection and assigning IP addresses to the subinterfaces rather than the physical interface, and by placing the subinterfaces into different subnets, split horizon will not see the "virtual" interfaces as the same interface and will allow these routing updates to be sent back out the same physical interface on which they arrived. It is important to map each subnet (or subinterface) to a remote Data Link Connection Identifier (DLCI) so that traffic to a remote network can be sent out the correct subinterface.

To summarize this discussion:

- Subinterfaces solve the NBMA split horizon issues.
- There should be one IP subnet mapped to each DLCI

Poison reverse is not the mechanism driving the requirement to create subinterfaces for each point-to-point connection in a partially meshed frame relay network. This mechanism requires a router to send an unreachable metric to the interface on which a network was discovered when it is learned from another interface that the network is no longer available.

Maximum hop count is not the mechanism driving the requirement to create sub interfaces for each point-to-point connection in a partially meshed frame relay network. Each routing protocol has a maximum hop count, which is the maximum number of hops allowed to a remote network before the network is considered "unreachable".

Feasible successor is not the mechanism driving the requirement to create sub interfaces for each point-to-point connection in a partially meshed frame relay network. This is a concept unique to EIGRP that represents a secondary route to a network that is considered the "best" route of possible backup routes.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure and verify Layer 2 protocols

References:

[Cisco>Home>Support>Technology Support>IP>IP Routing>Technology Information>technology Whitepaper>EIGRP> Split Horizon and Poison Reverse](#)

QUESTION 6

How is load balancing achieved when implementing HSRP?

- A. By configuring multiple gateways on the routers
- B. By using multiple HSRP groups
- C. By configuring the same priority on all HSRP group members
- D. By configuring multiple virtual router addresses

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

When implementing Hot Standby Router Protocol (HSRP), load balancing is achieved by using multiple HSRP groups. Routers configured for HSRP can belong to multiple groups and multiple VLANs. By configuring one group to be active for Router A and standby for Router B, and the second group to be active for Router B and standby for Router A, both routers A and B can be used to pass traffic, as opposed to one sitting idle.

Load balancing cannot be achieved by configuring multiple gateways on the routers. The routers have one IP address. Each group will have a virtual IP address. In the configuration below, line 4 configures the virtual IP address, and is therefore the address that clients will use as their gateway:

```
interface fastethernet 0/1
no switchport
ip address 192.168.5.5 255.255.255.0
standby 1 ip 192.168.5.10
```

Load balancing cannot be achieved by configuring the same priority on all HSRP group members. If that were done, one of the routers would become active and the others would remain inactive standbys. The active router will be the one with the highest IP address.

Load balancing cannot be achieved by configuring multiple virtual router addresses. Each HSRP group can only have one virtual address.

Objective:

Infrastructure Services

Sub-Objective:

Configure, verify, and troubleshoot basic HSRP

References:

[Internetworking Case Studies > Using HSRP for Fault-Tolerant IP Routing](#)

QUESTION 7

What command would you run to determine which switch is the root bridge for a particular VLAN?

- A. show spantree vlan
- B. show spanning tree
- C. show vlan spantree
- D. show spanning-tree vlan
- E. show spanning-tree interface

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The show spanning-tree vlan command provides Spanning Tree Protocol (STP) information on the root switch, including the bridge ID, root path, and root cost, as well as information on the local switch. The output of the command is as follows:

```
Switch# show spanning-tree vlan 1
VLAN0001
```

```
Spanning tree enabled protocol ieee
```

```
Root ID      Priority      0
Address      000c.00d3.5124
Cost         19
Port         2 (FastEthernet0/2)
Hello Time    2 sec Max Age 20 sec Forward Delay 15 sec
```

```
Bridge ID Priority 32769 (priority 32768 sys-id-ext 1)
Address 000c.14f5.b5c0
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
Aging Time 300
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Fa0/2	Root	FWD	19	128.2	P2p
Fa0/10	Desg	FWD	19	128.10	P2p
Fa0/1	Altn	BLK	19	128.1	P2p

This output indicates the following:

- The root switch has a bridge ID (Priority + MAC Address) of 0-000c.00d3.5124, while the local switch has a bridge ID of 32769-000c.14f5.b5c0. This indicates that the local switch is not the root switch for VLAN 1. Additional evidence that the local switch is not the root switch is the fact that the Fa0/1 port is blocking with a role listed as Altn. Only non- root bridges have blocking ports.
- For this switch, Fa0/1 represents the redundant link that needs to be blocked to prevent a switching loop.
- Interface Fa0/2 is the root port (the interface with the shortest path to the root switch).
- All three links have a cost of 19, which is the default cost of a single FastEthernet link.
- 802.1d is enabled in this switch, as indicated by the output Spanning tree enabled protocol ieee in line 2.

The show spanning-tree interface command will indicate the port role and state that a particular interface plays in each VLAN, but does not indicate the root bridge for a particular VLAN. Below is sample output from the show spanning-tree interface fastethernet0/1 command. In this example, RSTP is in use rather than 802.1d.

```
Switch# show spanning-tree interface fastethernet0/1
```

VLAN	Role	Sts	Cost	Prior.Nbr	Type
VLAN0001	Altn	BLK	19	128.2	P2P
VLAN0002	Root	FWD	19	128.1	P2P
VLAN0003	Root	FWD	19	128.1	P2P

In the above output, the Fa0/1 interface is not the root bridge for any of the three VLANs. It is the root port for VLANs 2 and 3. Root bridges have only designated ports. It is the alternate port for VLAN1, which means that Fa0/1 has a higher cost path to the root bridge than another interface in the topology, and will be in a blocking state as long as that other path is available.

The other options are incorrect because they are not valid Cisco IOS commands. The correct syntax would be show spanning-tree, not show spanning tree or show spantree.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot STP protocols

References:

[Cisco > Cisco IOS Switching Command Reference > show spanning-tree](#)

QUESTION 8

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.

Objective:

Routing Fundamentals

Sub-Objective:

Compare and contrast distance vector and link-state routing protocols

References:

[Cisco > Articles > Network Technology > General Networking > Dynamic Routing Protocols](#)

QUESTION 9

Which cable can suffer attenuation if it is bent beyond the minimum bend radius?

- A. UTP
- B. STP
- C. Co-axial
- D. Fiber optic

Correct Answer: D

Section: (none)**Explanation****Explanation/Reference:****Explanation:**

Fiber-optic cables can suffer attenuation if they are bent beyond the minimum bend radius. Fiber-optic cables work on the principle of total internal reflection. The fiber optic cable uses a laser and glass tubes with refractive internal coating to achieve total internal reflection. If a light ray travelling in the tube is bent at a certain angle, the light ray will be reflected inside the medium instead of passing through the medium. If the fiber optic cables are bent beyond the minimum bend radius, the signal will be lost and the cable will suffer attenuation. Fiber cables are expensive and are typically used for outdoor campus backbone. However, as the fiber cables use light to carry signals, they are not affected by the electro-magnetic interference (EMI) generated by electric cables.

Another advantage of fiber optic cabling is its applicability to situations where electrical issues may exist in the environment. Even in situations where the length of the cable run is well within the attenuation limits of STP (for example 55 meters), voltage differences between buildings can cause issues. That is a problem that can be solved by using fiber on the run, which is not impacted by electrical issue.

All other cables typically use copper to carry low voltage signals and are not affected by normal bending. However, even copper cables may suffer some signal loss if there are bends in the cable.

Objective:

Network Fundamentals

Sub-Objective:

Select the appropriate cabling type based on implementation requirements

References:

[Cisco > Articles > Network Technology > General Networking > Fiber-Optic Technologies](#)

QUESTION 10

Which type of network connection requires a straight-through cable?

- A. host to host
- B. switch to router
- C. switch to switch
- D. host to router's Ethernet port

Correct Answer: B**Section: (none)****Explanation****Explanation/Reference:****Explanation:**

A switch to router connection requires a straight-through cable. Straight-through cables are also used for host to switch communication.

A crossover cable is used to connect "like" devices, and a straight through cable is used when connecting "unlike" devices. The one exception to this rule is when connecting a computer NIC to an Ethernet port on a router, a crossover cable is used. In summary, the following list describes when to use crossover and straight through cables:


- Host to host Crossover
- Host NIC to router Crossover
- Host to switch Straight through
- Switch to Switch Crossover
- Switch to router Straight through

The difference between straight-through and crossover lies in the location of the wire termination on the two

ends of an RJ-45 cable. If the unshielded twisted-pair (UTP) cable wire connects Pin 1 of one side to Pin 1 of other side and Pin 2 to 2 through all eight Pins of the RJ-45 connector, the cable is said to be straight-through.

On the other hand, if the Pin 1 of one side RJ-45 cable connected to Pin 3 of other end and Pin 2 connects to Pin 6 of other side, it is called as crossover cable. The cable type to be used depends upon circuit connection on the hardware. Some devices have ports that are capable of identifying the cable type and automatically adjusting the port setting to be a standard or uplink port.

Host-to-host, switch-to-switch, and host-to-Ethernet-port would all use a crossover cable to connect in the network. The following figure shows the pin layout for a crossover cable:



Objective:

Network Fundamentals

Sub-Objective:

Select the appropriate cabling type based on implementation requirements

References:

[Cisco > Product Support > End-of-Sale and End-of-Life Products > Cisco 7000 Series Routers > Troubleshooting Technotes > Cabling Guide for Console and AUX Ports > Document ID: 12223](#)

QUESTION 11

Which of the following is a frame tagging method for identifying Virtual LAN (VLAN) memberships over trunk links?

- A. STP
- B. RIP
- C. CDP
- D. 802.1q

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

802.1q is a frame tagging method for identifying Virtual LAN (VLAN) memberships over trunk links. Frame tagging ensures identification of individual VLAN frames over a trunk link that carries frames for multiple VLANs. This frame tagging method is a standardized protocol developed by The Institute of Electrical and Electronics Engineers (IEEE). Cisco has also developed a proprietary frame tagging method, known as Inter-Switch Link (ISL).

When configuring a trunk link between a router and a switch, you must configure the physical interface on the router with one subinterface for each VLAN, and you must configure the physical ports on the router and the switch with the same encapsulation type, whether 802.1q or ISL.

Spanning Tree Protocol (STP) is not a frame tagging method, but a protocol used to remove switching loops in redundantly configured switched environments and create a single active Layer 2 path between any two network segments. Whenever a network segment can be handled by more than one switch, STP will elect one switch to take responsibility, and the other switches will be placed into a blocking state for the ports connected to that segment. In this way, only one switch receives and forwards data for this segment, removing the potential for generating multiple copies of the same frame. The benefits of STP include:

- Prevention of broadcast storms
- Prevention of multiple frame copies
- Media Access Control (MAC) address database stability

Routing Information Protocol (RIP) is not a frame tagging method, but a distance vector routing protocol. It populates routing tables dynamically about the topology changes.

Cisco Discovery Protocol is not a frame tagging method, but a Cisco proprietary protocol used to collect hardware and protocol information for directly connected Cisco devices. CDP has nothing to do with VLANs.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot VLANs (normal/extended range) spanning multiple switches

References:

[Cisco > Support > Technology Support > LAN Switching > Layer-Three-Switching and Forwarding > Configure > Configuration Examples and Technotes > Configuring InterVLAN Routing and ISL/802.1Q Trunking on a Catalyst 2900XL/3500XL/2950 Switch Using an External Router](#)

QUESTION 12

A packet is received with a destination IP address of 10.2.16.10.

```
Router# show ip route
<<output omitted>>
```

```
D 10.0.0.0 /8 [90/2172515] via 192.168.1.10, 00:00:44, Serial0/0
D 10.1.0.0 /16 [90/2144425] via 192.168.1.10, 00:01:03, Serial0/0
C 192.168.1.0 is directly connected, Serial0/0
C 192.168.4.0 is directly connected, Serial0/1
D 10.2.16.0 /24 [90/2162425] via 192.168.4.2, 00:00:25, Serial0/1
C 192.168.10.0 is directly connected, Serial1/0
D 10.2.32.0 /24 [90/2172425] via 192.168.10.254, 00:00:21, Serial1/0
    90/2172425] via 192.168.1.10, 00:03:33, Serial0/1
```

What would the next hop IP address be for this packet?

A. 192.168.1.10

- B. 192.168.4.2
- C. 192.168.10.254
- D. None; the packet will be dropped.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The packet will be routed to the next hop IP address of 192.168.4.2, since this routing table entry is the most specific match for the remote network. Packets are routed according to the most specific, or "longest," match in the routing table.

The packet in the scenario has a destination IP address of 10.2.16.10, which matches two entries in the routing table.

- 10.0.0.0 /8: this matches based on the /8 mask, where only the first byte has to match. The destination IP address of 10.2.16.10 has a first byte matching 10. If this were the only matching route table entry, it would be selected.
- 10.2.16.0 /24: The first 24 bits of this entry match the first 24 bits of the destination IP address of 10.2.16.10.

Therefore, the 10.2.16.0 /24 entry is selected for routing this packet because it most specifically matches the destination IP address, or has the longest number of matching bits.

The next hops of 192.168.1.10 and 192.168.10.254 will not be used, as these routes are not the most specific matches for the destination IP address of the packet.

It is interesting to note that packets that are destined for the 10.2.32.0 network will be load balanced across both serial 0/0 and serial 0/1 because the cost (2172425) is the same for both paths.


The packet will not be dropped because there is at least one routing table entry that matches the destination IP address of the packet.

To ensure that no packets are dropped, even if there is no matching route in the routing table, a default route could be configured as follows (next hop picked at random for illustration):

Router(config)# ip route 0.0.0.0 0.0.0.0 192.168.1.1

This configuration would instruct the router to send any packets that do match the existing routes to 192.168.1.1. For example, a packet destined for 201.50.6.8/24 would not match any routes in the table, and would thus be forwarded to 192.168.1.1.

If you understand how routing tables and routing advertisements work, it is relatively simple to describe the contents of a router's routing table without seeing the table directly. To do so, you would view the router's configuration and the configuration of its neighbors using show run, along with a diagram of its network connections. For example, examine the diagram of the two routers shown below along with their respective configurations:



```

hostname router 1      hostname router 2
router rip              router rip
network 192.168.54.64   network 192.168.110.128
ip route 0.0.0.0 0.0.0.0 192.35.87.5 <output omitted> <output omitted>

```

Based on this output and diagram, we can reconstruct the contents of the routing table for Router 1 as follows.

```

S*0.0.0.0/0 [1/0] via 192.35.87.5
R 192.168.110.128/26 [120/1] via 192.35.87.5 00:00:22, Serial 0/0
C 192.35.87.4/30 is directly connected, S0/0
C 192.168.54.64/26 is directly connected, Fa0/0

```

It will contain S*0.0.0.0/0 [1/0] via 192.35.87.5 because of the static default route indicated in line 4 of its configuration output.

It will contain R 192.168.110.128/26 [120/1] via 192.35.87.5 00:00:22, Serial 0/0 because Router 2 has a network 192.168.110.128 statement indicating that it will advertise this network to its neighbors.

It will contain the two routes C 192.35.87.4/30 is directly connected, S0/0 and C 192.168.54.64/26 is directly connected, Fa0/0 because all directly connected routes are automatically placed in the table.

Objective:

Routing Fundamentals

Sub-Objective:

Interpret the components of routing table

References:

[Cisco > Support > IP > IP Routing > Design > Design TechNotes > Route Selection in Cisco Routers > Document ID: 8651](#)

QUESTION 13

Which of the following statements are TRUE regarding carrier sense multiple access collision detection (CSMA/CD)? (Choose three.)

- A. Networks are segmented into multiple collision domains using switches for CSMA/CD networks.
- B. Networks are segmented into multiple broadcast domains using switches for CSMA/CD networks.
- C. CSMA/CD networks normally operate on half-duplex mode.
- D. CSMA/CD networks normally operate on full-duplex mode.
- E. Gigabit Ethernet uses CSMA/CD as the media access control method.
- F. Gigabit Ethernet uses carrier sense multiple access with collision avoidance (CSMA/CA) as the media access control method.

Correct Answer: ACE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The following statements are true:

- Networks are segmented into multiple collision domains using switches for CSMA/CD networks
- CSMA/CD networks normally operate on half-duplex mode
- Gigabit Ethernet uses CSMA/CD as its media access control method

CSMA/CD is a Local Area Network (LAN) access method used in Ethernet. In CSMA/CD, if a device or a node wants to send a packet in the network, it first determines if the network is free. If the network is not free, then the node will wait before sending the packet into a network. If the network is free, then the node sends the packet; if another device sends a packet simultaneously, their signals or packets collide. When the collision is detected, both packets wait for a random amount of time before retrying.

The option stating that networks are segmented into multiple broadcast domains using switches for CSMA/CD networks is incorrect because networks are segmented into multiple broadcast domains using routers for CSMA/CD networks.

The option stating that CSMA/CD networks normally operate on full-duplex mode is incorrect; these networks normally operate on half-duplex mode.

The option stating that gigabit Ethernet uses CSMA/CA as the media access control method is incorrect because gigabit Ethernet uses CSMA/CD as the media access control method.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Describe and verify switching concepts

References:

[Cisco > Internetworking Technology Handbook > Introduction to LAN Protocols > LAN Media-Access Methods](#)
[Cisco > The Internet Protocol Journal - Volume 2, No. 3 > Gigabit Ethernet](#)

QUESTION 14

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.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot VLANs (normal/extended range) spanning multiple switches

References:

[Cisco > Support > LAN Switching > Virtual LANS / VLAN Trunking Protocol \(VLANs/VTP\) > Configure > Configuration Examples and Technotes > Configuring VLAN Trunk Protocol \(VTP\) > Document ID: 98154](#)
[Cisco > Catalyst 4500 Series Switch Cisco IOS Software Configuration Guide, 12.2\(25\) > Understanding and Configuring VLANs, VTP, and VMPS](#)

CCNA ICND2 Official Exam Certification Guide (Cisco Press, ISBN 1-58720-181-X), Chapter 1: Virtual LAN Concepts, pp. 16-20.

CCNA ICND2 Official Exam Certification Guide (Cisco Press, ISBN 1-58720-181-X), Chapter 1: Virtual LAN Concepts, pp. 38-42.

QUESTION 15

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 0xEEE1 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 0xAE7
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.

Objective:

Routing Fundamentals

Sub-Objective:

Compare and contrast interior and exterior routing protocols

References:

[Cisco > Cisco IOS IP Routing Protocols Command Reference > OSPF Commands: show ip ospf through T > show ip ospf](#)

CCNA ICND2 Official Exam Certification Guide (Cisco Press, ISBN 1-58720-181-X), Chapter 9: OSPF, p. 375.

QUESTION 16

Which Cisco command keeps unauthorized users from viewing passwords in the router configuration file?

- A. enable secret
- B. enable password
- C. enable encryption
- D. service encryption
- E. service password-encryption

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The service password-encryption global configuration mode command keeps unauthorized users from viewing passwords in the router configuration file. The service password-encryption command encrypts all current and future passwords configured on the router, including the line password, virtual terminal password, console password, user name password, routing protocol passwords such as BGP neighbor passwords, the privileged command password, and authentication key passwords. Moreover, it encrypts any future passwords created on the router.

The encryption process occurs whenever the current configuration is built or a password is configured. The service password-encryption command will cause the router configuration file to display encrypted characters

instead of passwords when the running-configuration or startup-configuration files are viewed.

The enable password command creates a password that will be required to enter privileged EXEC mode, but the password will not be encrypted.

The enable secret command provides encryption to the enable mode passwords but does not apply globally to all passwords configured on the router. It also does not encrypt any future passwords created on the router.

The enable encryption and service encryption commands are invalid.

Objective:

Infrastructure Security

Sub-Objective:

Configure, verify, and troubleshoot basic device hardening

References:

[Cisco > Cisco IOS Security Command Reference > service password-encryption](#)

[Cisco Tech Notes > Cisco IOS Password Encryption Facts > Document ID: 107614](#)

QUESTION 17

Which of the following NAT terms refers to a registered address that represents an inside host to an outside network?

- A. inside global
- B. outside global
- C. inside local
- D. outside local

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:


Explanation:

An inside global address is the registered or public address assigned by the NAT server to an inside host. To the outside world it is only address seen for the host. The real (private) address, called the inside local, remains hidden.

The outside global address is the public address of a host that is outside the network connecting to the network. If the NAT server is also translating incoming addresses as well as outgoing addresses the address it assigns to the external host is called the outside local address.

The inside local address is the private IP address assigned to a host inside the network. The outside local is the address of a host outside the network as seen by hosts inside the network.

For example, in the diagram below, PC1 is sending a packet to the Web server. R1 is operating in NAT overload mode, which means that it maps all internal private IP addresses to a single public IP address.



Below is a listing of the names assigned to certain addresses in the diagram:

Inside local 192.168.5.6
Inside global 58.59.6.3
Outside local 192.168.5.1
Outside global 25.6.5.7

Objective:

Infrastructure Services

Sub-Objective:

Configure, verify, and troubleshoot inside source NAT

References:

[Cisco > Home > Support > Technology Support > IP > IP Addressing Services > Design > Design Technotes > NAT: Local and Global Definitions](#)

QUESTION 18

Which commands would you use to determine the IP address and hostname of a directly connected switch from which you received VLAN information? (Choose two. Each correct answer is part of the solution.)

- A. show vtp status
- B. show cdp neighbors detail
- C. show cdp neighbor status
- D. show vtp counters
- E. show cdp neighbor

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The VLAN Trunking Protocol (VTP) is used to synchronize VLANs between switches, and the question implies that VTP is being used in this environment. The show vtp status command will display the IP address of the switch that last updated your VLAN database. The output of this command is as follows:

```
Switch# show vtp status
VTP Version : 2
Configuration Revision : 247
Maximum VLANs supported locally : 1005
Number of existing VLANs : 33
VTP Operating Mode : Server
VTP Domain Name : Lab_Network
VTP Pruning Mode : Enabled
VTP V2 Mode : Disabled
VTP Traps Generation : Disabled
MD5 digest : 0x45 0x52 0xB6 0xFD 0x63 0xC8 0x49 0x80
Configuration last modified by 10.1.1.2 at 8-12-99 15:04:49
<output omitted>
```

The "Configuration last modified by 10.1.1.2" output reveals the IP address of the switch from which you received VLAN information. Once you know the IP address of the switch, you can use the show cdp neighbors detail command to determine the hostname associated with this IP address. The output of this command is as follows:

```
switch# show cdp neighbors detail
Device ID: RouterB
Entry address(es):
IP address: 172.20.52.254
Platform: cisco 2621, Capabilities: Router
Interface: FastEthernet0/1, Port ID (outgoing port):
FastEthernet0/0
Holdtime: 120 sec
<<output omitted>>
```

```
-----

Device ID: SwitchB
Entry address(es):
IP address: 10.1.1.2
Platform: cisco WS-C2950G-24, Capabilities: Switch IGMP
Interface: FastEthernet0/4, Port ID (outgoing port):
FastEthernet0/24
Holdtime: 101 sec
<<output omitted>>
```

The show cdp neighbors detail command provides detailed information about directly connected Cisco devices. The detail option is required to provide the IP address of the neighboring devices, and indicates here that IP address 10.1.1.2 is assigned to Device ID: SwitchB, which is the hostname for this device. SwitchB is the switch from which you received VLANs.

Although not offered as an option, the show cdp entry* command will also display all directly connected devices and will indicate the hostname and the IP address and platform, but will not indicate from which device VTP information was received. Its output is shown below:

```
switch#show cdp entry*
```

```
-----
```

```
Device ID: SwitchB  
Entry address(es):  
IP address: 10.1.1.2  
Platform: cisco WS-C2950G-24, Capabilities: Switch IGMP  
Interface: FastEthernet0/4, Port ID (outgoing port):  
FastEthernet0/24  
Holdtime: 101 sec  
<<output omitted>>
```

This command displays the same information as the show cdp neighbor detail command. It includes:

- The IP address of the neighbor (in this case 10.1.1.2)
- The port on which the CDP information was received (in this case FastEthernet0/4)
- The platform (in this case a Cisco WS-C2950G-24 Switch)

The show vtp counters command is incorrect because it does not display information about neighboring devices, nor information regarding from which switch VLANs were received.

The show cdp neighbor command is incorrect because the detail option is required to display the IP addresses of neighboring devices.

The show cdp neighbor status command is incorrect because this is not a valid Cisco IOS command.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot VLANs (normal/extended range) spanning multiple switches

References:

QUESTION 19

When the copy running-config startup-config command is issued on a router, where is the configuration saved?

- A. Random access memory (RAM)
- B. Flash
- C. Non-volatile random access memory (NVRAM)
- D. Read-only memory (ROM)

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

When the copy running-config startup-config command is issued on a router, the configuration is saved in the non-volatile random access memory (NVRAM) memory. The copy startup-config running-config command copies the version in RAM to NVRAM.

Note: For the copy startup-config running-config command to function, there must be a configuration already residing in RAM. For example, a brand-new router with no configuration created would have no startup configuration in RAM. If you attempted to execute the copy startup-config running-config command in that case, you would receive the following error message

% non-volatile memory configuration is invalid or not present

In addition to storing the running configuration in the NVRAM, you can also store it on a Trivial File Transfer Protocol (TFTP) server. When a router boots in the absence of a startup configuration, the router will look for a valid configuration on a TFTP server. In the case that the TFTP server also does not have a valid router configuration or is unreachable, the router will enter the setup dialog and prompt the user to provide initial configuration inputs.

The router does not store the startup configuration in random access memory (RAM). RAM only holds the running configuration that is loaded from the NVRAM or TFTP server during the boot process.

The router does not store the configuration in flash or read-only memory (ROM). ROM contains the bootstrap code, while flash memory contains the IOS image.

Objective:

Infrastructure Management

Sub-Objective:

Perform device maintenance


References:

[Cisco > Cisco IOS Configuration Fundamentals Command Reference > C > copy](#)

[Cisco > Cisco IOS Configuration Fundamentals Configuration Guide, Release 12.4 > Part 8: Managing Configuration Files > Managing Configuration Files](#)

QUESTION 20

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



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.

Objective:**Routing Fundamentals****Sub-Objective:**

Configure, verify, and troubleshoot single area and multi-area OSPFv2 for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub, virtual-link, and LSAs)

References:

[Cisco > Support > IP > IP Routing > Technology Information > Technology White Paper > OSPF Design Guide > Document ID: 7039 > DR Election](#)

QUESTION 21

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.

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
```

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot single area and multi-area OSPFv2 for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub, virtual-link, and LSAs)

References:

[Cisco > Cisco IOS IP Routing Protocols Configuration Guide, Release 12.4T > Part 6: OSPF > Configuring OSPF > OSPF Configuration Task List > Enabling OSPF](#)

QUESTION 22

You are the network administrator for your company. You have a Class B address range and are planning for a network that allows 150 hosts per subnet and at least 164 subnets.

Which subnet mask should you use to accomplish the task?

- A. 255.255.192.0
- B. 255.255.255.192
- C. 255.255.255.0
- D. 255.255.255.252

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should use 255.255.255.0 as the subnet mask to allow 150 hosts per subnet and at least 164 subnets. The formulas used to calculate the number of subnets and hosts are:

Number of subnets = $2^{\text{number-of-subnet-bits}}$

Number of hosts per subnet = $2^{\text{number-of-host-bits}} - 2$

Subnet mask in decimal: 255.255.255.0

Subnet mask in binary: 11111111.11111111.11111111.00000000

Number of subnet bits: 8 (binary 1s in the subnet octet of the subnet mask)

Number of host bits: 8 (binary 0s in the subnet mask)

In this scenario, we find that for 255.255.255.0:

Subnets that can be used: $2^8 = 256$

Hosts that can be used: $2^8 - 2 = 254$

The other options do not allow 150 hosts per subnet and at least 164 subnets.

If you use 255.255.192.0 as the subnet mask, then the total number of hosts that can be connected per subnet is 16382 ($2^{14} - 2 = 16382$). However, there will be 4 subnets ($2^2 = 4$).

If you use 255.255.255.192 as the subnet mask, there will be 62 hosts ($2^6 - 2 = 62$).

If you use 255.255.255.252 as the subnet mask, there will be two hosts per subnet ($2^2 - 2 = 2$).

Note: This mask is frequently used for a subnet that connects two routers. In that case, there are two interfaces in the subnet, and thus it is most efficient use of the addressing space. This is also the most efficient way to address a point-to-point serial link.

A note about the formulas: You will always subtract 2 from the number of hosts ($2^{\text{number-of-host-bits}} - 2$) because the all-zeroes bit address is reserved for the network address and the all-ones bit address is reserved for the broadcast address.

Before Cisco IOS Software Release 12.0, it was common practice to subtract 2 from the networks formula ($2^{\text{number-of-subnet-bits}}$) to exclude the all-ones subnet and subnet zero. Today that range is usable, except with some legacy systems. On certain networks with legacy software, you may need to use the previous formula ($2^{\text{number-of-subnet-bits}} - 2$) to calculate the number of valid subnets.

Objective:

Network Fundamentals

Sub-Objective:

Apply troubleshooting methodologies to resolve problems

References:

[Cisco > Technology Support > IP > IP Routing > Design TechNotes > IP Addressing and Subnetting for New Users > Document ID: 13788 > Understanding IP Addresses](#)
[Cisco > Technology Support > IP > IP Addressing Services > Design TechNotes > Subnet Zero and the All-Ones Subnet > Document ID: 13711](#)

QUESTION 23

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
<output omitted>
interface serial 0/0
encapsulation ppp

router2#show run
<output omitted>
interface serial 0/1
```

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.

Objective:

WAN Technologies

Sub-Objective:

Configure and verify PPP and MLPPP on WAN interfaces using local authentication

References:

[Internetworking Technology Handbook > WAN Technologies > Point-to-Point Protocol](#)

QUESTION 24

A user in your network is having trouble accessing resources and the Internet. You decide to examine the partial output of the ipconfig/all command on his machine. The output is shown below:

```
Microsoft Windows [Version 6.1.7601]
Copyright (c) 2009 Microsoft Corporation. All rights reserved.

C:\Users\TroyMcClure > ipconfig/all

Windows IP Configuration

Host Name . . . . . : KREMLIN0120
Primary Dns Suffix . . . . . : kappa.alpha.com
Node Type . . . . . : Hybrid
IP Routing Enabled. . . . . : No
WINS Proxy Enabled. . . . . : No
DNS Suffix Search List. . . . . : kappa.alpha.com


Ethernet adapter Local Area Connection:

Connection-specific DNS Suffix . : triad.rr.com
Description . . . . . : Intel(R) Dual Band Wireless-N 7260
Physical Address. . . . . : F8-16-54-12-E3-69
DHCP Enabled. . . . . : No
Autoconfiguration Enabled . . . . : Yes
IPv4 Address. . . . . : 192.168.1.3(Preferred)
Subnet Mask . . . . . : 255.255.255.0
Default Gateway . . . . . : 192.168.0.1
DNS Servers . . . . . : 192.168.0.50
```

Which of the following statements describes the user's problem?

A. The default gateway address is incorrect

- B. The IP address of the device is incorrect
- C. There is no DNS server configured
- D. IP routing is not enabled

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The IP address of the device is incorrect. It is not in the same subnet as the default gateway address. While it is possible that the default gateway address is incorrect, that is not as likely a reason, given the fact that the DNS server is also in the same IP subnet as the default gateway.

There is a DNS server configured and its IP address is 192.168.0.50. If a DNS server were not configured, this user would be unable to access the Internet, even if all IP addressing problems were resolved.

IP routing is NOT enabled. However, it is not required to be enabled because this device is not acting as a router. The device does not need IP routing enabled to access resources and the Internet if all other IP addressing issues are resolved.

Objective:

Infrastructure Services

Sub-Objective:

Describe DNS lookup operation

References:

[PChuck's Network > Microsoft Windows Networking, Security, and Support > Reading IPConfig and Diagnosing Network Problems](#)

QUESTION 25

Which type of Category 5 unshielded twisted-pair (UTP) cable is used to work as a trunk between two switches?

- A. RJ-45 straight-through
- B. RJ-41 crossover
- C. RJ-11 straight-through
- D. RJ-45 crossover

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

An RJ-45 crossover cable connects two switches. To act as a trunk a trunking protocol such as ISL or 802.1q must be configured on the link. . A trunk is a connection between two switches that is used to carry traffic from multiple VLANs.

In general, the rule to follow when choosing between a straight-through and a crossover cable is:

- When connecting like devices (i.e. router to router, switch to switch), use a crossover cable.
- When connecting dissimilar devices (i.e. switch to router), use a straight-through cable.

The one exception to this rule is when connecting a computer NIC to a router, in which case a crossover cable is used. Be aware, however, that many devices, including network cards in computers, now have the ability to sense automatically when they are connected to a like device and adapt to the connection, making crossover cables unnecessary in those situations.

You should not choose an RJ-45 straight-through cable. The cable type to be used depends on the circuit connection of the hardware. To connect two switches, a crossover cable is required. The difference between a straight-through cable and a crossover cable lies in the location of the wire termination on the two ends of an RJ-45 cable. If the UTP cable wire connects Pin 1 of one side to Pin 1 of other side and Pin 2 to 2 through all eight pins of the RJ 45 connector, the cable is said to be straight-through. On the other hand, if Pin 1 of one side of an RJ-45 cable connects to Pin 3 of the other end, and Pin 2 connects to Pin 6 of the other end, it is known as a crossover cable. A straight-through cable is used to connect a computer's network interface card (NIC) to a hub or switch.

You should not choose an RJ-41 crossover cable. RJ-41 is a single-line universal data jack normally associated with fixed-loss loop (FLL) or programmed (P) modems. It is not used between switches.

You should not choose an RJ-11 straight-through cable type. RJ-11 UTP cables have four pins and are used to connect voice instruments. RJ-11 UTP cables are not intended for connecting computers and transferring data. They are commonly used for telephones and modems.

Note: Cisco switches have an auto-mdix feature that notices when the wrong cabling pinouts are used, and readjusts the switch's logic so that the cable will work.

Objective:

Network Fundamentals

Sub-Objective:

Select the appropriate cabling type based on implementation requirements

References:

[Cisco Documentation > Internetwork Design Guide > Designing Switched LAN Internetworks > Technologies for Building Switched LAN Internetworks](#)

[Cisco > Troubleshooting Technotes > Cisco 7000 Series Routers > Cabling Guide for Console and AUX Ports > Types of RJ-45 Cabling](#)

QUESTION 26

A router is running a classful routing protocol. Which command will enable this router to select a default route when routing to an unknown subnet of a network for which it knows the major network?

- A. ip classless
- B. no ip classless
- C. auto-summary
- D. no auto-summary

Correct Answer: A

Section: (none)


Explanation

Explanation/Reference:

Explanation:

The ip classless command causes a routing protocol to change its default behavior of discarding any traffic that is bound for unknown subnets of a known classful network. If the command is enabled, the router tries to match the most number of bits possible against the route in its routing table. Alternatively, the router will use the default route rather than dropping the packet.

For an example of this behavior, examine the diagram below. The ip route 0.0.0.0 0.0.0.0 serial 0/0 command has been issued on Router B. If the 25.1.6.0/24 network is unknown to Router B, then under normal circumstances, Router B would NOT use its configured default route. Instead, it would drop any packets addressed to that unknown network, because when a router knows a route to a major classful network or its subnets (in this case, 25.1.5.0/30 and 25.1.1.0/24), it will not use a statically configured default route to forward traffic to an unknown subnet of that network (in this case 25.1.6.0/24). In the scenario described in the diagram, Router B will drop the packet. However, if the ip classless command has been executed, it will use the default route and send the traffic to Router A.



The `ip classless` command is a global configuration mode command enabled by default in Cisco IOS version 12.0 and later. If the default route is learned from IS-IS or OSPF, as opposed to being statically configured as in the above example, the `ip classless` command is not necessary for the router to use the default route.

The `no ip classless` command on routers will disable the forwarding of packets destined to an unknown subnet of a known classful network. Therefore, it is an incorrect option.

The `auto-summary` command is used to allow automatic summarization of subnet routes into network-level routes. This is a command executed in router configuration mode.

Classless routing protocols such as Routing Information Protocol version 2 (RIPv2) and Enhanced Interior Gateway Routing Protocol (EIGRP) perform automatic route summarization at classful boundaries. The `no auto-summary` command is used to turn off this route summarization.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot IPv4 and IPv6 static routing

References:

[Cisco > Articles > Cisco Certification > CCNP > CCNP Self-Study: Advanced IP Addressing](#)

[Cisco > Cisco IOS IP Addressing Services Command Reference > IP Addressing Commands > `ip classless`](#)

QUESTION 27

Which of the following commands will enable a global IPv6 address based on the Modified EUI-64 format interface ID?

- A. `ipv6 address 5000::2222:1/64`
- B. `ipv6 address autoconfig`
- C. `ipv6 address 2001:db8:2222:7272::72/64 link-local`
- D. `ipv6 enable`

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

To configure the interface to create a global IPv6 address based on the Modified EUI-64 format interface ID, you must enable stateless autoconfiguration. In stateless autoconfiguration, the interface will receive the network prefix from the router advertisement (RA) and generate a full IPv6 address by spreading the 48-bit MAC address of the interface across 64 bits to complete the address. This can all be done simply by executing the `ipv6 address autoconfig` command at the interface configuration prompt.

The command `ipv6 address 5000::2222:1/64` is used to manually assign a full IPv6 address to the interface without using stateless autoconfiguration or the `eui-64` keyword to manually specify the first 64 bits and allow the last 64 bits to be generated from the MAC address of the interface.

The command `ipv6 address 2001:db8:2222:7272::72/64 link local` is used to configure a link-local address manually without allowing the system to generate one from the MAC address, which is the default method.

The command `ipv6 enable` is used to allow the system to generate a link-local address from the MAC address. Because this is the default behavior, the command is not required if any other `ipv6` commands have been issued. Regardless of how many manual IPv6 addresses you configure, a link local address is always generated by default.

Objective:

Network Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot IPv6 addressing

References:

[Cisco > Product Support > Security > Cisco ASA 5500-X Series Firewalls > Configure > Configuration Guides > Cisco Security Appliance Command Line Configuration Guide, Version 7.2 > Chapter: Configuring IPv6 > Enforcing the Use of Modified EUI-64 Interface IDs in IPv6 Addresses](#)
[Cisco > Support > Cisco IOS IPv6 Command Reference > ipv6 address](#)

QUESTION 28

Which of the following commands is used to verify the link-local, global unicast, and multicast addresses of an IPv6 router?

- A. `show ipv6 neighbors` (only link-local addresses)
- B. `show ipv6 route`
- C. `show ipv6 protocols`
- D. `show ipv6 interface`

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The `show ipv6 interface` command is used to verify the link-local, global unicast, and multicast addresses assigned to an IPv6-enabled router interface. The `show ipv6 interface` command displays information regarding that interface, such as the physical state, MTU, and IPv6 enable/disable state.

Here is the partial output of the `show ipv6 interface` command on an IPv6-enabled router named `rtrA`:

```

rtrA# show ipv6 interface FastEthernet 0/1
FastEthernet0/1 is up, line protocol is up
IPv6 is enabled, link-local address is
FE80::6339:7BFF:FE5D:A031/64
Global unicast address(es):
2001:7067:90D1:1::1, subnet is 2001:7067:90D1:1/64
Joined group address(es):
FF02::1
FF02::2
FF02::1:FF5D:A031
MTU is 1500 bytes
<output omitted>

```

In the sample output, you can see that the Fa0/1 interface of rtrA has the link-local address FE80::6339:7BFF:FE5D:A031/64 and the global unicast address 2001:7067:90D1:1::1. The global unicast address is not in EUI-64 format because when the ipv6 address command was issued, the eui64 keyword was not used. If the EUI-64 format had been specified with the eui64 keyword, the global unicast address would have been 2001:7067:90D1:1:6339:7BFF:FE5D:A031.

An IPv6-enabled interface has not only a link-local and global unicast address, but also one or more multicast addresses. A multicast address is an IPv6 address that has the prefix FF00::/8. These addresses are assigned to interfaces of different nodes such that they appear as a logical group. This implies that when a packet is destined for a multicast address, that packet is delivered to all the interfaces that have the same multicast address. The various multicast groups are as follows:

- FF02::1 Indicates the group of all the nodes on the local segment
- FF02::2 Indicates the group of all the routers on the local segment
- FF02::1:FF00:0/104 Indicates a solicited-node multicast group for every unicast or anycast address assigned to the interface

You can also notice in the sample output that the Fa0/1 interface belongs to three multicast groups: FF02::1, FF02::2, and FF02::1:FF5D:A031. The first two multicast groups refer to the all-host and all-router multicast groups, respectively. The third group, FF02::1:FF5D:A031, is the solicited-node multicast address. This address is created for every unicast or anycast address. A solicited-node multicast address is determined by assigning the least significant 24 bits of the unicast address to the least significant 24 bits of the FF02::1:FF00:0 address.

The show ipv6 neighbors command displays the link-local /global unicast addresses of the neighbors, including other information such as state and the next-hop interface.

The show ipv6 route command is used to view the IPv6 routing table on the router. This command displays the prefixes, administrative distance, metric, and next-hop addresses for various IPv6 networks.

The show ipv6 protocols command is used to view the active routing protocols for IPv6 on the router. This command shows the interfaces, redistribution status, and summarization status about each of the routing protocols enabled on the router.

Objective:

Network Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot IPv6 addressing

References:

[Cisco IOS IPv6 Command Reference > show ipv6 eigrp topology through show ipv6 nat statistics > show ipv6 interface](#)
[Cisco IOS IPv6 Command Reference > show ipv6 nat translations through show ipv6 protocols > show ipv6 neighbors](#)
[Cisco IOS IPv6 Command Reference > show ipv6 nat translations through show ipv6 protocols > show ipv6 protocols](#)
[Cisco > Products & Services > Cisco IOS and NX-OS Software > Cisco IOS Technologies > IPv6 > Product Literature > White Papers > Cisco IOS IPv6 Multicast Introduction](#)
[Cisco > IPv6 Implementation Guide, Release 15.2M&T > Implementing IPv6 Multicast](#)

QUESTION 29

A host is powered up, but the connected switch port does not turn amber or green.

Which of the following methods would you use to troubleshoot the situation? (Choose three. Each answer is a complete solution.)

- A. Ensure the switch is powered up.
- B. Reinstall Windows on the workstation.
- C. Reseat the cable.
- D. Ensure that the cable is straight-through.
- E. Ensure that the cable is crossover.

Correct Answer: ACD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A black or unlit switch port LED is symptomatic of a Layer 1 problem. The port LED should first turn amber and then turn solid green when a host is powered up. The amount of time it takes to turn solid green will depend on the Spanning Tree Protocol configuration. If the LED is unlit, you should ensure that the switch is powered up and that a straight-through cable is used to connect a switch port to a host, such as a workstation or a printer. If the switch is powered up and a straight-through cable is used, reseat the cable to ensure a firm connection.

Reinstalling Windows on the workstation will not help because this is a Layer 1 problem having to do with the switch having power or the use of proper cabling.

You should not ensure that the cable is crossover, because straight-through (patch) cables are used to connect switch ports to hosts.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Troubleshoot interface and cable issues (collisions, errors, duplex, speed)

References:

[Cisco > Catalyst 2960 Switch Hardware Installation Guide > LEDs](#)

QUESTION 30

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.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot STP protocols

References:

[Cisco > Technology Support > LAN Switching > Spanning Tree Protocol > Technology White Paper > Understanding Rapid Spanning Tree Protocol \(802.1w\)](#)
[Cisco Support > Technology Support > LAN Switching > Spanning Tree Protocol > Troubleshoot and Alerts > Troubleshooting TechNotes > Understanding Rapid Spanning Tree Protocol \(802.1w\)](#)
[CCNA Routing and Switching Complete Study Guide: Exam 100-105, Exam 200-105, Exam 200-125, 2nd Edition](#), Chapter 2: LAN Switching Technologies - Configure, verify, and troubleshoot STP protocols

QUESTION 31

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 0
Serial0/0 0 0 0 0 0 0 0 0 0
FastEthernet0/1 0 0 0 0 0 0 0 0 0
Serial0/1 0 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.

Objective:

Routing Fundamentals

Sub-Objective:

Troubleshoot basic Layer 3 end-to-end connectivity issues

References:

[Cisco > Cisco IOS Interface and Hardware Component Command Reference > show interfaces summary](#)

QUESTION 32

Which of the following is NOT a VLAN Trunking Protocol (VTP) mode of operation?

- A. client
- B. server
- C. virtual
- D. transparent

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Virtual is not a valid VTP mode of operation. There are three different VTP modes of operation: client, server, and transparent.

In client mode, a switch can synchronize VLAN information with the domain and forward advertisements. However, VLANs cannot be created, deleted, or modified from a switch in client mode. Also, a client mode switch does not save VLAN information in non-volatile Random Access Memory (NVRAM). It is stored in Flash in a file called vlan.dat.

In server mode, a switch synchronizes the VLAN information with the domain, sends and forwards advertisements, and can create, delete, or modify VLANs. In server mode, VLAN information is stored in Flash in a file called vlan.dat.

In transparent mode, a switch does not synchronize its VLAN configuration with the domain, but it forwards advertisements. VLANs can be created, deleted, or modified locally and VLAN configuration is saved in both the running-config file in RAM and in flash in a file called vlan.dat.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot VLANs (normal/extended range) spanning multiple switches

References:

[Cisco > Catalyst 4500 Series Switch Cisco IOS Software Configuration Guide, 12.2\(25\) > Understanding and Configuring VLANs, VTP, and VMPS](#)

QUESTION 33

You are advising a client on the options available to connect a small office to an ISP.

Which of the following is an advantage of using an ADSL line?

- A. it uses the existing cable TV connection
- B. it uses the existing phone line
- C. you receive a committed information rate (CIR) from the provider
- D. the upload rate is as good as the download rate

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

xDSL lines, including the ADSL variant, use the existing phone line and as such make installing only a matter of hooking up the DSL modem to the line.

It does not use the use the existing cable TV connection. This is a characteristic of using a cable modem rather than ADSL.

You do not receive a committed information rate (CIR) from the provider. CIR is provided with a frame relay connection.

The upload rate is NOT as good as the download rate with asynchronous DSL (ADSL). The download rate is significantly better than the upload rate. Symmetric Digital Subscriber Line (SDSL) is a version of DSL that supplies an equal upload and download rate, but that is not the case with ADSL.

Objective:

WAN Technologies

Sub-Objective:
Describe WAN access connectivity options

References:

[Cisco Documentation > Internetworking Technology Handbook > Introduction to WAN Technologies > DSL](#)

QUESTION 34

Which of the following methods will ensure that only one specific host can connect to port F0/1 on a switch?

- A. Configure port security on F0/1 to forward traffic to a destination other than that of the MAC address of the host.
- B. Configure the MAC address of the host as a static entry associated with port F0/1.
- C. Configure port security on F0/1 to accept traffic only from the MAC address of the host.
- D. Configure an inbound access control list on port F0/1 limiting traffic to the IP address of the host.
- E. Configure port security on F0/1 to accept traffic other than that of the MAC address of the host.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

To limit connections to a specific host, you should configure port security to accept traffic only from the MAC address of the host. By default, an unlimited number of MAC addresses can be learned on a single switch port, whether it is configured as an access port or a trunk port. Switch ports can be secured by defining one or more specific MAC addresses that should be allowed to connect, and by defining violation policies (such as disabling the port) to be enacted if additional hosts try to gain a connection.

The following example secures a switch port by manually defining the MAC address of allowed connections:

```
switch(config-if)# switchport port-security  
switch(config-if)# switchport port-security mac-address 00C0.35F0.8301
```

The first command activates port security on the interface, while the second command statically defines the MAC address of 00c0.35f0.8301 as an allowed host on the switch port.

The mac-address-table static command assigns a permanent MAC address to the port, but does not prevent any other MAC addresses from being associated with the port. . The command below would assign the MAC address 0050.3e8d.62bb to port 15 on the switch:

```
switch(config)# mac-address-table static 0050.3e8d.6400 interface fastethernet0/15
```

You should not configure port security on F0/1 to forward traffic to a destination other than that of the MAC address of the host. Traffic from other hosts should be rejected, not forwarded or accepted. For the same reason, you should not configure port security on F0/1 to accept traffic other than that of the MAC address of the host.

You cannot configure an inbound access control list on port F0/1 limiting traffic to the IP address of the host. It is impossible to filter traffic based on IP addresses on a Layer 2 switch.

Objective:

Infrastructure Security

Sub-Objective:

Configure, verify, and troubleshoot port security

References:

[Cisco > Catalyst 6500 Release 12.2SXH and Later Software Configuration Guide > Configuring Port Security > Enabling Port Security](#)

QUESTION 35

Which two statements are TRUE of Internet Protocol (IP) addressing? (Choose two.)

- A. Public addresses are registered with the Internet Assigned Numbers Authority (IANA).
- B. These addresses are publicly registered with the Internet Service Provider (ISP).
- C. Through a public IP address, you can access another computer on the Internet, such as a Web server.
- D. The ranges of public IP addressing are 10.0.0.0 to 10.255.255.255, 172.16.0.0 to 172.31.255.255, and 192.168.0.0 to 192.168.255.255.
- E. Private addresses are allocated by the Internet Assigned Numbers Authority (IANA).

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Public addresses are publicly registered with the Internet Assigned Numbers Authority (IANA). Through a public IP address, you can access an Internet computer like a Web server.

The following statements are true of public IP addressing:

- These addresses are publicly registered with the Internet Assigned Numbers Authority (IANA)
- Through a public IP address, you can access another Internet computer, such as a Web server.
- Other people on the Internet can obtain information about or access to your computer via a public IP address.
- Public IP addresses are visible to the public.

The option stating that public IP addresses are publicly registered with the Internet Service Provider (ISP) is incorrect. Public IP addresses are registered with the Internet Assigned Numbers Authority (IANA). Since 1998, InterNIC has been primarily responsible for allocating domain names and IP addresses under the governance of the Internet Corporation for Assigned Names and Numbers (ICANN) body, a U.S. non-profit corporation that was created to oversee work performed by the Internet Assigned Numbers Authority (IANA).

The option stating that 10.0.0.0 to 10.255.255.255, 172.16.0.0 to 172.31.255.255, and 192.168.0.0 to 192.168.255.255 are the range of public IP addressing is incorrect. These ranges belong to private IP addressing.

The option stating that private addresses are allocated by the IANA is incorrect. Private IP addresses are not managed, but are used by private organizations as they see fit. The IANA is governed by ICANN, and its primary role is to allocate overseas global IP addresses from the pools of unallocated addresses, as well as DNS root zone management.

Objective:

Network Fundamentals

Sub-Objective:

Describe the need for private IPv4 addressing

References:

<http://www.debianadmin.com/private-and-public-ip-addresses-explained.html>

QUESTION 36

Which type of network uses Carrier Sense Multiple Access with Collision Detection (CSMA/CD) as an access method?

- A. Token Ring
- B. LocalTalk
- C. 100VG-AnyLan
- D. Ethernet

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Ethernet networks use CSMA/CD as an access method. In CSMA/CD, if a device wants to send a frame in the network, it first determines if the network is free. If the network is not free, the node will wait before sending the frame into a network. If the network is free, it sends the frame; if another device sends a frame simultaneously then their signals or frames collide. When the collision is detected, both packets wait for a random time before retrying.

The following statements are true regarding CSMA/CD:

- CSMA/CD is required for shared collision domains, such as when hosts are connected via hubs. (Hubs are Layer 1 devices, and thus do not create collision domains.)
- CSMA/CD networks normally operate in half-duplex mode, since in a shared collision domain, a host cannot send and receive data at the same time.
- CSMA/CD is not required when connected to non-shared (private) collision domains, such as when hosts are connected to dedicated switch ports.
- Switches create dedicated collision domains, so devices can operate in full-duplex mode.

Token Ring is incorrect because Token Ring uses token passing as the access method.

LocalTalk is incorrect because LocalTalk uses CSMA/CA (Collision Avoidance) as the access method.

100VG-AnyLan is incorrect because 100VG-AnyLan uses demand priority as the access method.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Interpret Ethernet frame format

References:

[Cisco > Internetworking Technology Handbook > Introduction to LAN Protocols > LAN Media-Access Methods](#)

QUESTION 37

Two catalyst switches on a LAN are connected to each other with redundant links and have Spanning Tree Protocol (STP) disabled.

What problem could occur from this configuration?

- A. It may cause broadcast storms.
- B. All ports on both switches may change to a forwarding state.
- C. It may cause a collision storm.
- D. These switches will not forward VTP information.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The configuration in the scenario may cause broadcast storms. When there are redundant links between two switches, it is recommended that you enable Spanning Tree Protocol to avoid switching loops or broadcast storms. Loops occur when there is more than one path between two switches. STP allows only one active path at a time, thus preventing loops. A broadcast storm occurs when the network is plagued with constant broadcasts. When the switches have redundant links, the resulting loops would generate more broadcasts, eventually resulting in a complete blockage of available bandwidth that could bring the complete network down.

This situation is referred to as a broadcast storm.

The option stating that all ports on both switches may change to a forwarding state is incorrect. Forwarding is a port state that is available when using STP. When STP is disabled, the switch cannot change the STP states of its ports.

The option stating that the switches will not forward VLAN Trunking Protocol (VTP) information is incorrect. Enabling or disabling STP does not have a direct effect on VTP messages.

The term collision storm is not a valid term.

Objective:

LAN Switching Fundamentals

Sub-Objective:


Configure, verify, and troubleshoot interswitch connectivity

References:

[Cisco > Support > Technology Support > LAN Switching > Ethernet > Design > Troubleshooting LAN Switching Environments > Document ID: 12006 > Spanning Tree Protocol](#)

QUESTION 38

Users on the LAN are unable to access the Internet. How would you correct the immediate problem?



```
Router# show ip interface brief
```

```
Interface IP-Address OK? Method Status Protocol
FastEthernet 0/0 unassigned YES unset down down
FastEthernet 0/1 172.16.1.254 YES NVRAM up up
Serial0/0 200.16.4.25 YES NVRAM administratively down down
Serial0/1 unassigned YES unset down down
```

- A. Configure a bandwidth on the serial interface.
- B. Perform a no shutdown command on the serial interface.
- C. Configure a private IP address on the Fastethernet0/0 LAN interface.
- D. Change the IP address on the serial interface.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The output indicates that the serial interface leading to the Internet is administratively down. All router interfaces are disabled by default due to the presence of a shutdown command in the running configuration. The no shutdown command removes this configuration, and the interface becomes active. The command sequence is:

```
Router(config)# interface serial0/0
Router(config-if)# no shutdown
```

Although it was not the problem in the scenario, the S0/0 interface could also cause an error if it is configured as shown in this output:

Interface IP-Address OK? Method Status Protocol

Serial0/0 200.16.4.25 YES NVRAM up down

In this example, the S0/0 interface has been enabled, and while there is Layer 1 connectivity (the Status column), Layer 2 is not functioning (the Protocol column). There are two possible reasons for this result:

- Interface S0/0 is not receiving a clock signal from the CSU/DSU (if one is present).
- The encapsulation type configured on S0/0 does not match the type configured on the other end of the link (if the other end is a router).

Configuring a bandwidth on the serial interface is incorrect because the output indicates the interface is administratively down, which does not pertain to bandwidth.

Configuring a private IP address on the Fastethernet0/0 LAN interface is incorrect because the output indicates the problem is with the disabled serial interface.

The IP address on the serial interface may or may not be valid, but it is not the immediate cause of the connectivity problem. The serial interface is disabled.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Troubleshoot interface and cable issues (collisions, errors, duplex, speed)

References:

[Cisco > Support > Administrative Commands > shutdown](#)

QUESTION 39

You are the network administrator for your company and have configured Cisco Discovery Protocol (CDP) in your network. You recently noticed that when devices send large numbers of CDP neighbor announcements, some devices are crashing. You decide to disable CDP on the router.

Which command should you use to achieve the objective?

- A. no cdp run
- B. set cdp disable
- C. no cdp enable
- D. no cdp advertise-v2

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should use the no cdp run command to disable CDP on the router. Due to a known vulnerability regarding

the handling of CDP by Cisco routers and switches when devices send large numbers of CDP neighbor announcements, some devices can crash or cause abnormal system behavior. To overcome this problem, you can disable CDP for the entire router by using the no cdp run command.

You cannot use the set cdp disable command to disable CDP on the router. This command disables CDP on an entire Catalyst switch.

You cannot use the no cdp enable command to disable CDP on the router. This command disables CDP on a specific interface.

You cannot use the no cdp advertise-v2 command to disable CDP on the router. This command disables CDPv2 advertisements.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure and verify Layer 2 protocols

References:

[Cisco > Support > Using Cisco Discovery Protocol](#)

[Cisco > Support > Technology Support > Network Management > Cisco's Response to the CDP Issue > Document ID: 13621](#)

QUESTION 40

Which of the following topologies is used in Wide Area Networks (WANs)?

- A. FDDI
- B. CDDI
- C. SONET
- D. Token Ring

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Synchronous Optical Network (SONET) is the standard topology for fiber optic networks. Developed in 1980s, SONET can transmit data at rates of up to 2.5 gigabits per second (Gbps).

All other options are incorrect because they are LAN topologies, not WAN topologies.

Fiber Distributed Data Interface (FDDI) specifies a 100-Mbps dual-ring fiber optics-based token-passing LAN. FDDI is typically implemented for high-speed LAN backbones because of its support for high bandwidth.

Copper Distributed Data Interface (CDDI) is copper version of FDDI. They differ only in that FDDI can span longer distances than CDDI due to the attenuation characteristics of copper wiring.

Token Ring/IEEE 802.5 LAN technology was developed by IBM in 1970. Token-ring LAN technology is based on token-passing, in which a small frame, called a token, is passed around the network. Possession of the token grants the node the right to transmit data. Once the data is transmitted, the station passes the token to the next end station.

Objective:

Network Fundamentals

Sub-Objective:

Compare and contrast network topologies

References:

QUESTION 41

What is the primary benefit of the Virtual Local Area Network (VLAN) Trunking Protocol (VTP)?

- A. broadcast control
- B. frame tagging
- C. inter-VLAN routing
- D. consistent VLAN configuration across switches in a domain

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

VTP manages configured VLANs across a switched network and maintains consistency of VLAN information throughout a VTP domain. When an administrator adds, deletes, or renames VLANs, VTP propagates this information to all other switches in the VTP domain. This makes the process of VLAN changes a plug-and-play activity. This protocol was developed by, and remains proprietary to Cisco Systems.

Broadcast control is not the primary benefit of VTP. Broadcast control is achieved by using VLANs. VLANs segment the network into logical broadcast domains. This helps in the reduction of unnecessary traffic over the network and optimizes the available bandwidth use. VTP pruning helps reduce broadcast and unknown unicast over VLAN trunk links. However, this is not the primary benefit of VTP.

Frame tagging is required for VLAN identification as frames traverse trunk links in a switch fabric. Inter-Switch Link (ISL) and IEEE 802.1q are the two methods of frame tagging available on Cisco devices. ISL is proprietary to Cisco, whereas IEEE 802.1q is a standard method. VTP is not a frame tagging method.

Inter-VLAN routing is achieved by an Open Systems Interconnect (OSI) Layer 3 device (Router). Inter-VLAN routing is not a benefit of VTP.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot VLANs (normal/extended range) spanning multiple switches

References:

[Cisco > Support > LAN Switching > Virtual LANS / VLAN Trunking Protocol \(VLANs/VTP\) > Configure > Configuration Examples and Technotes > Configuring VLAN Trunk Protocol \(VTP\) > Document ID: 98154](#)
[Cisco > Catalyst 4500 Series Switch Cisco IOS Software Configuration Guide, 12.2\(25\) > Understanding and Configuring VLANs, VTP, and VMPS](#)

QUESTION 42

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

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

EIGRP is a Cisco-proprietary routing protocol, and does not support multiple vendor environments.

EIGRP is a classless routing protocol, and thus supports variable length subnet masks (VLSM).

EIGRP routers build a neighbor table in memory, and use a multicast-based neighbor discovery mechanism.

EIGRP routers send partial updates when there are network events.

The following are features offered by EIGRP:

- Fast convergence
- Partial updates
- Neighbor discovery mechanism
- VLSM
- Route summarization
- Scalability

Objective:

Routing Fundamentals

Sub-Objective:

Compare and contrast distance vector and link-state routing protocols

References:

[Cisco > Support > IP > IP Routing > Design > Design TechNotes > Introduction to EIGRP > Document ID: 13669](#)

QUESTION 43

Which set of Cisco Internetwork Operating System (IOS) commands is used on Cisco routers to set a password for Telnet lines?

- A. `router(config-router)# line vty 0 4`
`router(config-line)# login`
`router(config-line)# password password`
- B. `router(config)# line telnet 0 4`
`router(config-line)# login`
`router(config-line)# password password`
- C. `router(config)# line aux 0`
`router(config-line)# login`
`router(config-line)# password password`
- D. `router(config)# line vty 0 4`
`router(config-line)# login`
`router(config-line)# password password`

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The following commands are used on Cisco routers to set a password for Telnet lines:

```
router(config)# line vty 0 4  
router(config-line)# login  
router(config-line)# password password
```

An explanation of the commands is below:

`router(config)# line vty 0 4`: Enters line configuration mode for virtual terminal lines 0 to 4.

`router(config-line)# login`: Ensures that any remote access is prompted for a password.

`router(config-line)# password password`: Sets a password of "password" for VTY lines.

Assigning a password to the VTY lines is required for remote connections to the device to be possible. If a

password has not been configured the following error message will be generated when the connection is attempted:

```
Password required but not set
```

```
[Connection to foreign host 106.5.5.1 closed by foreign host]
```

Configuring a VTY password and requiring the password (accomplished with the login command) is good first step in securing Telnet access to the device. Another step that can enhance the security of remote access to the device would be to apply an access list to the VTY lines with the access-class command.

The command sequence which begins with router(config-router)# line vty 0 4 is incorrect because the line vty 0 4 command should be executed in global configuration mode, not routing protocol configuration mode.

The line telnet 0 4 command is incorrect because this is not a valid Cisco IOS command.

The line aux 0 command is incorrect because this allows you to configure the properties of the Auxiliary port, as opposed to the incoming Telnet (VTY) lines.

Objective:

Infrastructure Management

Sub-Objective:

Configure and verify device management

References:

[Cisco > Support > Technology Support > IP > IP Addressing Services > Design > Design TechNotes > Cisco Guide to Harden Cisco IOS Devices > Document ID: 13608](#)

[Cisco > Support > End-of-sale and End-of-life Products > Cisco IOS Software Releases 11.0 > Configuration Examples and TechNotes > Telnet, Console and AUX Port Passwords on Cisco Routers Configuration Example](#)

QUESTION 44

In which of the following networks does the address 192.168.54.23/27 reside?

- A. 192.168.54.0
- B. 192.168.54.8
- C. 192.168.54.4
- D. 192.168.54.16

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

When a class C address such as 192.168.54.0 is subnetted with a /27 mask, the subnet mask in dotted decimal format is 255.255.255.224. This means that the interval between the network IDs of the resulting subnets is 32. The resulting network IDs are as follows:

192.168.54.0
192.168.54.32
192.168.54.64
192.168.54.96 and so on.

Therefore, the address 192.168.54.23 resides in the 192.168.54.0 subnet. The address 192.168.54.0 is called a network ID or, alternately, a subnet address. It represents the subnet as a group and will be used in the routing tables to represent and locate the subnet.

Neither the first address (192.168.54.0, the network ID) nor the last address (192.168.54.31, the broadcast

address) in any resulting subnet can be used. Therefore, the addresses in this range are 192.168.54.1 through 192.168.54.30, which includes the 192.168.54.23 address.

192.168.54.8 would only be a network ID if the mask were /29, which would result in an interval of 8 between network IDs. However, even if a /29 mask were used, the 192.168.54.23 address would not fall in its range. The address range for a /29 mask would be 192.168.54.9 through 192.168.54.14.

Similarly, 192.168.54.4 would only be a network ID for a /30 mask, which would result in an interval of 4 between network IDs. But even if a /30 mask were used, the 192.168.54.23 address would not fall in its range. The address range for a /30 mask would be 192.168.54.5 through 192.168.54.6.

192.168.54.16 could be a network ID if the mask were /28, /29 or /30, but not with a /27 mask.

Objective:

Network Fundamentals

Sub-Objective:

Apply troubleshooting methodologies to resolve problems

References:

[Cisco > Support > IP Routing > Design TechNotes > Document ID: 13788 > IP Addressing and Subnetting for New Users](#)

QUESTION 45

What is the purpose of frame tagging in Virtual LAN (VLAN) configurations?

- A. inter-VLAN routing
- B. encryption of network packets
- C. frame identification over trunk links
- D. frame identification over access links

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Frame tagging is used when VLAN traffic travels over a trunk link. Trunk links carry frames for multiple VLANs. Therefore, frame tags are used for identification of frames from different VLANs. Inter Switch Link (ISL) and Institute of Electrical and Electronics Engineers (IEEE) 802.1q are the two frame tagging methods supported on Cisco devices.

The purpose of frame tagging is not inter-VLAN routing. A Layer 3 device, such as a router or multilayer switch, is used for inter-VLAN routing. To configure inter-VLAN routing a logical or subinterface for each VLAN must be created on the single physical interface used to connect to the switch. An IP address is NOT applied to the physical interface; instead, each subinterface is configured with an IP address that will become the default gateway of all devices residing in that VLAN. Consequently, each subinterface and its VLAN devices must reside a different subnet as well. If a subinterface on the router is NOT configured with an IP address that resides in the same network as the hosts that reside in the VLAN that the subinterface serves, the hosts in that VLAN will be isolated from the other VLANs. The hosts in the VLAN served by the subinterface should also use this address as their default gateway, or the hosts in the VLAN will likewise be isolated from the other VLANs.

To verify the IP address of the subinterface, execute the show interfaces subinterface ID command. As shown below, the IP address will appear in line 3 of the output. Compare this IP address with the IP address set as the default gateway of each host in the VLAN served by the subinterface. They should be the same, and the IP address of the hosts should be in the same subnet as this address as well.

```
router# show interfaces fastEthernet 0/0.1
FastEthernet0/0.1 is up, line protocol is up
Hardware is AmdFE, address is 0003.e36f.41e0 (bia 0003.e36f.41e0)
```

```
Internet address is 10.10.10.1/24
MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ISL Virtual LAN, Color 1.
ARP type: ARPA, ARP Timeout 04:00:00
```

Frame tagging does not provide encryption of network packets. Packets are transmitted unencrypted unless the network device or the application uses an additional encryption mechanism. A Virtual Private Network (VPN) is a popular solution for providing encrypted network communication.

An access link is a connection between a switch and an end-user computer with a normal Ethernet Network Interface Card (NIC). On these links, Ethernet frames are transmitted without frame tagging.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot VLANs (normal/extended range) spanning multiple switches

References:

[Cisco > Cisco IOS LAN Switching Configuration Guide, Release 12.4 > Part 1: Virtual LANs > Routing Between VLANs Overview](#)

[Cisco > Technology Support > LAN Switching > Virtual LANs/VLAN Trunking Protocol \(VLANs/VTP\) > Design > Design Technotes > Inter-Switch Link and IEEE 802.1Q Frame Format > Document ID: 17056](#)

QUESTION 46

The output of the show ip route command is given:

```
Router# show ip route
```

```
Gateway of last resort is 10.119.254.240 to network 10.140.0.0
O 172.16.0.0 [110/5] via 10.19.24.6, 0:01:00, Ethernet2
B 172.17.12.0 [200/128] via 10.19.24.24, 0:02:22, Ethernet2
O 172.71.13.0 [110/5] via 10.19.24.6, 0:00:59, Ethernet2
O 10.13.0.0 [110/5] via 10.19.24.6, 0:00:59, Ethernet2
```

What does the value 110 in the output represent?

- A. The administrative distance of the information source
- B. The metric to the route
- C. The type of route
- D. The port number of the remote router

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The value 110 in the output represents the administrative distance (AD) of the information source.

Administrative distance is used by Cisco routers to select the most trustworthy source of routing information for a particular route. Every routing protocol has a default administrative distance, and if more than one routing protocol is providing route information about a route, the protocol with the lowest AD will be selected to populate the routing table. The following table shows the AD values for different routing protocols:

IP Route	Default AD value
Connected interface	0
Static route directed to an connected interface	0
Static route directed to an IP address	1
Enhanced Interior Gateway Routing Protocol (EIGRP) summary route	5
External Border Gateway Protocol (BGP) route	20
Internal Enhanced Interior Gateway Routing Protocol (EIGRP) route	90
Interior Gateway Routing Protocol (IGRP) route	100
Open Shortest Path First (OSPF) route	110
Intermediate System-to-Intermediate System (IS-IS) route	115
Routing Information Protocol (RIP) route	120
Exterior Gateway Protocol (EGP) route	140
On Demand Routing (ODR)	160
External Enhanced Interior Gateway Routing Protocol (EIGRP) route	170
Internal Border Gateway Protocol (BGP) route	200
Unknown origin routes	255

The following is the sample output for the show ip route command:

```
Router# show ip route
```

```
Gateway of last resort is 10.119.254.240 to network 10.140.0.0
O 172.16.0.0 [110/5] via 10.19.24.6, 0:01:00, Ethernet2
B 172.17.12.0 [200/128] via 10.19.24.24, 0:02:22, Ethernet2
O 172.71.13.0 [110/5] via 10.19.24.6, 0:00:59, Ethernet2
O 10.13.0.0 [110/5] via 10.19.24.6, 0:00:59, Ethernet2
```

The following are the fields in the output:

- O: Indicates that the route was discovered using Open Shortest Path First (OSPF).
- B: Indicates that the route was discovered using Border Gateway Protocol (BGP).
- 172.16.0.0: Indicates the address of the remote network.
- 110: Indicates the administrative distance of the route.
- 128: Indicates the metric for the route.
- Via 10.19.24.6: Specifies the address of the next router in the remote network.
- 0:02:22: Indicates the last time the route was updated.
- The metric for the route is also called the cost. In the case of the OSPF routes above, the cost is 5.

The administrative distance for any particular protocol can be changed if you would like to use a routing protocol that is normally not the preferred provider. For example, if you prefer that RIP routes be installed in the routing table rather than OSPF routes, you could change the administrative distance of RIP to a lower value than OSPF (110), as shown below.

```
Router(config)# router rip
Router(config)# distance 100
```

All the other options are incorrect because they do not represent the administrative distance.

Objective:

Routing Fundamentals

Sub-Objective:

Interpret the components of routing table

References:

QUESTION 47

Which of the following represents the correct method of assigning an IP address and default gateway to a switch?

- A. Switch(config)# interface vlan1
Switch(config-if)# ip address 10.0.0.1 255.0.0.0
Switch(config-if)# default-gateway 10.0.0.254
- B. Switch(config)# ip default-gateway 10.0.0.254
Switch(config)# interface vlan1
Switch(config-if)# ip address 10.0.0.1 255.0.0.0
- C. Switch(config)# ip address 10.0.0.1 255.0.0.0
Switch(config)# default-gateway 10.0.0.254
- D. Switch(config)# ip address 10.0.0.1 255.0.0.0
Switch(config)# interface vlan1
Switch(config)# ip default-gateway 10.0.0.254

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

IP addresses are assigned to switches by assigning the address to VLAN 1 using the ip address command, while the default gateway is configured in global configuration mode using the ip default-gateway command. A default gateway is assigned to a Layer 2 switch using the following command syntax, where h.h.h.h is the IP address of the default gateway:

Switch(config)# ip default-gateway h.h.h.h

An IP address is assigned to a Layer 2 switch using the following command syntax, where h.h.h.h is the IP address and m.m.m.m is the subnet mask:

Switch(config)# interface vlan1

Switch(config-if)# ip address h.h.h.h m.m.m.m

Configuring an IP address on a switch is usually accompanied by adding a default gateway as well. Switches do not require an IP address to perform their function on the network. IP addresses are added so that an administrator can make a Telnet connection to the switch to manage the switch. If this Telnet access does not occur on the same local subnet with the switch, which is unlikely, or if the administrator is trying to Telnet to the switch using a host that resides a VLAN other VLAN1 (the management VLAN) the absence of a gateway address will render the switch incapable of answering Telnet connection attempts. Therefore, a gateway address is usually required on the switch to make a telnet connection.

The following command set is incorrect because the command setting the default gateway must be executed in global configuration mode, not in configuration mode, for VLAN1:

Switch(config)# interface vlan1

Switch(config-if)# ip address 10.0.0.1 255.0.0.0

Switch(config-if)# default-gateway 10.0.0.254

The following command set is incorrect because the IP address must be configured in configuration mode for VLAN1, not global configuration mode:

Switch(config)# ip address 10.0.0.1 255.0.0.0

Switch(config)# default-gateway 10.0.0.254

The following command set is incorrect because an IP address must be configured in configuration mode for VLAN1. Also, if you executed the command interface vlan1, the prompt would change to Switch(config-if)#. Once it did, that would be an incorrect mode for entering the default gateway.

```
Switch(config)# ip address 10.0.0.1 255.0.0.0
Switch(config)# interface vlan1
Switch(config)# ip default-gateway 10.0.0.254
```

Objective:

Infrastructure Management

Sub-Objective:

Configure and verify device management

References:

QUESTION 48

Which statement best describes a converged network?

- A. a network with real-time applications
- B. a network with a mix of voice, video, and data traffic
- C. a network with a mix of voice and video traffic
- D. a network with mix of data and video traffic

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A converged network is a combination of voice, video, and data traffic. Network convergence is a migration from maintaining multiple service-specific networks, namely data voice and video, to a single IP-based network. All services are delivered on the same network, reducing infrastructure costs. Despite the benefits that network convergence provides, it is highly susceptible to network delays, especially for real-time traffic.

Converged networks frequently face the following problems:

- Bandwidth: As all the voice and video networks are combined into one universal converged network, bandwidth capacity becomes a priority.
- Packet loss: When links become congested, packets will be dropped. Voice and video traffic are intolerant of dropped packets.
- Delay: Delay represents the time it takes for packets to traverse the network and reach their destinations. While some delay is expected, delay increases when links are over-subscribed.

Voice and video traffic are intolerant of high or variable delay. A packet that arrives late is no better than a packet that does not arrive. Delays can be variable and fixed.

Fixed delays are constant and mostly induced by the computing software of the hardware devices, such as processing delay and packetization delay.

Variable delays, known as jitter, cause problems for voice and video.

Objective:

Network Fundamentals

Sub-Objective:


Compare and contrast network topologies

References:

[Cisco Documentation > Internetworking Technology Handbook > Multiservice Access Technologies](#)

QUESTION 49

Which of the following statements are true with regard to the network shown in the exhibit? (Click the Exhibit(s) button.)



- A. there is one broadcast domain and one collision domain
- B. there is one broadcast domain and four collision domains
- C. there are two broadcast domains and two collision domains
- D. there are two broadcast domains and four collision domains
- E. the hosts in VLAN1 could use IP addresses 192.168.5.4/24 and 192.168.5.5/24 and the hosts in VLAN2 could use IP addresses 192.168.6.1/24 and 192.168.6.2/24
- F. the hosts in VLAN2 could use IP addresses 192.168.5.5/24 and 192.168.6.5/24

Correct Answer: DE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

There are two broadcast domains and four collision domains in the network shown in exhibit. A Virtual LAN (VLAN) is a group of networking devices in the same broadcast domain. A broadcast domain is a group of devices such that when one device in the group sends a broadcast, all the other devices in the group will receive that broadcast. Because there are two VLANs shown in the exhibit, VLAN1 and VLAN2, there are two broadcast domains. A switch will not forward broadcast frames between VLANs.

A collision domain is a domain where two or more devices in the domain could cause a collision by sending frames at the same time. Each switch port is a separate collision domain. Because there are four switch ports in the exhibit, there are four collision domains.

The hosts in VLAN1 could use IP addresses 192.168.5.4/24 and 192.168.5.5/24 and the hosts in VLAN2 could use IP addresses 192.168.6.1/24 and 192.168.6.2/24. Hosts in different VLANs must have IP addresses that are in different subnets.

The other options that offer IP address plans are incorrect because they either place hosts from different VLANs in the same subnet, or place hosts in the same VLAN in different subnets.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot VLANs (normal/extended range) spanning multiple switches

References:

[Cisco > Support > Technology Support > LAN Switching > Layer-Three-Switching and Forwarding > Configure > Configuration Examples and TechNotes > How To Configure InterVLAN Routing on Layer 3 Switches](#)

QUESTION 50

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
Router(config-router)# network 192.168.1.128
- D. Router(config-router)# network 192.168.1.0

Correct Answer: D

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
<output omitted>
router eigrp 36
network 192.168.5.0
```

```
Router B#show run
<output omitted>
router eigrp 36
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.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot EIGRP for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub)

References:

[Cisco > Cisco IOS IP Routing Protocols Configuration Guide, Release 12.4T > Part 3: EIGRP > Configuring EIGR](#)

QUESTION 51

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

Objective:

Infrastructure Security

Sub-Objective:


Verify ACLs using the APIC-EM Path Trace ACL analysis tool

References:

[Performing Path Traces](#)

QUESTION 52

You are the switch administrator for InterConn. The network is physically wired as shown in the diagram. You are planning the configuration of STP. The majority of network traffic runs between the hosts and servers within each VLAN.



You would like to designate the root bridges for VLANS 10 and 20. Which switches should you designate as the root bridges?

- A. Switch A for VLAN 10 and Switch E for VLAN 20
- B. Switch A for VLAN 10 and Switch B for VLAN 20
- C. Switch A for VLAN 10 and Switch C for VLAN 20
- D. Switch D for VLAN 10 and Switch B for VLAN 20
- E. Switch E for VLAN 10 and Switch A for VLAN 20
- F. Switch B for VLAN 10 and Switch E for VLAN 20

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should designate Switch A for VLAN 10 and Switch B for VLAN 20. The STP root bridge for a particular VLAN should be placed as close as possible to the center of the VLAN. If the majority of network traffic is between the hosts and servers within each VLAN, and the servers are grouped into a server farm, then the switch that all hosts will be sending their data to is the ideal choice for the STP root. Cisco's default

implementation of STP is called Per-VLAN Spanning Tree (or PVST), which allows individual tuning of the spanning tree within each VLAN. Switch A can be configured as the root bridge for VLAN 10, and Switch B can be configured as the root bridge for VLAN 20, resulting in optimized traffic flow for both.

None of the other switches is in the traffic flow of all data headed towards the VLAN 20 or VLAN 10 server farms, so they would not be good choices for the root bridge for either VLAN. Care should be taken when adding any switch to the network. The addition of an older, slower switch could cause inefficient data paths if the old switch should become the root bridge.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot STP protocols

References:

[Cisco > Support > Technology Support > LAN Switching > Spanning Tree Protocol > Configure > Configuration Examples and TechNotes > Understanding and Configuring Spanning Tree Protocol \(STP\) on Catalyst Switches](#)

QUESTION 53

Which switch port will be in a blocking state? (Click the Exhibit(s) button to view the switch port diagram.)



- A. SwitchA Fa0/1
- B. SwitchA Fa0/2
- C. SwitchB Fa0/1
- D. SwitchB Fa0/2

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

SwitchB will be forwarding on F0/1, and blocking on F0/2.

SwitchA will become the STP root bridge due to its lower MAC address. All ports on the root bridge will become designated ports in a forwarding state. SwitchB has redundant connectivity to the root bridge, and must block one of its interfaces to prevent a switching loop. STP will use its operations to determine which of the redundant interfaces on SwitchB to block to prevent a switching loop

Both interfaces are the same speed (FastEthernet), and thus their cost to the root is the same.

Finally, the interface with the lowest number will become the forwarding port. F0/1 has a lower port number than F0/2, so F0/1 becomes a forwarding port, and F0/2 becomes a blocking port.

Note: Unlike STP, Rapid Spanning Tree Protocol (RSTP) uses the term "discarding" for a switch port that is not forwarding frames.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Describe and verify switching concepts

References:

[Cisco > Support > LAN Switching > Spanning Tree Protocol > Technology White Paper > Understanding Rapid Spanning Tree Protocol \(802.1w\) > Document ID: 24062](#)

QUESTION 54

Which type of IP address is a registered IP address assigned by the Internet Service Provider (ISP), and represents one or more inside local IP addresses externally?

- A. Inside local address
- B. Outside local address
- C. Inside global address
- D. Outside global address

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

An inside global address is a registered IP address assigned by the ISP that represents internal local IP addresses externally.

An inside local address is an IP address (usually private) assigned to a host on the internal network. The inside local address is usually not assigned by the service provider, nor used to represent one or more inside local IP addresses externally

An outside local address is the IP address of an outside host as it appears to the internal network. It is not used to represent one or more inside local IP addresses externally

An outside global address is the IP address assigned to a host on the external network by the host owner. The address is allocated from a globally routable address space. It is not used to represent one or more inside local IP addresses externally

Objective:

Infrastructure Services

Sub-Objective:

Configure, verify, and troubleshoot inside source NAT

References:

[Cisco > Support > Technology Support > IP > IP Addressing Services > Design > Design TechNotes > NAT: Local and Global Definitions](#)

[Cisco > Articles > Network Technology > General Networking > Network Address Translation](#)

QUESTION 55

Which three statements are TRUE regarding Network Address Translation (NAT)? (Choose three.)

- A. It connects different Internet Service Providers (ISPs).
- B. It can act as an address translator between the Internet and a local network.
- C. It conserves IP addresses.

- D. It creates additional IP addresses for the local network.
- E. It helps the local network connect to the Internet using unregistered IP addresses.

Correct Answer: BCE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

NAT can act as an address translator between the Internet and the local network, conserve Internet Protocol (IP) addresses, and help the local network connect to the Internet using unregistered IP addresses.

The following statements are also TRUE regarding NAT:

- It can be used to present a single address for the entire network to the outside world when used in dynamic mode.
- It enhances network security by not disclosing the internal network addresses to the outside world.

It is not true that NAT connects different Internet Service Providers (ISPs). A gateway is used to connect different ISPs.

It is not true that NAT creates additional IP addresses for the local network. It only enables the use of unregistered addresses on the local area network.

Objective:

Infrastructure Services

Sub-Objective:

Configure, verify, and troubleshoot inside source NAT

References:

QUESTION 56

What is the default sequence in which a router searches for the Internetwork Operating System (IOS) image upon power on?

- A. TFTP, Flash, ROM
- B. ROM, Flash, TFTP
- C. Flash, TFTP, ROM
- D. Flash, TFTP, NVRAM
- E. NVRAM, Flash, TFTP

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The default sequence in which a router searches for the IOS image is in Flash memory, on a Trivial File Transfer Protocol (TFTP) server, and in read-only memory (ROM). The router will first search for the IOS image in the Flash memory. If there is no image in the Flash, the router will try to contact a TFTP server. If the router cannot find the IOS image on the TFTP server, it will load a limited version from the ROM.

The sequence that begins with TFTP and the sequence that begins with ROM are both incorrect sequences because the router will begin searching for the IOS image in Flash memory.

The sequences that include Non-volatile random access memory (NVRAM) are both incorrect because a router does not store the IOS image in NVRAM. The startup configuration is stored in NVRAM.

Objective:


Infrastructure Management
Sub-Objective:
Perform device maintenance

References:

[Cisco > Cisco IOS Configuration Fundamentals Configuration Guide, Release 12.4 > Part 9: Loading and Maintaining System Images > Specifying the Startup System Image in the Configuration File](#)

QUESTION 57

R1 and R2 are connected as shown in the diagram and are configured as shown in output in the partial output of the show run command.



```
R1#show run
```

```
version 12.0
hostname R1

interface s1
ip address 192.168.5.5 255.255.255.252

ip host R1 192.168.5.6
```

```
R2#show run
```

```
version 12.0
hostname R2
interface s1
ip address 192.168.5.6 255.255.255.252
ip host R1 192.168.5.5
```

The command ping R2 fails when executed from R1. What command(s) would allow R1 to ping R2 by name?

- A. R1(config)#int S1
R1(config-if)#no ip address 192.168.5.5
R1(config-if)# ip address 192.168.5.9 255.255.255.252
- B. R1(config)#no ip host R1
R1(config)# ip host R2 192.168.5.6 255.255.255.252
- C. R1(config)#no hostname R2
R1(config)# hostname R1
- D. R2(config)#int S1
R1(config-if)#no ip address 192.168.5.5
R1(config-if)# ip address 192.168.5.9 255.255.255.0

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Both routers have been configured with the ip host command. This command creates a name to IP address mapping, thereby enabling the pinging of the device by address. On R1, the mapping is incorrect and needs to be corrected. Currently it is configured as ip host R1 192.168.5.6. It is currently mapping its own name to the IP address of R2.

To fix the problem, you should remove the incorrect IP address mapping and create the correct mapping for R2, as follows:

```
R1(config)#no ip host R1
R1(config)# ip host R2 192.168.5.6 255.255.255.252
```

Once this is done, the ping on R2 will succeed.

The IP address of the S1 interface on R1 does not need to be changed to 192.168.5.9 /30. In fact, if that is done the S1 interface on R1 and the S1 interface in R2 will no longer be in the same network. With a 30-bit mask configured, the network they are currently in extends from 192.168.5.4 - 192.168.5.7. They are currently set to the two usable addresses in that network, 192.168.5.5 and 192.168.5.6.

The hostnames of the two routers do need to be set correctly using the hostname command for the ping to function, but they are correct now and do not need to be changed.

The subnet mask of the S1 interface on R2 does not need to be changed to 255.255.255.0. The mask needs to match that of R1, which is 255.255.255.252.

Objective:

Infrastructure Services

Sub-Objective:

Troubleshoot client connectivity issues involving DNS

References:

QUESTION 58

You are attempting to add an IP address to an interface on a router with which you are unfamiliar. You type the following command and receive the following error:

```
Router78(config)#interfce Serial0
                        ^
%invalid input detected at '^' marker.
```

Which of the following could be a reason for receiving this message?

- A. the command syntax is incorrect
- B. the interface type does not exist on this router
- C. the command is entered at the wrong prompt
- D. the interface is configured already

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The command has a syntax error. The word interface is misspelled as indicated by the marker.

The interface type may not exist on the router, but that is not the problem with this specific error message. If you attempt to access an interface that is not present on the router, it will elicit this same message, but the marker will be placed at the beginning of the interface type as shown below. The interface information is in lines 14-19.

```
Router78(config)#interface Serial0
```

```
%invalid input detected at '^' marker.
```

When you are unfamiliar with a router, it is best to execute the show version command, which will indicate the type and number of interfaces on the router as shown below:

```
Router78# show version
Cisco IOS Software, 3800 Software (C3845-IPBASE-M), Version 12.3(11)T7, RELEASE SOFTWARE (fc3)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2005 by Cisco Systems, Inc.
Compiled Sat 30-Jul-05 03:12 by dchih

ROM: System Bootstrap, Version 12.3(11r)T2, RELEASE SOFTWARE (fc1)

Router78 uptime is 10 weeks, 6 days, 9 hours, 30 minutes
System returned to ROM by power-on
System restarted at 05:09:07 CDT Thu Oct 20 2005
System image file is "flash:c3845-ipbase-mz.123-11.T7.bin"

Cisco 3845 (revision 1.0) with 419839K/104448K bytes of memory.
Processor board ID FTX0938A5PE
2 Gigabit Ethernet interfaces
27 Serial interfaces
1 ISDN Basic Rate interface
6 terminal lines
2 Channelized T1/PRI ports
1 Subrate T3/E3 port
DRAM configuration is 64 bits wide with parity enabled.
479K bytes of NVRAM.
62720K bytes of ATA System CompactFlash (Read/Write)

Configuration register is 0x2102
```

The command is not entered at the wrong prompt. It should be entered at the global configuration prompt.

If the interface were already configured, it would still allow you to access the interface and make changes.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Troubleshoot interface and cable issues (collisions, errors, duplex, speed)

References:

QUESTION 59

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.

Objective:

LAN Switching Fundamentals


Sub-Objective:

Configure, verify, and troubleshoot STP protocols

References:

QUESTION 60

Examine the diagram below and assume that routing is configured properly.



Why is Host A unable to ping Host B?

- A. The IP address of Switch A is incorrect
- B. The gateway address of Host B is incorrect
- C. The IP address of Host A is incorrect
- D. The Fa0/2 and Fa0/1 interfaces on R1 and R2 are not in the same subnet

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The IP address of Host A is incorrect. The Fa0/1 interface on R1 (Host A's default gateway) is in the 192.168.1.64/27 network, and Host A's IP address is in the 192.168.1.32/27 network. With a 27-bit mask against the 192.168.1.0 classful network, the resulting subnets are:

192.168.1.0
192.168.1.32
192.168.1.64
192.168.1.96

And so it would continue, increasing the fourth octet in intervals of 32. By only going this far we can see that they are in different subnets.

The IP address of Switch A is correct for its subnet because it needs to be in the same subnet as the Fa0/1 interface on R1. Even if it were incorrect or missing altogether it would have no impact on Host A. Switches

merely switch frames based on MAC addresses and only need an IP address for management purposes.

The gateway address of Host B is correct. It is in the same subnet (15.0.0.0/8) with the Fa0/2 interface on R2, its gateway.

The Fa0/2 and Fa0/1 interfaces on R1 and R2 are in the same subnet. Using a 25-bit mask against the 192.18.5.0/24 classful network yields the following subnets:

192.18.5.0
192.168.5.128

Both router interfaces in question are in the 192.18.5.0 subnet.

Objective:

Network Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot IPv4 addressing and subnetting

References:

[Cisco > Support > IP Routing > Design TechNotes > Document ID: 13788 > IP Addressing and Subnetting for New Users](#)

QUESTION 61

You are configuring a serial link between a Cisco router and a router produced by another vendor.

What would be the advantages of using Point to Point Protocol (PPP) over High Level Data Link Control (HDLC) in this scenario?

- A. HDLC has a proprietary "type" field that may be incompatible with equipment from other vendors.
- B. HDLC is not available on non-Cisco routers.
- C. PPP is faster.
- D. PPP performs error checking.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

High Level Data Link Control (HDLC) has a proprietary "type" field that may be incompatible with equipment from other vendors. It is recommended that PPP always be used when combining equipment from multiple vendors because this Data Link layer WAN protocol is an industry standard. PPP is implemented in the same manner on all PPP-capable equipment.

HDLC is available on non-Cisco routers. However, the Cisco implementation has a "type" field that may prevent the connection from working.

PPP is not faster than HDLC.

PPP performs error checking, but so does HDLC.

Objective:

WAN Technologies

Sub-Objective:

Describe WAN access connectivity options

References:

[Cisco > Internetworking Technology Handbook > Point to Point Protocol \(PPP\)](#)

QUESTION 62

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)#.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot single area and multi-area OSPFv2 for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub, virtual-link, and LSAs)

References:

[Cisco > Home > Support > Technology Support > IP > IP Routing > Technology Information > Technology White paper > OSPF Design Guide](#)

QUESTION 63

Which command would be used to establish static translation between an inside local address 192.168.144.25 and an inside global address 202.56.63.102?

- A. router(config)#ip nat inside source static 192.168.144.25 202.56.63.102
- B. router(config)#ip source nat inside static local-ip 192.168.144.25 global-ip 202.56.63.102
- C. router(config)#ip nat static inside source 192.168.144.25 202.56.63.102
- D. router(config)#ip nat inside static source 192.168.144.25 202.56.63.102

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

To establish a static translation between an inside local address 192.168.144.25 and an inside global address 202.56.63.102, you would use the ip nat inside source static 192.168.144.25 202.56.63.102 command executed in global configuration mode. The correct format of the command is:

ip nat inside source static local-ip global-ip

This static configuration can be removed by entering the global no ip nat inside source static command.

Simply executing the ip nat inside source command will not result in NAT functioning. The NAT process also has to be applied correctly to the inside and outside interfaces. For example, if, in this scenario the Fa0/0 interface hosted the LAN and the S0/0 interface connected to the Internet the following commands would complete the configuration of static NAT.

```
Router(config)#interface F0/0
Router(config-if)#ip nat inside
Router(config-if)#exit
Router(config)#interface S0/0
Router(config-if)#ip nat outside
```

The other options are incorrect because they are not valid Cisco IOS configuration commands. They all contain syntax errors.

Objective:

Infrastructure Services

Sub-Objective:

Configure, verify, and troubleshoot inside source NAT

References:

QUESTION 64

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.

Objective:

WAN Technologies

Sub-Objective:

Describe WAN access connectivity options

References:

QUESTION 65

Refer to the partial output of the show interfaces command:

```
Serial 0 is administratively down, 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 runts, 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 administratively down, line protocol is down line indicate with certainty?

- A. There is no problem with the physical connectivity.
- B. There is a configuration problem in the local or remote router.
- C. There is a problem at the telephone company's end.
- D. The shutdown interface command is present in the router configuration.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Serial 0 is administratively down, line protocol is down line in the output of the show interfaces command indicates the following:

- The shutdown interface command is present in the router configuration. This indicates that the administrator might have manually shut down the interface by issuing the shutdown command.
- A duplicate Internet Protocol (IP) address might be in use.

This line does not show that there is no problem with the physical connectivity. Since the interface is administratively shut down, there is no way of determining the operational status of the physical layer.

The Serial 0 is administratively down, line protocol is down line does not indicate a configuration problem in the local or remote router. A problem in the configuration of local or remote router would be indicated by the Serial 0 is up, line protocol is down message.

This line does not show that there is a problem at the telephone company's end. Since the interface is administratively shut down, there is no way of determining the operational status of the physical layer or protocol layer on the other end of the line.

Objective:

Infrastructure Management


Sub-Objective:

Use Cisco IOS tools to troubleshoot and resolve problems

References:

QUESTION 66

The following exhibit displays the MAC address table of a switch in your network, along with the location of each device connected to the switch:



Which of the following frames will be flooded to all ports after it is received by the switch?

- A. source MAC: 12-34-56-78-9A-BD, destination MAC: 12-34-56-78-9A-BF
- B. source MAC: 12-34-56-78-9A-BF, destination MAC: 12-34-56-78-9A-BD
- C. source MAC: 12-34-56-78-9A-BF, destination MAC: 12-34-56-78-9A-BC
- D. source MAC: 12-34-56-78-9A-BC, destination MAC: 12-34-56-78-9A-BF

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The frame with a source MAC of 12-34-56-78-9A-BF and a destination MAC of 12-34-56-78-9A-BC would be sent to all ports because the destination MAC address is not already in the MAC address table.

The frame with a source MAC of 12-34-56-78-9A-BD and a destination MAC of 12-34-56-78-9A-BF would not be sent to all ports because the destination MAC address is in the MAC address table.

The frame with a source MAC of 12-34-56-78-9A-BF and a destination MAC of 12-34-56-78-9A-BD would not be sent to all ports because the destination MAC address is in the MAC address table.

The frame with a source MAC of 12-34-56-78-9A-BC and a destination MAC of 12-34-56-78-9A-BF would not be sent to all ports because the destination MAC address is in the MAC address table.

Objective:
LAN Switching Fundamentals
Sub-Objective:
Interpret Ethernet frame format

References:
[Cisco Press > Articles > Cisco Certification > CCNA Routing and Switching > Basic Data Transmission in Networks: MAC Tables and ARP Tables](#)
[How do Switches Work?](#)

QUESTION 67

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.

Objective:
Routing Fundamentals
Sub-Objective:
Configure, verify, and troubleshoot single area and multi-area OSPFv2 for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub, virtual-link, and LSAs)

References:

QUESTION 68

Which layer in the Open Systems Interconnection (OSI) model defines an Internet Protocol (IP) address that helps in selecting the route to the destination?

- A. Data Link
- B. Network
- C. Application
- D. Transport

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Network layer in the OSI model defines a logical address that helps select the route to the destination. Logical addresses, such as IP addresses or IPX-SPX addresses, are used by routers to forward the packets to the destination. This is accomplished systematically by comparing the destination address with the network addresses listed in the routing table. The logical layout of the network is also defined at this layer. The Network layer is primarily concerned with logical addressing, routing, and path determination.

Protocol data units (PDUs) are called packets at the Network layer. The information that is applied at this layer, which consists of IP addresses, is used in the routing process.

The Data Link layer does not define an IP address. This layer ensures the reliable transmission of data across a network and defines the Media Access Control (MAC) address, which defines the physical device addressing. This layer also defines the format of the header and trailer.

Protocol data units (PDUs) are called frames at the Data Link layer. The information that is applied at this layer, which consists of MAC addresses, is used in the switching process.

The application layer does not define an IP address. The application layer is responsible for interacting directly with the application and provides application services, such as e-mail, FTP, and Telnet. It also defines the user authentication process.

The Transport layer does not define an IP address. The Transport layer is responsible for the error-free and sequential delivery of data. This layer is used to manage data transmission between devices, a process known as flow control.

Protocol data units (PDUs) are called segments at the Transport layer, where the two protocols TCP and UDP operate. Windowing, which is the real-time management of the number of packets that can be received without an acknowledgement, is handled by TCP at this layer.

Objective:

Network Fundamentals

Sub-Objective:


Compare and contrast OSI and TCP/IP models

References:

[Cisco Documentation > Internetworking Technology Handbook > Internetworking Basics>Internet Protocols](#)

QUESTION 69

Examine the network diagram.



Which switch port(s) will be in a forwarding state? (Choose two.)

- A. SwitchA - Fa0/1 and Fa0/2
- B. SwitchA - Fa0/1
- C. SwitchA - Fa0/2
- D. SwitchB - Fa0/1
- E. SwitchB - Fa0/2

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Both switch ports on Switch A and Fa0/1 on Switch B will be in a forwarding state. Switch A will become the STP root bridge due to its lower MAC address. All ports on the root bridge will become designated ports in a forwarding state. Switch B has redundant connectivity to the root bridge, and must block one of its interfaces to prevent a switching loop. Both interfaces are the same speed (FastEthernet), and thus their cost to the root is the same. Finally, the interface with the lowest number will become the forwarding port. F0/1 has a lower port number than F0/2, so F0/1 becomes a forwarding port, and F0/2 becomes a blocking port.

In this scenario there are only two switches in the diagram. However, if there were more switches and Switch A were not the root bridge, the result would be the same with regard to the ports between Switch A and B. Whenever there are redundant links between switches, one of the four ports involved will be set to a blocking (or in the case of RSTP, discarding) mode. The logic will still be the same, since the cost to get to the root bridge will still be equal if the port speeds are equal.

Without STP (which can be disabled) operating on switches with redundant links, such as those in the figure, loops can and almost surely will occur. For example, if a host connected to SwitchA were to send an ARP request for the MAC address of a host connected to SwitchB, the request could loop and cause a broadcast storm, slowing performance dramatically. This would probably occur when any host connected to either switch sends a broadcast frame, such as a DHCP request.

Rapid Spanning Tree Protocol (RSTP) uses the term discarding for a switch port that is not forwarding frames.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot STP protocols


References:

[Cisco > Support > Technology Support > LAN Switching > Spanning Tree Protocol > Configure > Configuration Examples and TechNotes > Understanding and Configuring Spanning Tree Protocol \(STP\) on Catalyst](#)

[Switches](#)

QUESTION 70

The diagram below shows the state of the switch interfaces after STP has converged.



Based on the interface states, which of the following statements are true? (Choose all that apply.)

- A. The Fa0/2 interface on SW 2 is a designated port
- B. SW 3 is the root bridge
- C. SW 2 is the root bridge
- D. The Fa0/0 interface on SW 2 is a designated port
- E. The Fa0/0 interface on SW 2 is a root port

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Convergence has occurred in a spanning-tree network when all switch ports are in either a forwarding state or a blocking state (known as discarding state in RSTP). You can use the location of these blocked and forwarding ports to infer the location of the root bridge and the state of any unlabeled ports in the diagram.

SW3 is the root bridge and the Fa0/0 interface on SW2 is a designated port. It can be determined that SW3 is the root bridge because all of its ports are in a forwarding state. Any switch that has at least one port blocking (such as SW1 and SW2) are non-root bridges. As there must be a root bridge, that leaves SW3 as the only candidate.

After establishing that SW3 is the root, it can be determined that the connection between SW1 and SW2 is a segment that does not have a direct connection to the root bridge. These segments must have one end set as a designated port and thus set to forward. Since the Fa0/0 interface on SW2 is forwarding, it is the designated port for that segment.

The Fa0/2 interface on SW2 is not a designated port. The interface on each non-root switch with the lowest cost path to the root bridge will be the root port. Since SW3 is the root bridge, the connection to SW3 via Fa0/2 is the lowest cost path to the root bridge for SW1 and thus is a root port, not a designated port. Moreover, designated ports only exist on segments that do not have a direct connection to the root bridge.

SW2 is not the root bridge. One of its ports is blocking, which will not occur on a root bridge.

The Fa0/0 interface on SW2 is a not root port. It is the designated port for the segment between SW1 and SW2.

The process of determining these port states occurs in this order:

1. Selection of the root bridge. When all bridge priorities have been left to their default, all switches will have same bridge priority. When that is the case, the switch with the lowest MAC address will be selected root bridge. ALL ports are in a forwarding state on the root bridge, which explains why all of the ports on SW3 will be in a forwarding state.
2. Determination of the root ports on each non-root bridge. Each non-root bridge will select the interface it possesses with the least cost path to the root bridge. Once selected, that port will be placed in a forwarding state.
3. Determination of the designated port on each segment that does not connect directly to the root bridge. There is one such segment in the diagram (SW1 to SW2). The interface on either end of the segment that has the least cost path to the root bridge will be the designated port for that section. It may have several paths, but the least cost path is used in the determination of the designated port for the segment.

Once determined, the designated ports will be set to forwarding, and all ports that are neither root nor designated ports will be set to blocking.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot STP protocols

References:

[Cisco > Internetworking Technology Handbook > Bridging and Switching > Transparent Bridging > Spanning-Tree Algorithm](#)

QUESTION 71

Which command would be used to establish static translation between an inside local address 192.168.144.25 and an inside global address 202.56.63.102?

- A. router(config)#ip nat inside source static 192.168.144.25 202.56.63.102
- B. router(config)#ip source nat inside static local-ip 192.168.144.25 global-ip 202.56.63.102
- C. router(config)#ip nat static inside source 192.168.144.25 202.56.63.102
- D. router(config)#ip nat inside static source 192.168.144.25 202.56.63.102

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

To establish a static translation between an inside local address 192.168.144.25 and an inside global address 202.56.63.102, you would use the ip nat inside source static 192.168.144.25 202.56.63.102 command executed in global configuration mode. The correct format of the command is:

ip nat inside source static local-ip global-ip

This static configuration can be removed by entering the global no ip nat inside source static command.

Simply executing the ip nat inside source command will not result in NAT functioning. The NAT process also has to be applied correctly to the inside and outside interfaces. For example if, in this scenario the Fa0/0 interface hosted the LAN and the S0/0 interface connected to the Internet the following commands would complete the configuration of static NAT.

```
Router(config)#interface F0/0
Router(config-if)#ip nat inside
Router(config-if)#exit
Router(config)#interface S0/0
Router(config-if)#ip nat outside
```

The other options are incorrect because they are not valid Cisco IOS configuration commands. They all contain syntax errors.

Objective:

Infrastructure Services

Sub-Objective:

Configure, verify, and troubleshoot inside source NAT

References:

QUESTION 72

How many IP addresses can be assigned to hosts in subnet 192.168.12.64/26?

- A. 32
- B. 62
- C. 128
- D. 256

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Subnet 192.168.12.64/26 has 62 IP addresses that can be assigned to hosts.

The formula to calculate the available number of hosts is:

$$2^n - 2 = x$$

Where n = the number of host bits in the subnet mask and x = the number of possible hosts.

You will subtract 2 from the hosts calculation to remove the first address (the network ID) and the last address (the broadcast ID) from the valid hosts range. These addresses are reserved as the network ID and the broadcast address, respectively, in each subnet.

An IP address has 32 available bits divided into four octets. In this scenario, the /26 indicates that the subnet mask is 26 bits long, or that 26 bits are reserved for the network portion of the address. This leaves 6 bits for the host addresses ($32 - 26 = 6$). The number of host addresses would be calculated as follows:

$$\text{Number of hosts} = 2^6 - 2$$

$$\text{Number of hosts} = 64 - 2 = 62$$

Another simple way of determining the number of hosts in a range, when the subnet mask extends into the last octet, is to determine the decimal value of the last bit in the subnet mask after converting it to binary notation. This process only works when the subnet extends into the last octet, meaning that the subnet is greater than /24. The /26 subnet mask equals 26 network bits and 6 hosts bits, written as follows:

11111111.11111111.11111111.11000000

The 1s represent network bits and the 0s represent host bits.

In this example, the 26th bit (read from left to right) has a decimal value of 64, indicating that this subnet has 64

addresses. Subtract 2 to represent the network and broadcast addresses ($64 - 2 = 62$). This shows that this subnet range can be used to address 62 hosts.

Network address: 192.168.12.0
Subnet Mask in decimal: 255.255.255.192
Subnet Mask in binary: 11111111.11111111.11111111.11000000
Hosts: $64 - 2 = 62$

For subnet 192.168.12.64, the valid host range will start from 192.168.12.65 to 192.168.12.126. For the next subnet 192.168.12.128, the valid host range will start from 192.168.12.129 to 192.168.12.190.

To construct a subnet that would contain 32 addresses would require using a mask of 255.255.255.224. This mask would leave 5 host bits, and $2^5 - 2 = 32$.

To construct a subnet that would contain 128 addresses would require using a mask of 255.255.255.128. This mask would leave 7 host bits, and $2^7 - 2 = 128$.

To construct a subnet that would contain 256 addresses would require using a mask of 255.255.255.0. This mask would leave 8 host bits, and $2^8 - 2 = 256$.

Objective:
Network Fundamentals
Sub-Objective:
Configure, verify, and troubleshoot IPv4 addressing and subnetting

References:
[Cisco > Design Tech Notes > IP Routing > IP Addressing and Subnetting for New Users > Understanding IP Addresses > Document ID: 13788](#)
[Nooning, Thomas. "TechRepublic Tutorial: Subnetting a TCP/IP Network." TechRepublic, 20 May 2003.](#)

QUESTION 73

Which option lists the given applications in the correct sequence of increasing bandwidth consumption?

- A. an interactive Telnet session on a server running an SAP application
a voice conversation between PC-based VoIP services
a voice conversation between two IP phones while accessing an online video site
- B. a voice conversation between two IP phones while accessing an online video site
an interactive Telnet session on a server running an SAP application
a voice conversation between PC-based VoIP services
- C. a voice conversation between PC-based VoIP services
a voice conversation between two IP phones while accessing an online video site
an interactive Telnet session on a server running an SAP application
- D. an interactive Telnet session on a server running an SAP application
a voice conversation between two IP phones while accessing an online video site
a voice conversation between PC-based VoIP services

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The correct sequence of increasing bandwidth consumption in the given scenario would be, from lowest to highest:

1. an interactive Telnet session on a server running an SAP application
2. a voice conversation between PC-based VoIP services
3. a voice conversation between two IP phones while accessing an online video site

An interactive Telnet session uses the least amount of bandwidth of the three application examples because it

mainly involves the transfer of text.

A voice conversation between IP phones, also known as voice over IP (VoIP) traffic, requires more bandwidth than Telnet. Voice traffic is delay-sensitive and benefits from Quality of Service (QoS) to ensure service quality.

A voice conversation between two IP phones while accessing an online video site would consume the most bandwidth. A voice conversation with real-time video exchange is the equivalent of real-time video traffic. Video traffic is real-time and benefits from dedicated bandwidth with QoS implementation to ensure quality.

Objective:

WAN Technologies

Sub-Objective:

Describe basic QoS concepts

References:

[Cisco Documentation > Internetworking Technology Handbook > Voice/Data Integration Technologies](#)

QUESTION 74

Consider the following output of the show ip interface brief command:

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

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

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

- A. R1(config-router)#network 192.168.12.64 0.0.0.127 area 0
- B. R1(config-router)#network 192.168.12.64 0.0.0.63 area 0
- C. R1(config-router)#network 192.168.12.64 0.0.0.66 area 0
- D. R1(config-router)#network 192.168.12.64 255.255.255.192 area 0
- E. R1(config-router)#network 192.168.12.64 0.0.0.63 area1

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

Wildcard masks in OSPF network statements are expressed inversely, and not as a regular subnet masks. For example, if the network you are configuring for OSPF operation is 192.168.10.0/24, then the inverse version of a /24 mask (or 255.255.255.0) would be 0.0.0.255.

The network ID is the starting point and the wildcard mask specifies where the network will end or the range of

the network. In this case, the network begins at 192.168.12.64. The value in the last octet of the mask indicates the number of values (including 64) that will be included in the network, which means that it will range from 192.168.12.64 - 192.168.12.127. 64 to 127 equals 64 values if you include the endpoints 64 and 127.

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

The command R1(config-router)#network 192.168.12.64 0.0.0.127 area 0 is incorrect because the resulting network would range from 192.168.12.64 - 192.168.12.191. This would include all of the required interfaces, but would also include E1 (192.168.12.129) and S0 (192.18.12.187), which is not desired.

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

The command R1(config-router)#network 192.168.12.64 255.255.255.192 area 0 is incorrect because the mask, while correct in its breadth and the exact inverse of the wild card mask 0.0.0.63, is not stated in wildcard mask format.

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

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot single area and multi-area OSPFv2 for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub, virtual-link, and LSAs)

References:

[Cisco > Home > Support > Technology Support > IP > IP Routing > Technology Information > Technology White paper > OSPF Design Guide > Enabling OSPF on the Router](#)

QUESTION 75

Which Cisco Internetwork Operating System (IOS) command is used to save the running configuration to non-volatile random access memory (NVRAM)?

- A. copy startup-config running-config
- B. move startup-config running-config
- C. copy running-config startup-config
- D. move startup-config running-config

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The copy running-config startup-config command is used to save the running configuration to NVRAM. This command will should always been run after making changes to the configuration. Failure to do so will result in the changes being discarded at the next restart of the router. When the router is restarted, the startup configuration file is copied to RAM and becomes the running configuration.

The copy startup-config running-config command is incorrect because this command is used to copy the startup configuration to the running configuration. The command would be used to discard changes to the configuration without restarting the router.

The move startup-config running-config and move startup-config running-config commands are incorrect

because these are not valid Cisco IOS commands. There is no move command when discussing the manipulation of configuration files.

Objective:

Infrastructure Management

Sub-Objective:

Perform device maintenance

References:

[Cisco Documentation > RPM Installation and Configuration > IOS and Configuration Basics](#)

QUESTION 76

DRAG DROP

Click and drag the features on the left to their corresponding frame tagging method on the right.

Select and Place:

Features	ISL	IEEE 802.1Q
Cisco standard		
Industry standard		
Adds a 4-byte tag in the middle of original Ethernet frame		
Adds a 26-byte header and 4-byte trailer		
Does not modify Ethernet frame		
Native VLAN frames are not tagged while traversing over trunk links		

Correct Answer:

Features

ISL	IEEE 802.1Q
Cisco standard	Industry standard
Adds a 26-byte header and 4-byte trailer	Adds a 4-byte tag in the middle of original Ethernet frame
Does not modify Ethernet frame	Native VLAN frames are not tagged while traversing over trunk links

Section: (none)

Explanation

Explanation/Reference:

Explanation:

ISL and IEEE 802.1Q are VLAN frame tagging methods.

ISL:

- Is Cisco proprietary
- Adds a 26-byte header and 4-byte trailer
- Does not modify Ethernet frame

IEEE 802.1Q frame tagging method:

- Is a standard method
- Adds a 4-byte tag in the middle of original Ethernet frame
- Has a concept called native VLAN. Native VLAN frames are not tagged while traversing over a trunk link.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot interswitch connectivity

References:

[Cisco > Technology Support > LAN Switching > Virtual LANs/VLAN Trunking Protocol \(VLANs/VTP\) > Design > Design Technotes > Inter-Switch Link and IEEE 802.1Q Frame Format > Document ID: 17056](#)

QUESTION 77

Which of the following is the correct command to define a default route using a gateway address of 172.16.0.254?

- A. ip default-route 172.16.0.254 255.255.0.0
- B. ip route 0.0.0.0 0.0.0.0 172.16.0.254
- C. default-gateway 172.16.0.254

D. ip route default 172.16.0.254

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The ip route command is used to manually define a static route to a destination network. The syntax of the command is as follows:

ip route [destination_network] [mask] [next-hop_address or exit interface] [administrative_distance] [permanent]

The attributes of the command are as follows:

- destination_network: Defines the network that needs to be added in the routing table.
- mask: Defines the subnet mask used on the network.
- next-hop_address: Defines the default gateway or next-hop router that receives and forwards the packets to the remote network.
- administrative_distance (AD): States the administrative distance. Static routes have an AD of 1, which can be changed to change the priority of the route.

Creating a default route is accomplished by substituting 0.0.0.0 for both the [destination_network] and [mask] fields, yielding the following command to create a default route through host 172.16.0.254:

router(config)# ip route 0.0.0.0 0.0.0.0 172.16.0.254

Any route configured manually is considered a static route. Another example of a command that creates a non-default route is shown below:

router(config)# ip route 192.168.12.0 255.255.255.0 172.65.3.1

This command would instruct the router on which the command was executed to send any traffic for the 192.168.12.0/24 network to the router located at 172.65.3.1.

You can also affect the route by changing the administrative distance of the route. By default, all static routes have an AD of 1, making them preferable to routes learned from routing protocols. However, you can add the AD parameter at the end of the command as shown below, making the static route less desirable than one learned from a routing protocol such as RIP:

router(config)# ip route 192.168.12.0 255.255.255.0 172.65.3.1 150

One reason to configure the routes this way could be to make the static route a backup route to the route learned by RIP, such as when the static route is a less desirable route through a distant office.

Once the ip route command has been used to add either a static route or a static default route to a router, the routes should appear in the routing table. They will be indicated with an S next to a static route and an S* for a default static route. The first two examples from the explanation above would appear in the routing table as follows:

```
S*0.0.0.0/0 [1/0] via 172.16.0.254
S 192.168.12.0/24 [1/0] via 172.65.3.1
```

The ip default-route, default-gateway, and ip route default commands are incorrect because they are not valid Cisco IOS commands.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot IPv4 and IPv6 static routing

References:

[Cisco > Cisco ASDM User Guide, 6.1 > Configuring Dynamic And Static Routing > Field Information for Static Routes](#)

[Cisco > Support > IP > IP Routing > Design > Design TechNotes > Specifying a Next Hop IP Address for Static Routes > Document ID: 27082](#)

QUESTION 78

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.

Objective:

Infrastructure Management

Sub-Objective:

Describe network programmability in enterprise network architecture

References:

[Software Defined Networking: The Cisco approach](#)

QUESTION 79

Which show interfaces command output indicates that the link may not be functional due to a Data Link layer issue, while the Physical layer is operational?

- A. Ethernet 0/0 is up, line protocol is up
- B. Ethernet 0/0 is up, line protocol is down
- C. Ethernet 0/0 is down, line protocol is up
- D. Ethernet 0/0 is down, line protocol is down

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The first or left-hand column (Ethernet 0/0 is up) indicates the Physical layer state of the interface, while the second or right-hand column (line protocol is down) indicates the Data Link layer state of the interface. The following command output excerpt indicates that the link is not functional due to a Data Link layer (or "line

protocol") issue, while the Physical layer is operational:

Ethernet 0/0 is up, line protocol is down

If the problem were at the Data Link layer while the Physical layer is operational, the show interfaces command output will indicate that the interface is up, but the line protocol is down.

In the normal operation mode, when both Physical layer and Data Link layer are up, the show interfaces output will display the following message:

Ethernet0/0 is up, line protocol is up

The message Ethernet 0/0 is down, line protocol is up is not a valid output.

The message Ethernet 0/0 is down, line protocol is down indicates that both the Physical layer and the Data Link layer are down. Therefore, this is an incorrect option.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Troubleshoot interface and cable issues (collisions, errors, duplex, speed)

References:

QUESTION 80

Which of the following topologies is used in Wide Area Networks (WANs)?

- A. FDDI
- B. CDDI
- C. SONET
- D. Token Ring

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Synchronous Optical NETwork (SONET) is the standard topology for fiber optic networks. Developed in 1980s, SONET can transmit data at rates of up to 2.5 gigabits per second (Gbps).

All other options are incorrect because they are LAN topologies, not WAN topologies.

Fiber Distributed Data Interface (FDDI) specifies a 100-Mbps dual-ring fiber optics-based token-passing LAN. FDDI is typically implemented for high-speed LAN backbones because of its support for high bandwidth.

Copper Distributed Data Interface (CDDI) is copper version of FDDI. They differ only in that FDDI can span longer distances than CDDI due to the attenuation characteristics of copper wiring.

Token Ring/IEEE 802.5 LAN technology was developed by IBM in 1970. Token-ring LAN technology is based on token-passing, in which a small frame, called a token, is passed around the network. Possession of the token grants the node the right to transmit data. Once the data is transmitted, the station passes the token to the next end station.

Objective:

Network Fundamentals

Sub-Objective:

Compare and contrast network topologies

References:

[Cisco>Home>Cisco Documentation > Internetworking Technology Handbook>WAN Technologies](#)

QUESTION 81

You are the network administrator for your company. You have implemented VLAN Trunking Protocol (VTP) in your network. However, you have found that VTP is not synchronizing VLAN information.

Which of the following items should be verified to resolve the problem? (Choose three.)

- A. Ensure that switches in the VTP domain are configured with VTP version 1 and version 2.
- B. Ensure that VLANs are active on at least one switch on the VTP domain.
- C. Ensure that all of the ports that interconnect switches are configured as trunks and are trunking properly.
- D. Ensure that the VTP domain name is the same on all switches in the domain.
- E. Ensure that identical passwords are configured on all VTP switches.

Correct Answer: CDE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The following is a list of the steps to take if VTP fails to exchange VLAN information:

- Ensure that all of the ports that interconnect switches are configured as trunks and are trunking properly.
- Ensure that VLANs are active in all the devices.
- Ensure that at least one switch is acting as a VTP server in the VTP domain.
- Ensure that the VTP domain name is the same for all switches in the domain. The VTP domain name is case-sensitive.
- Ensure that the VTP password is the same for all switches in the domain.
- Ensure that the same VTP version is used by every switch in the domain. VTP version 1 and version 2 are not compatible on switches in the same VTP domain.

You should not ensure that switches are configured with VTP version 1 and version 2 in the domain, because VTP version 1 and version 2 are incompatible. VTP version 1 is the default on all Cisco switches.

You should not ensure that VLANs are active on at least one switch in the VTP domain, because VLANs should be active in all of the devices in a VTP domain.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot VLANs (normal/extended range) spanning multiple switches

References:

QUESTION 82

Which of the following is NOT a possible component of Enhanced Interior Gateway Routing Protocol's (EIGRP) composite metric?

- A. Cost
- B. Load
- C. Delay
- D. Bandwidth

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Cost is not a component of EIGRP's composite metric. The cost, or efficiency, of a path is used as a metric by the Open Shortest Path First (OSPF) routing protocol.

Enhanced IGRP (EIGRP) is Cisco Systems' proprietary routing protocol. It can use bandwidth, delay, load, reliability, and maximum transmission unit (MTU) to calculate the metric. Of these five metrics, by default, only minimum bandwidth and delay are used to compute the best path.

The metric for EIGRP can be calculated with this formula:

$$\text{Metric} = [K1 * \text{Bandwidth} + (K2 * \text{Bandwidth}) / (256 - \text{load}) + K3 * \text{Delay}] * [K5 / (\text{reliability} + K4)]$$

The default constant values for Cisco routers are K1 = 1, K3 = 1, and K2 = 0, K4 = 0, K5 = 0. In the default setting, K1 and K3 have non-zero values, and therefore, by default, the metric is dependent on bandwidth and delay.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot EIGRP for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub)

References:

[Cisco > Support > IP > IP Routing > Design > Design TechNotes > Introduction to EIGRP > Document ID: 13669](#)

QUESTION 83

Which are among the valid steps in the process of recovering a password on a Cisco router? (Choose all that apply.)

- A. Restart the router.
- B. Configure the enable secret password.
- C. Enter the router diagnostic mode.
- D. Enter user mode.
- E. Answer the security question to recover the password.

Correct Answer: ABC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Three of the steps that should be performed while recovering a password on a Cisco router are to restart the router in ROMMON mode, enter ROMMON mode (router diagnostic mode) and reset the enable secret password. The complete password recovery process on a Cisco Router is as follows:

Configure the router so that it starts without reading the non-volatile random access memory (NVRAM). This is also referred to as the system test mode, which you enter by changing the configuration register. You must first restart the router and within 60 seconds press Break on the terminal keyboard. Then the router will skip normal reading of the startup configuration file and will go to the ROMMON prompt (shown below this text section). At this command prompt, type confreg 0x2142 to instruct the router to boot to flash memory at the next reboot. When it does, it will ignore the startup configuration file again and will behave as if it had no configuration, as a new router would.

rommon 1> confreg 0x2142

Type reset to reboot the router.

Enter enable mode through the test system mode.

View the existing password (if it can be viewed, it may be encrypted), configure a new password, or delete the configuration.
Configure the router to start by reading the NVRAM, which is done by resetting the configuration register to its normal value. Run these commands:

Router#configure terminal

Enter configuration commands, one per line. End with CNTL/Z.

Router(config)#config

Router(config)#config-register 0x2102

Restart the router.

You will proceed through user mode but to make any changes you make must be at the global configuration prompt.

Finally, there is no way to recover a password by answering a security question.

Objective:

Infrastructure Management

Sub-Objective:

Use Cisco IOS tools to troubleshoot and resolve problems

References:

[Cisco > Home>Support>Product Support>End-of-Sale and End-of-Life Products>Cisco IOS Software Releases 12.1 Mainline>Troubleshoot and Alerts> Troubleshooting TechNotes> Password Recovery Procedures](#)

QUESTION 84

Which of the following is NOT a characteristic of private Internet Protocol (IP) addressing?

- A. These addresses are not routable through the public Internet.
- B. These addresses are publicly registered with the Internet Network Information Center (InterNIC).
- C. These addresses are reserved by the Internet Assigned Numbers Authority (IANA).
- D. The ranges of private IP addressing are 10.0.0.0 to 10.255.255.255, 172.16.0.0 to 172.31.255.255, and 192.168.0.0 to 192.168.255.255.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

It is NOT correct to state that private IP addresses are publicly registered with the Internet Network Information Center (InterNIC). Only public IP addresses are registered with the InterNIC.

The following characteristics are TRUE regarding private IP addressing:

- Private addresses are not routable through the public Internet.
- Private addresses are reserved by the Internet Assigned Numbers Authority (IANA).
- The ranges of private IP addressing are 10.0.0.0 to 10.255.255.255, 172.16.0.0 to 172.31.255.255 and 192.168.0.0 to 192.168.255.255.
- The use of private IP addressing conserves the use of the public IPv4 address space

Private addresses cannot be registered publicly.

Private IP addresses conserve public addressing space and can improve network security.

Objective:

Network Fundamentals

Sub-Objective:
Describe the need for private IPv4 addressing

References:

QUESTION 85

From which of the following attacks can Message Authentication Code (MAC) shield your network?

- A. DoS
- B. DDoS
- C. spoofing
- D. SYN floods

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Message Authentication Code (MAC) can shield your network from spoofing attacks. Spoofing, also known as masquerading, is a popular trick in which an attacker intercepts a network packet, replaces the source address of the packets header with the address of the authorized host, and reinserts fake information which is sent to the receiver. This type of attack involves modifying packet contents. MAC can prevent this type of attack and ensure data integrity by ensuring that no data has changed. MAC also protects against frequency analysis, sequence manipulation, and ciphertext-only attacks.

MAC is a secure message digest that requires a secret key shared by the sender and receiver, making it impossible for sniffers to change both the data and the MAC as the receiver can detect the changes.

A denial-of-service (DoS) attack floods the target system with unwanted requests, causing the loss of service to users. One form of this attack generates a flood of packets requesting a TCP connection with the target, tying up all resources and making the target unable to service other requests. MAC does not prevent DoS attacks. Stateful packet filtering is the most common defense against a DoS attack.

A Distributed Denial of Service attack (DDoS) occurs when multiple systems are used to flood the network and tax the resources of the target system. Various intrusion detection systems, utilizing stateful packet filtering, can protect against DDoS attacks.

In a SYN flood attack, the attacker floods the target with spoofed IP packets and causes it to either freeze or crash. A SYN flood attack is a type of denial of service attack that exploits the buffers of a device that accept incoming connections and therefore cannot be prevented by MAC. Common defenses against a SYN flood attack include filtering, reducing the SYN-RECEIVED timer, and implementing SYN cache or SYN cookies.


Objective:
Infrastructure Security
Sub-Objective:
Configure, verify, and troubleshoot basic device hardening

References:

[Cisco > The Internet Protocol Journal, Volume 10, No. 4 > IP Spoofing](#)

QUESTION 86

In the diagram below, when a packet sent from the PC at 10.0.1.3 to the PC at 10.1.1.3 leaves the Fa0/1 interface of R1, what will be the source and destination IP and MAC addresses?



- A. source IP 10.1.1.2 destination IP 10.1.1.3
Source MAC ad.ad.ad.ad.ad.ad destination MAC ab.ab.ab.ab.ab.ab
- B. source IP 10.1.1.1 destination IP 10.1.1.3
Source MAC ad.dd.dd.dd.dd.dd destination MAC ab.ab.ab.ab.ab.ab
- C. source IP 10.0.1.3 destination IP 10.1.1.3
Source MAC ad.ad.ad.ad.ad.ad destination MAC ae.ae.ae.ae.ae.ae
- D. source IP 10.0.6.3 destination IP 10.1.1.3
Source MAC ad.ad.ad.ad.ad.ad destination MAC ae.ae.ae.ae.ae.ae

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The source IP address will be 10.0.1.3 and the destination IP address will be 10.1.1.3. The source MAC address will be ad.ad.ad.ad.ad.ad and the destination MAC address will be ae.ae.ae.ae.ae.ae.

The source and destination IP addresses never change as the packet is routed across the network. The MAC address will change each time a router sends the packet to the next router or to the ultimate destination. The switches do not change either set of addresses in the header; they just switch the frame to the correct switch port according to the MAC address table. Therefore, when the packet leaves R1, the source MAC address will be that of R1 and the destination MAC address will be that of the Fa0/0 interface of R2. The IP addresses will be those of the two workstations, 10.0.1.3 and 10.1.1.3.

When the workstation at 10.0.1.3 starts the process, it will first determine that the destination address is in another subnet and will send to its default gateway (10.0.1.2). It will perform an ARP broadcast for the MAC address that goes with 10.0.1.2, and R1 will respond with its MAC address, bb.bb.bb.bb.bb.bb.

After R2 determines the next-hop address to send to 10.0.1.3 by parsing the routing table, it will send the packet to R1 at 10.0.6.2. When R2 receives the packet, R2 will determine that the network 10.0.1.0/24 is directly connected and will perform an ARP broadcast for the MAC address that goes with 10.0.1.3. The workstation at 10.0.1.3 will respond with its MAC address, ab.ab.ab.ab.ab.ab.

Objective:

Routing Fundamentals
Sub-Objective:
Describe the routing concepts

References:
[Cisco > IOS Technology Handbook > Routing Basics](#)

QUESTION 87

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.

Objective:

Network Fundamentals

Sub-Objective:

Describe the effects of cloud resources on enterprise network architecture

References:
[IaaS, PaaS, SaaS \(Explained and Compared\)](#)

QUESTION 88

What command produced the following as a part of its output?

```
1 14.0.0.2 4 msec 4 msec 4 msec
2 63.0.0.3 20 msec 16 msec 16 msec
3 33.0.0.4 16 msec * 16 msec
```

- A. Ping
- B. Traceroute
- C. Tracert
- D. Extended ping

Correct Answer: B

Section: (none)**Explanation****Explanation/Reference:****Explanation:**

The output displayed is a part of the output from executing the traceroute command. The traceroute command finds the path a packet takes while being transmitted to a remote destination. It is also used to track down routing loops or errors in a network. Each of the following numbered sections represents a router being traversed and the time the packet took to go through the router:

```
1 14.0.0.2 4 msec 4 msec 4 msec
2 63.0.0.3 20 msec 16 msec 16 msec
3 33.0.0.4 16 msec * 16 msec
```

The output would not be displayed by the ping command. This command is used to test connectivity to a remote ip address. The output from the ping command is as follows:

```
router1# ping 10.201.1.11
```

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

```
Sending 5, 100-byte ICMP Echos to 10.201.1.11, timeout is 2 seconds:
```

```
.....
```

```
Success rate is 0 percent (0/5)
```

The ping in this output was unsuccessful, as indicated by the Success rate is 0 percent output.

The output would not be displayed by the tracert command. The tracert command is used by Microsoft Windows operating systems, not the Cisco IOS command line interface. However, the purpose of the tracert command is similar to the Cisco traceroute utility, which is to test the connectivity or "reachability" of a network device or host. The tracert command uses Internet Control Message Protocol (ICMP).

The output would not be displayed by the extended version of the ping command. This command can be issued on the router to test connectivity between two remote routers. A remote execution means that you are not executing the command from either of the two routers you are interested in testing, but from a third router.

To execute an extended ping, enter the ping command from the privileged EXEC command line without specifying the target IP address. The command takes the router into configuration mode, where you can define various parameters, including the destination and target IP addresses. An example is below:

```
Protocol [ip]:
Target IP address: 10.10.10.1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 12.1.10.2
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.
```

Each line is a menu question allowing you to either accept the default setting (in parenthesis) of the ping or apply a different setting. The real value of this command is that you can test connectivity between two remote routers without being physically present at those routers, as would be required with the standard version of the ping command.

Objective:

Routing Fundamentals

Sub-Objective:
Troubleshoot basic Layer 3 end-to-end connectivity issues

References:
[Cisco > Cisco IOS Command Fundamentals Reference, Release 12.4 > ping](#)
[Cisco > Tech Notes > Using the Extended ping and Extended traceroute Commands > Document ID: 13730 > The Extended ping Command](#)

QUESTION 89

Which command can be issued at the following prompt?

Router(config-router)#

- A. show interface
- B. network
- C. interface
- D. ip default-gateway

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The network command can be issued at the Router(config-router)# prompt, which also indicates that the router is in router configuration mode. The network command is used to configure the network upon which a routing protocol is functioning.

The router configuration mode is accessed by issuing the router command in the global configuration mode along with a parameter indicating the routing protocol to be configured. For example:

```
R4(config)#router eigrp 1
```

changes the prompt to:

```
R4(config-router)#
```

which then allows you to specify the network as follows:

```
R4(config-router)#network 192.18.5.0
```

All other options are incorrect as these commands can be issued only in the global configuration command mode (which would be indicated by the R4(config)# prompt).

Objective:

Infrastructure Management

Sub-Objective:

Configure and verify initial device configuration

References:
[Cisco > Support > Cisco IOS Software > Using the Command-Line Interface in Cisco IOS Software](#)

QUESTION 90

Which Cisco Internetwork Operating System (IOS) command would be used to set the privileged mode password to "cisco"?

- A. router(config)# enable password cisco
- B. router# enable secret cisco
- C. router(config)# line password cisco
- D. router(config-router)# enable password cisco

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The enable password command is used to set the local password to control access to privileged levels. This command is executed on the global configuration mode, as in `router(config)# enable password cisco`. The syntax of the command is:

`router(config)# enable password [level level] {password | [encryption-type] encrypted-password}`

The parameters of the command are as follows:

- `level level`: An optional parameter to set the privilege level at which the password applies. The default value is 15.
- `password`: Specifies the password that is used to enter enable mode.
- `encryption-type`: An optional parameter to specify the algorithm used to encrypt the password.
- `encrypted-password`: Specifies the encrypted password that is copied from another router configuration.

The `router# enable secret cisco` command is incorrect because the `enable secret` command must be executed from global configuration mode, not privileged EXEC mode. In fact, this is the password for which you will be prompted when you attempt to enter privilege exec mode.

The `line password` command is incorrect because this command is not a valid Cisco IOS command.

The `router(config-router)# enable password cisco` command is incorrect because the `enable password` command must be entered in global configuration mode.

Objective:

Infrastructure Security

Sub-Objective:

Configure, verify, and troubleshoot basic device hardening


References:

[Cisco > Cisco IOS Security Command Reference > E > enable password](#)

[Cisco > Cisco IOS Security Configuration Guide, Release 12.4 > Part 7: Secure Infrastructure > Configuring Security with Passwords, Privilege Levels and, Login Usernames for CLI Sessions on Networking Devices](#)

QUESTION 91

R1 and R2 are connected as shown in the diagram and are configured as shown in output in the partial output of the show run command.



```

R1#show run

version 12.0
hostname R1

interface s1
ip address 192.168.5.5 255.255.255.252

ip host R1 192.168.5.6

R2#show run

version 12.0
hostname R2
interface s1
ip address 192.168.5.6 255.255.255.252
ip host R1 192.168.5.5

```

The command ping R2 fails when executed from R1. What command(s) would allow R1 to ping R2 by name?

- A. R1(config)#int S1
R1(config-if)#no ip address 192.168.5.5
R1(config-if)# ip address 192.168.5.9 255.255.255.252
- B. R1(config)#no ip host R1
R1(config)# ip host R2 192.168.5.6 255.255.255.252
- C. R1(config)#no hostname R2
R1(config)# hostname R1
- D. R2(config)#int S1
R1(config-if)#no ip address 192.168.5.5
R1(config-if)# ip address 192.168.5.9 255.255.255.0

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Both routers have been configured with the ip host command. This command creates a name to IP address mapping, thereby enabling the pinging of the device by address. On R1, the mapping is incorrect and needs to be corrected. Currently it is configured as ip host R1 192.168.5.6. It is currently mapping its own name to the IP address of R2.

To fix the problem, you should remove the incorrect IP address mapping and create the correct mapping for R2, as follows:

```

R1(config)#no ip host R1
R1(config)# ip host R2 192.168.5.6 255.255.255.252

```

Once this is done, the ping on R2 will succeed.

The IP address of the S1 interface on R1 does not need to be changed to 192.168.5.9 /30. In fact, if that is done the S1 interface on R1 and the S1 interface in R2 will no longer be in the same network. With a 30-bit mask configured, the network they are currently in extends from 192.168.5.4 - 192.168.5.7. They are currently set to the two usable addresses in that network, 192.168.5.5 and 192.168.5.6.

The hostnames of the two routers do need to be set correctly using the hostname command for the ping to function, but they are correct now and do not need to be changed.

The subnet mask of the S1 interface on R2 does not need to be changed to 255.255.255.0. The mask needs to match that of R1, which is 255.255.255.252.

Objective:

Infrastructure Services

Sub-Objective:

Troubleshoot client connectivity issues involving DNS

References:

QUESTION 92

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.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot interswitch connectivity

References:

[Cisco > Support > Cisco IOS IP Addressing Services Command Reference > show ip interface](#)

QUESTION 93

You and your team are evaluating the use of OSPFv3 in your IPv6 network.

Which of the following statements is true of OSPFv3?

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

- Router IDs should not match.
- Router IDs should reflect the correct router ID for each device.
- Area IDs must match.
- Area types must match.

Objective:

Routing Fundamentals

Sub-Objective:


Configure, verify, and troubleshoot single area and multi-area OSPFv3 for IPv6 (excluding authentication, filtering, manual summarization, redistribution, stub, virtual-link, and LSAs)

References:

[Home > Network Infrastructure > IPv6 Integration and Transition > Troubleshooting OSPFv3 Neighbor Adjacencies](#)

QUESTION 94

You have established a console session with R1 and you are attempting to download an IOS image from the TFTP server in the diagram below.



However, you are unable to make the connection to 150.140.6.5. What is the problem?

- A. The IP address of the management station is incorrect
- B. The IP address of the TFTP server is incorrect
- C. The interfaces between R1 and R2 are not in the same subnet
- D. The IP address of Switch B is incorrect

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The IP address of the TFTP server is incorrect. The TFTP server, Switch B and the Fa0/2 interface on R3 should all be in the same subnet. With a 27-bit mask (255.255.255.224) against the 150.140.0.0 classful network the resulting subnets are:

150.140.0.0
150.140.0.32
150.140.0.64

and so on, incrementing in intervals of 32 in the last octet until it reaches the 150.140.6.0 subnet.

150.140.6.0
150.140.6.32
150.140.6.64

At this point, we can see that Switch B and the router interface are in the 150.140.6.32 subnet, while the TFTP server is in the 150.140.6.0 subnet. The IP address of the TFTP server needs to be in the 150.140.6.33-150.140.6.62 range, while avoiding the addresses already used on R1 and the switch.

The IP address of the management station does not appear to be in any of the networks listed in the diagram, but that doesn't matter since the connection to the router is through the console cable which does not require a correct IP address.

The Fa0/2 and Fa0/1 interfaces on R1 and R2 are in the same subnet. Using a 25-bit mask against the

192.18.5.0/24 classful network yields the following subnets:

192.18.5.0
192.168.5.128

Both router interfaces in question are in the 192.18.5.0 subnet.

As we have already determined, the IP address of Switch B is correct. Even if it were incorrect or missing altogether, it would have no impact on connecting to the TFTP server. Switches merely switch frames based on MAC addresses and only need an IP address for management purposes.

Objective:

Routing Fundamentals

Sub-Objective:


Troubleshoot basic Layer 3 end-to-end connectivity issues

References:

[Cisco > Support > IP Routing > Design TechNotes > Document ID: 13788 > IP Addressing and Subnetting for New Users](#)

QUESTION 95

Consider the following diagram:



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

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

advertisements.

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

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

Objective:

Routing Fundamentals

Sub-Objective:


Compare and contrast distance vector and link-state routing protocols

References:

[Cisco > Home > Support > Technology Support > IP > IP Routed Protocols > Design > Design TechNotes > Why Don't IGRP and RIP v1 support VLSM?](#)

QUESTION 96

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



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

Correct Answer: ADFGH

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

1. Selection of the root bridge. All ports on the root bridge become designated ports and are set to forward.
2. Determination of the root ports on each non-root bridge.
3. Determination of the designated port on each segment that does not connect directly to the root bridge.
4. Designated and root ports will be set to forwarding, and all other ports will be set to discarding.

For step 1, when all bridge priorities have been left to their default, all switches will have same bridge priority. When that is the case, as in this scenario, the switch with the lowest MAC address will be selected as root bridge. In this case, SW3 has the lowest MAC address and becomes the root bridge. ALL ports are in a forwarding state on the root bridge, which explains why all of the ports on SW3 will be in a forwarding state.

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

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

SW1 to SW3 (100 Mbps) cost = 19

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

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

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

SW1 will use the lowest cost path (SW1 to SW4 to SW3) as its root path. It will set the SW1 connection to SW4 to forwarding and the connection from SW1 to SW3 to blocking. The status of its third interface (SW1 to SW2) will be determined in Step 3, since it is a shared segment with SW2 that does not have a direct connection to the root bridge.

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

SW2 to SW3 (1 Gbps) cost = 4

SW2 to SW1 to SW3 (100 Mbps + 100 Mbps) cost = 19 + 19 = 38

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

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

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

SW4 to SW3 (1 Gbps) cost = 4

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

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

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

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


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

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

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

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

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



Once STP has converged, the port states will determine the path used when sending traffic from a host connected to one switch to a host connected to another switch. For example, if a host connected to SW3 were destined for a host connected to SW2, the path taken would be SW3 to SW2. It would not take SW3-SW1-SW2 or SW3-SW4-SW2 because on both of those paths, STP is blocking at least one port in the path.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot STP protocols

References:

[Home > Support > Technology Support > LAN Switching > Spanning Tree Protocol > Technology Information>Technology White Paper>Understanding Rapid Spanning Tree Protocol \(802.1w\)](#)

QUESTION 97

Which of the following characteristics are NOT shared by RIPv1 and RIPv2?

- A. They share an administrative distance value
- B. They use the same metric
- C. They both send the subnet mask in routing updates
- D. They have the same maximum hop count

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

RIPv1 and RIPv2 do NOT both send the subnet mask in routing updates. RIPv1 is classful, while RIPv2 is classless. This means the RIPv1 does not send subnet mask information in routing updates, while RIPv2 does.

Both versions have the same administrative distance of 120.

Both versions have the same metric, which is hop count.

Both versions have the same maximum hop count, which is 15.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot RIPv2 for IPv4 (excluding authentication, filtering, manual summarization, redistribution)

References:

[Home > Knowledgebase > Cisco Certified Network Associate \(CCNA\) > Difference between RIPv1 and RIPv2](#)
[Cisco Press > Articles > Cisco Certification > CCDA > CCDA Self-Study: RIP, IGRP, and EIGRP Characteristics and Design](#)

QUESTION 98

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.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot EIGRP for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub)

References:

[Internetworking Technology Handbook > Enhanced Interior Gateway Routing Protocol \(EIGRP\)](#)

QUESTION 99

Which of the following statements is TRUE about trunk ports?

- A. A trunk port connects an end-user workstation to a switch.
- B. A trunk port uses 802.1q to identify traffic from different VLANs.
- C. A trunk port supports a single VLAN.
- D. A trunk port uses a straight-through Ethernet cable when connecting two switches.

Correct Answer: B

Section: (none)


Explanation

Explanation/Reference:

Explanation:

A switch port can operate as an access port or a trunk port. An access port is used to connect to an end-user device, such as a workstation, server, or printer, while a trunk port is used to connect to neighboring switches or routers. The trunk link is responsible for carrying data between workstations connected to different switches, or a switch and a router configured for inter-VLAN routing. For example, in the diagram below where VLANs are in use on both switches and inter-VLAN routing is configured, the interfaces will operate as follows:

- SW1 - Fa0/1 and Fa0/2 are trunk links, Fa0/0 is an access link
- SW2 - Fa0/0 is trunk link and Fa0/1 is an access link
- RTR - Fa0/2 is a trunk link



With the exception of frames traveling on the native VLAN, data frames crossing a trunk link must be frame tagged over the link to identify the VLAN that sourced the frame. The receiving switch sees the VLAN ID, and uses this information to forward the frame appropriately. 802.1q and ISL are the two possible frame tagging methods between Cisco switches. In summary, some facts about access and trunk ports:

Access ports:

- Carry traffic for a single VLAN
- Connect end user workstations to the switch
- Use a straight-through cable to connect to the device

Trunk ports:

- Facilitate inter-VLAN communication when connected to a Layer 3 device
- Carry traffic from multiple VLANs
- Use 802.1q to identify traffic from different VLANs

When a new trunk link is created on a switch, all VLANs are allowed to use the trunk, by default.

Trunk ports are used between switches and routers, and do not connect to end-user workstations.

Trunk ports support all VLANs known to the switch by default, so that devices in the same VLAN can communicate across multiple switches. Trunk ports are not limited to a single VLAN, as access ports are.

Trunk ports connected between switches using crossover Ethernet cables, not straight-through Ethernet cables. Trunk ports between switches and routers use straight-through cables.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot interswitch connectivity

References:

[Home > Articles > Network Technology > General Networking > VLANs and Trunking](#)

QUESTION 100

In which two situations would it be appropriate to issue the ipconfig command with the /release and /renew options? (Choose two.)

- A. When the result of running the ipconfig /all command indicates a 169.254.163.6 address
- B. When recent scope changes have been made on the DHCP server
- C. When no IP helper address has been configured on the router between the client and the DHCP server
- D. When the no ip directed-broadcast command has been issued in the router interface local to the client, and no IP helper address has been configured on the router between the client and the DHCP server

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Explanation:

It would be appropriate to issue the ipconfig command with the /release and /renew options when the result of running the ipconfig /all command indicates a 169.254.163.6 address, or when recent scope changes have been made on the DHCP server. When a computer has an address in the 169.254.0.0 network, it indicates that the computer has not been issued an address from the DHCP server. Instead, the computer has utilized Automatic Private IP Addressing (APIPA) to issue itself an address. If the reason for this assignment is a temporary problem with the DHCP server or some other transitory network problem, issuing the ipconfig /release command followed by the ipconfig /renew command could allow the computer to receive the address from the DHCP server.

Similarly, if changes have been made to the settings on the DHCP server, such as a change in the scope options (such as gateway or DNS server), issuing this pair of commands would update the DHCP client with the new settings when his address is renewed.

These commands will have no effect when no IP helper address has been configured on the router between the client and the DHCP server. An IP helper address can be configured on the local interface of a router when no DHCP server exists on that subnet and you would like to allow the router to forward DHCP DISCOVER packets to the DHCP server on a remote subnet. DHCP DISCOVER packets are broadcast, and routers do not pass on broadcast traffic by default.

These commands also will be of no benefit if the no ip directed-broadcast command has been issued in the router interface local to the client and no IP helper address has been configured on the router between the client and the DHCP server. The no ip directed-broadcast command instructs the router to deny broadcast traffic (which is the default). Under those conditions, the command will not result in finding the DHCP server or receiving an address.

Objective:

Infrastructure Services

Sub-Objective:

Troubleshoot client- and router-based DHCP connectivity issues

References:

QUESTION 101

What data structure is pictured in the graphic?

0-15	16-31
Source Port Number	Destination Port Number
Length	Checksum
Data	

- A. TCP segment
- B. UDP datagram
- C. IP header
- D. Http header

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The data structure pictured in the graphic is an UDP datagram. It uses a header (not shown) that contains the source and destination MAC address. It has very little overhead as compared to the TCP segmented (shown later in this explanation) as any transmission that uses UDP is not provided the services of TCP.

It is not a TCP segment, which has much more overhead (shown below). The TCP header contains fields for sequence number, acknowledgment number, and windows size, fields not found in a UDP header because UDP provides none of the services that require use of these fields. That is, UDP cannot re-sequence packets that arrive out of order, nor does UDP acknowledge receipt (thus the term non-guaranteed to describe UDP). Furthermore, since UDP does not acknowledge packets there is no need to manage the window size (the window size refers to the number of packets that can be received without an acknowledgment).

Bit 0				Bit 15				Bit 16				Bit 31			
Source Port (16)								Destination Port (16)							
Sequence Number (32)															
Acknowledgement Number (32)															
Header length (4)				Reserved				Code Bits (6)				Window (16)			
Checksum (16)								Urgent (16)							
Options (0 or 32 if any)															
Data (Varies)															

It is not an IP header. An IP header contains fields for the source and destination IP address. The IP header, like the UDP segment, does not contain fields for sequence number, acknowledgment number, and windows size, fields not found in a TCP header because TCP provides none of the services that require use of these fields. IP provides best-effort user data. This does not cause a delivery problem, however, as IP relies on TCP to provide those services when the transmission is a unicast.

An HTTP header does not include fields for HTTP requests and responses.

Objective:

Network Fundamentals

Sub-Objective:

Compare and contrast TCP and UDP protocols

References:

[Cisco > Home > Internetworking Technology Handbook > Internet Protocols > User Datagram Protocol \(UDP\)](#)

QUESTION 102

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.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot EIGRP for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub)

References:

[Cisco > Cisco IOS IP Routing Protocols Command Reference > EIGRP Commands: M through V > show ip eigrp topology](#)

QUESTION 103

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.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot EIGRP for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub)

References:

[Cisco > Support > IP > IP Routing > Technology Information > Technology White Paper > Enhanced Interior Gateway Routing Protocol > Document ID: 16406 > Feasible Distance, Reported Distance, and Feasible Successor](#)

QUESTION 104

Which two are NOT valid Cisco IOS commands used for Cisco Discovery Protocol (CDP)? (Choose two.)

- A. show cdp
- B. show cdp entry *
- C. show cdp neighbor entries
- D. show cdp neighbors detail

E. show cdp devices

Correct Answer: CE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The show cdp neighbor entries command and the show cdp devices command are not valid Cisco IOS commands.

The Cisco IOS commands used for CDP are as follows:

show cdp: This command is used to view global CDP information, such as timer and hold time.

show cdp entry *: This command is used to view information regarding all neighboring devices.

show cdp neighbors detail: This command is used to view the details regarding the neighboring devices which are discovered by the CDP. This command is used to view details such as network address, enabled protocols, and hold time. The complete syntax of this command is:

show cdp neighbors [type number] [detail]

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure and verify Layer 2 protocols

References:

[Cisco > Cisco IOS Network Command Reference, Release 12.4 > show cdp neighbors](#)

QUESTION 105

Which of the following fields are in a Transmission Control Protocol (TCP) header? (Choose three.)

- A. Length
- B. Sequence Number
- C. Data Offset
- D. Type-of-Service
- E. Window

Correct Answer: BCE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

- Sequence Number, Data Offset, and Window are the fields found in a TCP header. TCP hosts create a connection-oriented session with one another. The following are the fields found in a TCP header:
- Sequence Number: Refers to the first byte of data in the current message. This field helps TCP to reassemble the packets in the correct order. For example, when data is transferred between an FTP server and FTP client, the receiver uses this field to reassemble the packets into the original file.
- Data Offset: Refers to the number of 32-bit words in the TCP header.
- Window: Refers to the size of the available space for the incoming data.
- Source Port and Destination Port: Refer to the point where upper-layer source and destination processes receive TCP services. Both TCP and UDP packets contain these fields.
- Acknowledgment Number: Refers to the sequence number of the next byte of data which the sender will receive.
- Reserved: Reserved for future use.
- Flags: Contains control information, such as the SYN and ACK bits which are used to establish and acknowledge communication, and the FIN bit which is used to terminate the connection.

- Checksum: An indicator of any damage to the header while being in transit. Both TCP and UDP packets contain this field.
- Urgent Pointer: Refers to the first urgent data byte in the packet.
- Options: Used to specify TCP options. Only TCP packets contain this field.
- Data: Has upper-layer information.

TCP is used for unicast transmissions and provides connection-oriented services for upper layer protocols. It will establish a state of connection between two devices before any data is transferred; for example, before a workstation can exchange HTTP packets with Web server, a TCP connection must be established between the workstation and the Web server.

The Length field is found in a User Datagram Protocol (UDP) header, where it specifies the length of the UDP header and data. UDP headers contain the Source Port, Destination Port, Length, and Checksum fields.

Sequence number, acknowledgment number, and windows size are fields not found in a UDP header because UDP provides none of the services that require use of these fields. That is, UDP cannot re-sequence packets that arrive out of order, nor does UDP acknowledge receipt (thus the term non-guaranteed to describe UDP). Furthermore, since UDP does not acknowledge packets, there is no need to manage the window size, which refers to the number of packets that can be received without an acknowledgment.

The Type-of-Service field is found in an Internet Protocol (IP) header, where it specifies the handling of a current datagram by an upper-layer protocol.

Objective:

Network Fundamentals

Sub-Objective:

Compare and contrast TCP and UDP protocols

References:

[Cisco Documentation > Internetworking Technology Handbook > Routing Basics > Internet Protocols > TCP Packet Format](#)

QUESTION 106

Which Cisco IOS command disables Cisco Discovery Protocol Version 2 (CDPv2) advertisements?

- A. no cdp advertise-v2
- B. no cdp v2-advertise
- C. no cdp run
- D. no cdp enable

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The no cdp advertise-v2 command disables CDPv2 advertisements. It is the reverse of the cdp advertise-v2 command, which enables CDPv2 advertisements on a device.

The no cdp v2-advertise command is not a valid Cisco IOS command.

The no cdp run command disables CDP, not CDPv2 advertisements.

The no cdp enable command disables CDP on an interface.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure and verify Layer 2 protocols

References:

[Cisco > Support > Using Cisco Discovery Protocol](#)

QUESTION 107

Which command would you use to see which switch interface is associated with a particular MAC address?

- A. show interface mac
- B. show mac
- C. show mac-address-table
- D. show ip interface

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The show mac-address-table command displays a table of every learned MAC address, and the switch port associated with the MAC address. Sample output is as follows:

```
Switch# show mac-address-table
Mac Address Table
-----
Vlan Mac Address Type Ports
----
1 0040.63d8.ba0a DYNAMIC Fa0/1
1 0004.274c.9ca0 DYNAMIC Fa0/3
1 0040.63d8.bab8 DYNAMIC Fa0/10
1 000f.1fd3.d85a DYNAMIC Fa0/7

Total Mac Addresses for this criterion: 4
```

This output indicates that four MAC addresses have been learned by this switch, and the last column indicates the switch port over which each MAC address was learned, and for which frames destined for each MAC address will be forwarded. The MAC address table is built dynamically by examining the source MAC address of received frames. If the switch receives a MAC address not listed in this table, it will send the frame out all ports except the one from which it was originated.

The show ip interface command is a router command, and displays no information on MAC address tables.

The show interface mac and show mac commands are incorrect because they are not valid Cisco IOS commands.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Describe and verify switching concepts

References:

QUESTION 108

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.

Objective:

Infrastructure Services

Sub-Objective:

Configure, verify, and troubleshoot basic HSRP

References:

[Cisco > Cisco IOS IP Application Services Command Reference > show standby through show udp > show standby](#)

QUESTION 109

What is the significance of the following BECN packet statistics?

```
Router# show frame-relay pvc 16

PVC Statistics for interface serial0 (Frame Relay DTE)

DLCI = 16, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE = Serial0
input pkts 0 output pkts 0 in bytes 0
out bytes 0 dropped pkts 0 in FECN pkts 0
in BECN pkts 100 out FECN pkts 0 out BECN pkts 0
<<output omitted>>
```

- A. The router is experiencing congestion in sending frames.
- B. The router is experiencing congestion in receiving frames.
- C. The Frame Relay mapping table is missing an entry.
- D. The Frame Relay mapping table is corrupt.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

When frames arrived at a router with the Backwards Explicit Congestion Notification (BECN) bit set, congestion was encountered in the opposite direction from which the frame was traveling. This bit is set by the Frame Relay switch. If an incoming packet has the BECN bit set, then this indicates congestion in outgoing packets, so the router will experience congestion in sending frames.

When a Frame Relay switch encounters congestion, it will mark packets being sent in both directions on a PVC with either the Forward Explicit Congestion Notification (FECN) or the BECN bit set. It will set the BECN bit on packets headed in the opposite direction of the congestion and FECN in the same direction as the congestion. When a packet with the FECN bit is received by a router, it means there will be congestion when the receiving router receives packets.

A third type of marking is the Discard Eligibility (DE) bit. When this bit is set on a packet, it ensures that if congestion occurs and packets need to be discarded, the packet with the DE bit set should be discarded first.

ALL packets in excess of the committed information rate (CIR) are marked with the DE bit.

Frame Relay mapping tables have nothing to do with congestion in the Frame Relay network.

Objective:

WAN Technologies

Sub-Objective:

Describe basic QoS concepts

References:

[Cisco > Home > Support > Technology Support > WAN > Frame Relay > Design > Design TechNotes > show Commands for Frame Relay Traffic Shaping](#)

QUESTION 110

In the following partial output of the show ip route command, what does the letter D stand for?

```
D 192.1.2.0/24 via 5.1.1.71 [w:0 m:0]
```

```
C 192.8.1.1/32 directly connected to loopback 0
```

- A. This is a default route
- B. This is an EIGRP route
- C. This is static route
- D. This is a directly connected route

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The letter D indicates that it was a route learned by the EIGRP routing protocol. In the output of the show ip route command, each route will have a letter next to it that indicates the method by which the route was learned. At the beginning of the output will be a legend describing the letters as shown below:

```
Router# show ip route
Codes: I - IGRP derived, R - RIP derived, O - OSPF derived,
C - connected, S - static, E - EGP derived, B - BGP derived,
* - candidate default route, IA - OSPF inter area route,
i - IS-IS derived, ia - IS-IS, U - per-user static route,
o - on-demand routing, M - mobile, P - periodic downloaded static route,
D - EIGRP, EX - EIGRP external, E1 - OSPF external type 1 route,
E2 - OSPF external type 2 route, N1 - OSPF NSSA external type 1 route,
N2 - OSPF NSSA external type 2 route
```

The letter does not indicate that it is a default route. The default route (if configured) will appear at the end of the legend as follows:

```
Gateway of last resort is 10.119.254.240 to network 10.140.0.0
```

The letter does not indicate that it is a static route. Static routes will have an "S" next to them.

The letter does not indicate that it is a directly connected route. Directly connected routes will have a "C" next to them.

Objective:

Routing Fundamentals

Sub-Objective:

Interpret the components of routing table

References:

[Cisco > Support > Cisco IOS IP Routing: Protocol-Independent Command Reference > show ip route](#)

QUESTION 111

The conference room has a switch port available for use by the presenter during classes. Each presenter uses the same PC attached to the port. You would like to prevent any other PCs from using that port. You have completely removed the former configuration in order to start anew.

Which of the following steps are required to prevent any other PCs from using that port?

- A. make the port a trunk port
- B. enable port security
- C. make the port an access port
- D. assign the MAC address of the PC to the port
- E. make the port a sticky port
- F. set the maximum number of MAC addresses on the port to 1

Correct Answer: BCD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should create the port as an access port, enable port security, and statically assign the MAC address of the PC to the port. Creating the port as an access port ensures that the PC can use the port and port security can be enabled on the port. The second step is to enable port security, which is required to use the third command. The third command sets the MAC address of the PC as the statically assigned address on that port, meaning that only that address can send and receive on the port.

You should not make the port a trunk port. There is no need to make this a trunk port because it will not be carrying multiple VLAN traffic, only the traffic of the PC.

You should not make the port a sticky port. The sticky keyword, when used with switchport port-security command, is used to allow a port to dynamically learn the first MAC address it sees in the port, add it to the MAC address table, and save it to the running configuration of the switch. It will not limit the MAC addresses allowed on the port to that of the PC.

You should not set the maximum number of MAC addresses on the port to 1. That would prevent the attachment of a hub or switch to the port, but would not restrict the MAC addresses allowed on the port to the MAC address of the PC.

Objective:

Infrastructure Security

Sub-Objective:

Configure, verify, and troubleshoot port security

References:

[Catalyst 4500 Series Switch Cisco IOS Software Configuration Guide, 12.2\(20\)EWA > Configuring Port Security](#)

QUESTION 112

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

S0/0 - 192.168.5.1/24
S0/1 - 10.0.0.6/8

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

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

- A. Router5(config)#ipv6 ospf 25
Router5(config)# network 192.168.5.0
- B. Router5(config)#ipv6 ospf 25
Router5(config)#router-id 192.168.5.1
- C. Router5(config)#ipv6 unicast-routing
Router5(config)#ipv6 router ospf 25
Router5(config-rtr)#router-id 1.1.1.1
Router5(config)#interface S0/0
Router5(config-if)#ipv6 ospf 25 area 0
- D. Router5(config)#ipv6 unicast-routing
Router5(config)#ipv6 ospf 25
Router5(config-rtr)#router-id 1.1.1.1

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The correct command sequence would be as follows:

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

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

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

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

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

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

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

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

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

```
Router5(config)# ipv6 unicast-routing
Router5(config)# ipv6 ospf 25
```

Router5(config-rtr)# router-id 1.1.1.1

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot single area and multi-area OSPFv3 for IPv6 (excluding authentication, filtering, manual summarization, redistribution, stub, virtual-link, and LSAs)

References:

[Cisco > Implementing OSPF for IPv6 > How to Implement OSPF for IPv6](#)

[Cisco > Cisco IOS IPv6 Command Reference > ipv6 unicast-routing](#)

[Cisco > Cisco IOS IPv6 Command Reference > ipv6 ospf area](#)

QUESTION 113

Which command is used to view the entire routing table?

- A. show route-map
- B. show ip mroute
- C. show ip route
- D. show ip protocols

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The show ip route command is used to view the entire routing table. The output of this command consists of codes, gateway of last resort, directly connected networks, and routes learned through different protocols working on the network. The syntax of the show ip route command is as follows:

show ip route [address [mask] [longer-prefixes]] | [protocol [process-id]]

The parameters of the show ip route command are as follows:

- address: Specifies the address for which the routing information should be displayed.
- mask: Specifies the subnet mask.
- longer-prefixes: Specifies the combination of mask and address.
- protocol: Specifies the name of the routing protocols such as Routing Information Protocol (RIP), or Open Shortest Path First (OSPF).
- protocol-id: Specifies the protocol ID used to identify a process of a particular protocol.

The show route-map command is incorrect because this command is used to view the route-maps configured on the router.

The show ip mroute command is incorrect because this command is used to view the contents of the IP multicast routing table.

The show ip protocols command is incorrect because this command is used to view the routing protocols parameters, and the current timer values.

Objective:

Routing Fundamentals

Sub-Objective:

Interpret the components of routing table

References:

QUESTION 114

DRAG DROP

Click and drag the VLAN Trunking Protocol (VTP) mode descriptions on the left to their corresponding VTP modes on the right. (The descriptions on the left can be used more than once.)

Select and Place:

Descriptions	Server Mode	Client Mode	Transparent Mode
Switch can add, modify, or delete VLANs.			
Switch can generate VTP messages.			
Switch can forward VTP messages.			
Switch can synchronize VTP information.			

Correct Answer:

- Switches forward advertisements.
- Does not synchronize VTP information.
- VLAN information is saved in NVRAM.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot inter-VLAN routing


References:

[Support > Technology Support > LAN Switching > Virtual LANs/VLAN Trunking Protocol \(VLANs/VTP\) > Configure > Configuration Examples and TechNotes > Configuring VLAN Trunk Protocol \(VTP\)](#)

[CCNA Routing and Switching Complete Study Guide: Exam 100-105, Exam 200-105, Exam 200-125, 2nd Edition](#), Chapter 2: LAN Switching Technologies - Configure, verify, and troubleshoot VLANs (normal/extended range) spanning multiple switches

QUESTION 115

Which device in the given network diagram has as its primary responsibility the regulation of network traffic flow based on different trust levels for different computer networks?



- A. the router
- B. the switch
- C. the hub(s)

D. the firewall

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The firewall has as its primary responsibility the regulation of network traffic flow based on different trust levels for different computers or networks. In the network diagram shown in the exhibit, a firewall protects the network from unauthorized access attempts. A firewall can be implemented in hardware or software. Firewalls permit, deny, or filter data packets coming into and going out of the network. This helps prevent unauthorized access attempts from outside the network.

The primary function of a router is to perform routing between two subnets or between dissimilar network technologies. Routers can provide limited firewall functionality, but a firewall is a dedicated hardware or software solution with the primary responsibility of securing the network. A router does not have as its primary responsibility the regulation of network traffic flow based on different trust levels.

Switches work at Layer 2 in the Open System Interconnection (OSI) model and perform the function of separating collision domains. A switch does not have as its primary responsibility the regulation of network traffic flow based on different trust levels.

A hub is a device that provides a common connection point for network devices. The primary responsibility of a hub is not to regulate network traffic flow based on different trust levels.

Objective:

Network Fundamentals

Sub-Objective:

Describe the impact of infrastructure components in an enterprise network


References:

[Cisco > Home > Internetworking Technology Handbook > Internetworking Basics> Bridging and Switching Basics](#)

QUESTION 116

The exhibit displays the MAC address table of a switch in your network, along with the location of each device connected to the switch.

MAC Address Table	
Port	MAC Address
1	12-34-56-78-9A-BF
3	12-34-56-78-9A-BD



Which of the following frames will cause the switch to add a new MAC address to its table and forward the frame to all ports when the frame is received?

- A. source MAC: 12-34-56-78-9A-BC, destination MAC: ff-ff-ff-ff-ff
- B. source MAC: ff-ff-ff-ff-ff, destination MAC: 12-34-56-78-9A-BC
- C. source MAC: 12-34-56-78-9A-BF, destination MAC: 12-34-56-78-9A-BC
- D. source MAC: 12-34-56-78-9A-BC, destination MAC: 12-34-56-78-9A-BF

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The only frame that will be handled in the specified way is the one with a source MAC of 12-34-56-78-9A-BC and a destination MAC of ff-ff-ff-ff-ff. Since the source address 12-34-56-78-9A-BC is not already in the MAC table, the switch will add it. It will forward the frame to all ports because the destination is the broadcast MAC address of ff-ff-ff-ff-ff.

A frame with a source MAC of ff-ff-ff-ff-ff and a destination MAC of 12-34-56-78-9A-BC is an impossible combination. That would mean that the frame is coming from all devices, which is not possible.

The frame with a source MAC of 12-34-56-78-9A-BF and a destination MAC of 12-34-56-78-9A-BC would be sent to all ports because the destination MAC address is not in the MAC address table. However, the switch would not add a new MAC address to the table because the source address is already in the table.

The frame with a source MAC of 12-34-56-78-9A-BC and a destination MAC of 12-34-56-78-9A-BF would not be forwarded to all ports because the destination MAC address is in the table. The switch would add a new MAC address to the table because the source MAC address is not currently in the MAC address table.

Objective:

LAN Switching Fundamentals

Sub-Objective:
Interpret Ethernet frame format

References:

[Cisco Press > Articles > Cisco Certification > CCNA Routing and Switching > Basic Data Transmission in Networks: MAC Tables and ARP Tables](#)
[How do Switches Work?](#)

QUESTION 117

You have been assigned a network ID of 172.16.0.0/26. If you utilize the first network resulting from this ID, what would be the last legitimate host address in this subnet?

- A. 172.16.0.64
- B. 172.16.0.63
- C. 172.16.0.62
- D. 172.16.0.65

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

When a class B address such as 172.16.0.0 is subnetted with a /26 mask, the subnet mask in dotted decimal format is 255.255.255.192. This means that the interval between the network IDs of the resulting subnets is 64. The resulting network IDs are as follows:

172.16.0.0
172.16.0.64
172.16.0.128
172.16.0.192
172.16.1.0
and so on.

For the network ID 172.16.0.0, the last address in the range is 172.16.0.63, which is the broadcast address. Neither the network ID nor the broadcast address for any subnet can be assigned to computers. This means that the addresses that can actually be assigned range from 172.16.0.1 to 172.16.0.62. The last legitimate host address, therefore, is 172.16.0.62.

172.16.0.63 cannot be used because it is the broadcast address for the 172.16.0.0 network.

172.16.0.64 is the network ID for the 172.16.0.64 network, and 172.16.0.65 is the first address in the second network.

Objective:

Network Fundamentals

Sub-Objective:

Apply troubleshooting methodologies to resolve problems

References:

[Cisco > Support > IP Routing > Design TechNotes > Document ID: 13788 > IP Addressing and Subnetting for New Users](#)

QUESTION 118

Which Cisco IOS command enables a router to copy IOS images to a router?

- A. copy tftp flash

- B. copy flash tftp
- C. copy running-config tftp
- D. copy running-config startup-config
- E. copy tftp running-config

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The copy tftp flash command enables a router to copy an IOS image (the router operating system) to a router from a TFTP server. One router can act as a TFTP server to the other in this process.

The following example illustrates the steps to copy an image from Router A to Router B:

- Verify the connectivity between Router A and Router B using the ping command.
- Check the image size on both of the routers with the show flash command to verify that enough space exists on Router B.
- Configure Router A as the TFTP server using the configure terminal command. Use the tftp-server flash [partition-number:]filename1 [alias filename2] [access-list-number] command to define the path to system image that needs to be transferred. There can be multiple entries for multiple images.
- Copy the image from Router A to Router B using the copy tftp flash command.
- Verify the flash for the copied new image on Router B with the show flash command.

The copy flash tftp command is used to copy an IOS image from the router to a TFTP server.

The copy running-config tftp command is used to copy the active or running configuration file from RAM to a TFTP server.

The copy running-config startup-config command copies the active or running configuration from RAM to NVRAM. This command creates the configuration file that will be used as the startup configuration at reboot. This should always be done after making changes to the router so that the changes are saved when the router is rebooted.

The copy tftp running-config command merges a backup configuration with the currently active running configuration in RAM.

Objective:

Infrastructure Management

Sub-Objective:

Perform device maintenance

References:

QUESTION 119

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.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot single area and multi-area OSPFv2 for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub, virtual-link, and LSAs)

References:

[Cisco > Home > Support > Technology Support > IP > IP Routing > Design > Design TechNotes > OSPF Neighbor Problems Explained](#)

QUESTION 120

Which two are the limitations of the service password-encryption command? (Choose two.)

- A. It uses the MD5 algorithm for password hashing.
- B. It uses the Vigenere cipher algorithm.
- C. An observer cannot read the password when looking at the administrator's screen.
- D. The algorithm used by this command cannot protect the configuration files against detailed analysis by attackers.

Correct Answer: BD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The following are limitations of the service password-encryption command:

- It uses the Vigenere cipher algorithm, which is simple in nature.
- A cryptographer can easily crack the algorithm in a few hours.
- The algorithm used by this command cannot protect the configuration files against detailed analysis by attackers.

The service password-encryption command does not use the MD5 algorithm for password hashing. The MD5 algorithm is used by the enable secret command.

The option stating that an observer cannot read the password when looking at the administrator's screen is incorrect because this is an advantage of the service password-encryption command.

Objective:

Infrastructure Security

Sub-Objective:

Configure, verify, and troubleshoot basic device hardening

References:

[Cisco Documentation > Cisco IOS Security Command Reference, Release 12.4 > service password-encryption](#)
[Cisco > Tech Notes > Cisco Guide to Harden Cisco IOS Devices > Document ID: 13608](#)

QUESTION 121

DRAG DROP

Group the special DHCP messages exchanged over the network, on the left, into the different transmission types, on the right.

Select and Place:

DHCP Messages	Unicast	Multicast	Broadcast
DHCPACK			
DHCPOFFER			
DHCPREQUEST			
DHCPDISCOVER			

Correct Answer:

Sub-Objective:

Configure and verify DHCP on a router (excluding static reservations)

References:

[Cisco Documentation > Cisco IOS IP Configuration Guide, Release 12.2 > Part 1: IP Addressing and Services > Configuring DHCP](#)

QUESTION 122

Which command will save a dynamically learned MAC address in the running-configuration of a Cisco switch?

- A. switchport port-security mac-address
- B. switchport port-security
- C. switchport port-security sticky mac-address
- D. switchport port-security mac-address sticky
- E. switchport mac-address sticky

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Issuing the switchport port-security mac-address sticky command will allow a switch to save a dynamically learned MAC address in the running-configuration of the switch, which prevents the administrator from having to document or configure specific MAC addresses. Once the approved MAC addresses have all been learned, the network administrator simply saves the running-configuration file to NVRAM with the copy running-config startup-config command.

Switches dynamically build MAC address tables in RAM, which allow the switch to forward incoming frames to the correct target port. By default, an unlimited number of MAC addresses can be learned on a single switch port, whether it is configured as an access port or a trunk port. Switch ports can be secured by defining one or more specific MAC addresses that should be allowed to connect, and by defining violation policies (such as disabling the port) if additional hosts try to gain a connection. The following command secures a switch by manually defining an allowed MAC address:

switch(config-if)# switchport port-security mac-address 00C0.35F0.8301

This command statically defines the MAC address of 00c0.35f0.8301 as an allowed host on the switch port. Manually configuring all of your switch ports in this way, however, would require documenting all of your existing MAC addresses and configuring them specifically per switch port, which could be an extremely time-consuming task.

An example of the use of the switchport port-security mac-address sticky command is shown below:

Switch(config)#interface fastethernet0/16

Switch(config-if)#switchport port-security

Switch(config-if)#switchport port-security mac-address sticky

Switch(config-if)#switchport port-security maximum 1

With the above configuration, if a computer with a MAC address of 0000.00bb.bbbb were plugged into the switch, the following two things would occur:

- Only the host with MAC address 000.00bb.bbbb will be allowed to transmit on the port. This is a result of the port-security mac-address-sticky command, which instructs the switch to learn the next MAC address it sees on the port, and of the port-security maximum 1 command, which further instructs the switch that the address learned is the only address allowed on the port.
- All frames arriving at the switch with a destination address of 0000.00bb.bbb will be forwarded out on Fa0/16.

The switchport port-security mac-address sticky command can also be used in combination with the interface-range command to make every port on the switch behave in this fashion as shown below for a 24-port switch.

```
Switch(config)#interface range fastethernet0/1-24
Switch(config-if)#switchport port-security
Switch(config-if)#switchport port-security mac-address sticky
Switch(config-if)#switchport port-security maximum 1
```

The switchport port-security mac-address command is incorrect since this command requires an additional argument to be valid (either a statically configured MAC address or the sticky option).

The switchport port-security command activates port security on the switch port, but does not configure sticky MAC address learning.

The switchport port-security sticky mac-address and switchport mac-address sticky options are incorrect because these are not valid Cisco IOS commands.

Objective:

Infrastructure Security

Sub-Objective:

Configure, verify, and troubleshoot port security

References:

[Cisco > Catalyst 6500 Release 12.2SXH and Later Software Configuration Guide > Configuring Port Security > Enabling Port Security with Sticky MAC Addresses on a Port](#)
[Cisco > Cisco IOS Security Command Reference > show vtemplate through switchport port-security violation > switchport port-security mac-address](#)

QUESTION 123

Which of the following statements are true of Class C IP addresses?

- A. The decimal values of the first octet can range from 192 to 223
- B. The decimal values of the first octet can range from 1 to 126
- C. The first octet represents the entire network portion of the address
- D. The first three octets represent the entire network portion of the address
- E. The value of the first binary place in the first octet must be 0
- F. The value of the first two binary places in the first octet must be 11

Correct Answer: ADF

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A class C IP addresses will have the following characteristics:

- The decimal values of the first octet can range from 192 to 223
- The first three octets represent the entire network portion of the address
- The value of the first two binary place in the first octet must be 11

Class B IP addresses will have the following characteristics:

- The decimal values of the first octet can range from 128 to 191
- The first two octets represent the entire network portion of the address
- The value of the first two binary place in the first octet must be 10

Class A IP addresses will have the following characteristics:

- The decimal values of the first octet can range from 1 to 126
- The first octet represents the entire network portion of the address

- The value of the first binary place in the first octet must be 0

Objective:

Network Fundamentals

Sub-Objective:

Compare and contrast IPv4 address types

References:

[Cisco > IP Routing > IP Addressing and Subnetting for New Users](#)

QUESTION 124

Which Cisco Internetwork Operating System (IOS) command would be used to define a static route for network 192.168.11.0 through default gateway 192.168.43.1?

- A. router(config)# ip route 192.168.11.0 255.255.255.0 192.168.43.1
- B. router# ip route 192.168.11.0 255.255.255.0 192.168.43.1
- C. router(config)# ip classless 192.168.43.1
- D. router(config)# ip default gateway 192.168.11.0 255.255.255.0 192.168.43.1
- E. router# ip default gateway 192.168.43.1

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:


The router(config)# ip route 192.168.11.0 255.255.255.0 192.168.43.1 command would be used to define a static route for network 192.168.11.0 through default gateway 192.168.43.1. Static routing is used to manually configure routes to remote networks. The syntax of the ip route command is as follows:

ip route [destination_network] [mask] [next-hop_address or exit interface] [administrative_distance] [permanent]

The parameters of the command are as follows:

- destination_network: Defines the network that needs to be added in the routing table.
- mask: Defines the subnet mask used on the network.
- next-hop_address: Defines the default gateway or next hop router that receives and forwards the packets to the remote network.
- administrative_distance (AD): Static routes have an AD of 1, which can be changed to change the priority of the route.

Static routing is often implemented in small yet stable networks where the number of routes is small and manageable, and the network can benefit from the elimination of the traffic that dynamic routing protocols would introduce. If this is the case, it is important that all routes be statically created, or else networking problems can occur. For example, if in the diagram below no route to the 192.168.110.128/26 network on Router 2 exists on Router 1, Host 1 will be unable to ping Host 2. The fact that Host 1 would still be able to ping the S0/0 interface on Router 2 could obscure this missing route.



Host 1 will be able to ping the S0/0 interface of Router 2 because the 192.35.87.4/30 network will be in the routing table of Router 1, being directly connected to Router 1. Directly connected routes are automatically placed in the routing table. However, if you executed the show run command on Router 1, the output would indicate that no route to the 192.168.110.128/26 exists:

```

<output omitted>
interface Fa0/1
  ip address 192.168.54.65 255.255.255.192
  no shutdown
interface S0/0
  ip address 192.35.87.5 255.255.255.252
  no shutdown

```

The option router# ip route 192.168.11.0 255.255.255.0 192.168.43.1 is incorrect because the ip route command should be configured in the global configuration mode.

The option router(config)# ip classless 192.168.43.1 is incorrect because the ip classless global configuration mode command allows a router to accept and forward packets for subnets that are not directly connected. The packets are forwarded to the best available supernet route.

The option router(config) # ip default gateway 192.168.11.0 255.255.255.0 192.168.43.1 is incorrect because the ip default gateway command is used to define the default gateway address when IP routing is disabled in the network.


Objective:
 Routing Fundamentals
 Sub-Objective:
 Configure, verify, and troubleshoot IPv4 and IPv6 static routing

References:
[Cisco > Cisco ASDM User Guide, 6.1 > Configuring Dynamic And Static Routing > Field Information for Static Routes](#)
[Cisco > Support > IP > IP Routing > Design > Design TechNotes > Specifying a Next Hop IP Address for Static Routes > Document ID: 27082](#)
[Cisco > Cisco IOS IP Routing Protocols Command Reference > IP Routing Protocol-Independent Commands:](#)

[A through R > ip route](#)

QUESTION 125

View the following network diagram:



Which switch will become the root bridge?

- A. SwitchA
- B. SwitchB
- C. SwitchC
- D. The root bridge cannot be determined from the given information.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

SwitchA will become the root bridge. 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

The switch with the lowest bridge ID is selected 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, and 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.

Neither SwitchB nor SwitchC will become the root bridge. Although both have an equal priority value to SwitchA (32768), the MAC addresses of SwitchB and SwitchC are higher than that of SwitchA.

The root bridge can be determined with the information given. If the diagram did not indicate MAC addresses, then the root bridge would not be able to be determined, since the priorities are equal.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot STP protocols

References:

[Cisco Documentation > Cisco 7600 Series Cisco IOS Software Configuration Guide, 12.2SX > Configuring STP and IEEE 802.1s MST > Understanding the Bridge ID Support > Technology Support > LAN Switching > Spanning Tree Protocol > Configure > Configuration Examples and TechNotes > Understanding and Configuring Spanning Tree Protocol \(STP\) on Catalyst Switches](#)

QUESTION 126

Which protocol is used in redundant network topologies to avoid receiving multiple copies of the same frame?

- A. 802.1q
- B. Spanning Tree Protocol
- C. Cisco Discovery Protocol
- D. Routing Information Protocol

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Spanning Tree Protocol (STP) is used to remove switching loops in redundantly configured switched environments, and to create a single active Layer 2 path between any two network segments. This eliminates the chance of multiple copies of the same unicast frame being sent in the LAN. It also prevents broadcast packets from creating a broadcast storm when redundant connections exist between the switches. The benefits of STP include:

- Prevention of broadcast storms
- Prevention of multiple frame copies
- Media Access Control (MAC) address database stability

Whenever a network segment can be handled by more than one switch, STP will elect one switch to take responsibility, and the other switches will be placed into a blocking state for the ports connected to that segment. In this way, only one switch receives and forwards data for this segment, which removes the potential for multiple copies of the same frame being generated. For STP to provide this functionality it must be running on all of the switches. Therefore, a properly implemented redundant topology STP is required in order to prevent multiple copies of the same unicast frame from being transmitted.

802.1q is a frame tagging method for identifying Virtual LAN (VLAN) memberships over trunk links. Frame tagging ensures identification of individual VLAN frames over a trunk link carrying frames for multiple VLANs. This frame tagging method is a standardized protocol that was developed by The Institute of Electrical and Electronics Engineers (IEEE). Cisco has also developed a proprietary frame tagging method known as Inter-Switch Link (ISL). 802.1q does not mitigate loops or the reception of multiple copies of frames. The IEEE specification for STP is 802.1d.

Cisco Discovery Protocol is a Cisco proprietary protocol used to collect hardware and protocol information for directly connected Cisco devices. CDP has nothing to do with redundant network topologies.

Routing Information Protocol (RIP) is a distance vector routing protocol. It populates routing tables dynamically about the topology changes. However, RIP does not control the receipt of multiple copies of frames in redundant network topologies.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot STP protocols

References:

[Cisco > Support > Configuring Spanning Tree-Protocol > How STP Works](#)

QUESTION 127

Which trunk encapsulation defines one VLAN on each trunk as a native VLAN?

- A. ISL
- B. IEEE 802.1q
- C. IEEE 802.11a
- D. auto

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

IEEE 802.1q defines one VLAN on each trunk as the native VLAN.

The default value of a native VLAN is VLAN1. The IEEE 802.1q method does not encapsulate frames when forwarded over a trunk in a native VLAN; that is, IEEE 802.1q does not add its header information while transmitting frames in the native VLAN. This traffic is called untagged traffic. Frames originating from other VLANs, however, will have a 4-byte 802.1q header inserted into the frame to identify the VLAN number.

The native VLAN number can be changed if desired. If done it should be done on both ends of the connection. Otherwise, traffic that uses the native VLAN (untagged traffic) will not be able to cross the link. The command to change the native VLAN is

Switch(config)#switchport trunk native vlan vlan number

Inter Switch Link (ISL) does not define one VLAN on each trunk as a native VLAN. ISL is the Cisco proprietary trunk encapsulation, and it can only be used between two Cisco switches.

IEEE 802.11a is a wireless standard defined by the IEEE, and has nothing to do with VLANs.

Auto is not an encapsulation method. The auto trunking mode is a method for negotiating an encapsulation method over trunk links.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure and verify Layer 2 protocols

References:

QUESTION 128

Which command will display the Virtual LAN (VLAN) frame tagging method for a switch link?

- A. show vlan
- B. show vlan encapsulation
- C. show vtp status
- D. show interfaces trunk

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The show interfaces trunk command displays the list of trunk ports and the configured VLAN frame tagging methods.

Sample output of the show interfaces trunk command would be as follows:

```
SwitchB# show interfaces trunk
Port Mode Encapsulation Status Native vlan
Fa0/1 on 802.1q trunking 1
Fa0/2 on 802.1q trunking 1
Fa0/3 on 802.1q trunking 1
<<output omitted>>
```

The show vlan command displays the VLAN number, name, status, and ports assigned to individual VLANs. Although the command cannot be used to determine the frame tagging method used for each trunk, it can be used to determine which ports are trunk ports by the process of elimination.

In the output below, generated from a six-port switch, the missing port (Fa0/6) is a trunk port. For communication to be possible between the two VLANs configured on the switch, Fa0/6 must be connected to a router, and trunking must be configured on the router end as well. The command is also useful for verifying that a port has been assigned to the correct VLAN as it indicates in the VLAN column the VLAN to which each port belongs.

```
Switch# show vlan
```

Vlan name	Status	Ports
1 default	active	Fa0/1, Fa0/2, Fa0/3, Fa0/4
58 vlan 58	active	Fa0/5

The show vlan encapsulation command is not a valid command for Cisco switches.

The show vtp status command does not display VLAN frame tagging method. The command is used to verify the status of VTP. The output of the show vtp status command would be as follows:

```
SwitchB# show vtp status
VTP Version : 2
Configuration Revision : 1
Maximum VLANs supported locally : 64
Number of existing VLANs : 16
VTP Operating Mode : Client
VTP Domain Name : MARKETING
VTP Pruning Mode : Disabled
VTP V2 Mode : Disabled
VTP Traps Generation : Disabled
MD5 digest : 0x4D 0x60 0xA3 0x5E 0xC7 0x41 0x8C 0x47
```

Line 6 of the given output indicates that the switch is operating in VTP Client mode. There are three possible VTP modes in which a switch can operate: Server, Client, and Transparent.

- In Server mode, any changes made in the switch, such as adding a VLAN, will be recorded in the local database and also passed on to the other switches, where the change will be added.
- In Client mode, the switch will accept and record changes from switches in Server mode, but will not accept changes made on the local switch.
- In Transparent mode, the switch adds changes made locally to the database, but will not send or accept changes sent from other switches.

The mode in use could be a useful piece of information during troubleshooting. For example, if you were unsuccessfully attempting to add a VLAN to the database, the reason would be that the switch is in VTP Client mode. If you were adding a VLAN in Transparent mode, the VLAN would be added to the local database but fail to appear on the other switches. If the switch were in Transparent mode, Line 6 in the above output would appear as follows:

```
VTP Operating Mode: Transparent
```

Only switches operating in VTP Server mode can accept changes to the VLAN database. This situation could be corrected easily and a VLAN 50 could be successfully added at two different configuration prompts by executing the following commands:

At global configuration mode:

```
switchB# config t
switchB(config)# vtp mode server
switchB(config)# vlan 50
```

At VLAN configuration mode:

```
switchB# vlan database
switchB(vlan)# vtp server
switchB(vlan)# vlan 50
```

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot VLANs (normal/extended range) spanning multiple switches

References:

[Cisco Press Home > Articles > Cisco Certification > CCNA > CCNA Self-Study \(ICND Exam\): Extending Switched Networks with Virtual LANs](#)

QUESTION 129

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.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot EIGRP for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub)

References:

[Cisco > Support > IP > IP Routing > Technology Information > Technology White Paper > Enhanced Interior Gateway Routing Protocol > Document ID: 16406 > Feasible Distance, Reported Distance, and Feasible Successor](#)

QUESTION 130

Examine the output of the show ip route command below:

Gateway of last resort is not set

```
20.0.0.0/24 is subnetted, 1 subnets
O E2    20.20.20.0 [110/20] via 192.168.1.2, 00:05:10, FastEthernet0/0
O IA 172.16.0.0/16 [50/21] via 192.168.4.1, 00:05:10, FastEthernet0/1
          [50/21] via 192.168.1.2, 00:05:10, FastEthernet0/0
C    192.168.4.0/24 is directly connected, FastEthernet0/1
10.0.0.0/32 is subnetted, 1 subnets
C    10.10.10.10 is directly connected, Loopback0
C    192.168.1.0/24 is directly connected, FastEthernet0/0
O    192.168.2.0/24 [110/20] via 192.168.1.2, 00:05:10, FastEthernet0/0
O    192.168.3.0/24 [110/20] via 192.168.4.1, 00:05:10, FastEthernet0/1
30.0.0.0/32 is subnetted, 1 subnets
O    30.30.30.30 [50/21] via 192.168.4.1, 00:05:10, FastEthernet0/1
          [50/21] via 192.168.1.2, 00:05:10, FastEthernet0/0
```

Which of the following statements is FALSE?

- A. The route to 30.30.30.30 uses a cost of 21
- B. The command `ip route 192.168.2.0 255.255.255.0 172.16.14.2 200` will replace the current route to 192.168.2.0/24
- C. The route to 192.168.2.0/24 uses the default administrative distance
- D. Traffic will be load balanced across two routes to 30.30.30.30

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The command `ip route 10.10.10.0 255.255.255.0 172.16.14.2 200` will NOT replace the current route to 10.0.0.0/24.

When you execute the `ip route` command to enter a static route, the administrative distance can be altered by adding the desired distance value to the end of the command. In this scenario, the administrative distance value was set to 200. The route to the 10.10.10.0/24 network that is currently in the table was learned by OSPF and is using the default administrative distance of 110. Since 110 is lower than 200, the new static route will not be added to the routing table UNLESS the current route becomes unavailable.

The route to 30.30.30.30 does use a cost of 21, as is indicated by the value on the right side of the forward slash within the brackets found in the route entry, [50/21].

The route to 192.168.2.0/24 uses the default administrative distance. It was learned from OSPF, which has a default distance of 110. Its administrative distance is indicated by the value on the left side of the forward slash within the brackets found in the route entry, [110/20].

Traffic will be load balanced across two routes to 30.30.30.30 because they have equal cost of 21. This cost is indicated by the value on the right side of the forward slash within the brackets found in the route entry, [50/21].

Objective:

Routing Fundamentals

Sub-Objective:

Describe how a routing table is populated by different routing information sources

References:


[Cisco > Support > Cisco IOS IP Routing: Protocol-Independent Command Reference > ip route](#)

[Cisco Press > Articles > Cisco Network Technology > General Networking > Cisco Networking Academy's](#)

[Introduction to Routing Dynamically](#)

QUESTION 131

Examine the diagram below:



You attempt to make a Telnet connection from PC1 to the switch connected to the Web server, but the connection fails. After making a console connection to the switch connected to the Web server and executing the show run command, you see the following information:

<output omitted>

```
interface vlan 1
ip address 58.59.6.2 255.0.0.0
!
ip default gateway 192.168.5.1
!
line vty 04
password ajax
login
```

Which value is NOT correct?

- A. the default gateway
- B. the VLAN number
- C. the password
- D. the login command

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The switch is connected to the Fa0/0 interface on the router R1. The address of Fa0/0 should be the default gateway for the switch. This means it should be 58.59.6.4 rather than 192.168.5.1.

The VLAN number is correct. The IP address of a switch is set on the VLAN 1 interface of the switch.

The password can be anything you desire, so that is correct.

The login command is correct. This command instructs the switch to prompt for a password. Since there is a password configured, this will not prevent a connection to the switch.

Objective:

Infrastructure Management

Sub-Objective:

Configure and verify device management

References:

QUESTION 132

You have been asked to troubleshoot the NTP configuration of a router named R70. After executing the show run command, you receive the following partial output of the command that shows the configuration relevant to NTP:

```
clock timezone PST -8
clock summer-time PDT recurring
ntp update-calendar
ntp server 192.168.13.57
ntp server 192.168.11.58
interface Ethernet 0/0
ntp broadcast
```

Based on this output, which of the following statements is true?

- A. the time zone is set to 8 hours less than Pacific Standard time
- B. the router will listen for NTP broadcasts on interface E0/0
- C. the router will send NTP broadcasts on interface E0/0
- D. the router will periodically update its software clock

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The router will send NTP broadcast on its E0/0 interface. The command ntp broadcast, when executed under an interface, instructs the router to send NTP broadcast packets on the interface. Any devices on the network that are set with the ntp broadcast client command on any interface will be listening for these NTP broadcasts. While the clients will not respond in any way, they will use the information in the NTP broadcast packets to synchronize their clocks with the information.

The time zone is not set to 8 hours less than Pacific Standard Time. The value -8 in the command clock timezone PST -8 represents the number of hours of offset from UTC time, not from the time zone stated in the clock timezone command.

The router will not listen for NTP broadcasts on the interface E0/0. The ntp broadcast command, when executed under an interface, instructs the router to send NTP broadcast packets on the interface. To set the interface to listen and use NTP broadcasts, you would execute the ntp broadcast client command on the interface.

The router will not periodically update its software clock. The command ntp update-calendar configures the system to update its hardware clock from the software clock at periodic intervals.

Objective:

Infrastructure Services

Sub-Objective:

Configure and verify NTP operating in a client/server mode

References:

[Basic System Management > Setting Time and Calendar Services > Configuring NTP](#)

QUESTION 133

As you are training a new junior technician, the trainee is examining the routing table. He tells you that there are four different routes to the same network in different routing databases. He asks you which of the routes will be used to populate the routing table.

What will your answer be, assuming that all routing protocols are set at the default administrative distance?

- A. The route with an R next to it
- B. The route with an S next to it
- C. The route with a C next to it
- D. The route with an I next to it

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:


Explanation:

The route with a C next to it is a directly connected route and has an administrative distance of 0, which means it will be preferred over any routes with a larger value for administrative distance. Each routing protocol has a default administrative distance assigned. Administrative distance is used by the router to determine the preferred route when a route is learned from different routing protocols. This process can be manipulated by the administrator by using the distance command to alter the default assignments.

It is significant to note that routers with no static routes and no routing protocols enabled will populate all directly connected routes to the routing table with no action on the part of the administrator. Routes that are NOT directly connected will not be in the routing table unless one of two things occurs:

- A static route is created by the administrator
- A routing protocol is enabled that allows the router to learn about the network and its route from another router running the same routing protocol

For example, in the diagram below, R3 will have routes to the 192.168.3.0/24 ,192.168.1.0/24 and the 192.168.0.0/24 networks in its routing table by default. It will only have routes to the 192.168.2.0/24, 192.168.5.0/24, and 192.168.6.0/24 networks if a routing protocol is used or if an administrator creates static routes for each network.



When a packet is received by a router interface, the router de-encapsulates the frame or removes the layer two information (MAC data for Ethernet or DLCIs for frame relay) and then performs a lookup for the network ID of the network in which the destination IP address resides. When multiple routes exist, it will choose the one with the lowest administrative distance. The router only places the route with the lowest distance in the table.

The route with an R next to it is a route learned from Routing Information Protocol (RIP). It has a default administrative distance of 120, so it will not prefer over a directly connected route.

The route with an S next to it is a static route or one configured manually. It has an administrative distance of 1, so it will not be preferred over a directly connected route.

The route with an I next to it is a route learned from Internal Gateway Routing Protocol (IGRP). It has an administrative distance of 100, so it will not be preferred over a directly connected route.

Objective:

Routing Fundamentals

Sub-Objective:

Interpret the components of routing table

References:

[Cisco > Support > IP > IP Addressing Services > Design Technotes > What Is Administrative Distance? > Document ID: 15986](#)

QUESTION 134

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.

Objective:

Infrastructure Services

Sub-Objective:

Configure, verify, and troubleshoot basic HSRP

References:

[Cisco > Cisco IOS IP Application Services Command Reference > show standby](#)

[Cisco > Cisco IOS IP Application Services Configuration Guide, Release 12.4 > Part 1: First Hop Redundancy Protocols > Configuring HSRP](#)

QUESTION 135

Enhanced Interior Gateway Routing Protocol (EIGRP) uses which algorithm to select the best path to the destination?

- A. Diffusing Update Algorithm (DUAL)
- B. Dijkstra algorithm
- C. Bellman-Ford algorithm
- D. Shortest Path First (SPF) algorithm

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

EIGRP uses the Diffusing Update Algorithm (DUAL) to select the best path to the destination. EIGRP is a classless protocol that allows the use of variable length subnet masks (VLSM), and supports classless interdomain routing (CIDR) for the allocation of IP addresses.

EIGRP is characterized by these components:

- DUAL: EIGRP implements DUAL to select paths free of routing loops. DUAL selects the best path and the second best path to the destination. The terminology used in DUAL is as follows:
 - Successor: Best path selected by DUAL.
 - Feasible successor: Second best path selected by DUAL. This is a backup route stored in the topology table.
 - Feasible distance: The lowest calculated metric of a path to destination.
- Protocol-dependent modules: Different modules are used by EIGRP to independently support Internet Protocol (IP), Internetwork Packet Exchange (IPX), and AppleTalk routed protocols. These modules act as a logical interface between DUAL and routing protocols.
- Neighbor discovery and recovery: Neighbors are discovered and information about neighbors is maintained

by EIGRP. A hello packet is multicast on 224.0.0.10 every five seconds and the router builds a table with the information. EIGRP also enables proper operation over a Non-Broadcast Multiple Access (NBMA) point-to-multipoint network. EIGRP multicasts a hello packet every 60 seconds on the multipoint Wide Area Network (WAN) interfaces (X.25, frame relay, or Asynchronous Transfer Mode).

- Reliable Transport Protocol (RTP): RTP is used by EIGRP to manage EIGRP packets. Reliable and ordered delivery of route updates is ensured using RTP.

EIGRP updates about routes can contain five metrics: minimum bandwidth, delay, load, reliability, and maximum transmission unit (MTU). Of these five metrics, by default, only minimum bandwidth and delay are used to compute the best path.

The Dijkstra algorithm and Shortest Path First (SPF) algorithm are used by the Open Shortest Path First (OSPF) routing protocol for selecting the best path to the destination, not by EIGRP.

The Bellman-Ford algorithm is used by Routing Information Protocol (RIP).

Objective:

Routing Fundamentals

Sub-Objective:

Compare and contrast interior and exterior routing protocols

References:

[Cisco > Internetworking Technology Handbook > Enhanced Interior Gateway Routing Protocol \(EIGRP\)](#)

QUESTION 136

Examine the following output from SwitchD.

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

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

- A. switchport mode trunk
- B. switchport mode nonegotiate
- C. switchport trunk encapsulation 802.1q
- D. switchport trunk native VLAN

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The command switchport mode trunk must be executed for a trunk to form. The output indicates that the Administrative Mode of the port is "static access," which means the port has been configured as a static (fixed) access port. Access mode disables trunking on an access port.

Below is a sample of the configuration required to allow a router to provide inter-VLAN routing between two VLANs residing on the switch:

```
Router(config)#interface fa0/0
Router(config)#no shut down
Router(config)#interface fa0/0.1
Router(config-subif)#encapsulation dot1q
Router(config-subif)#ip address 192.168.10.1 255.255.255.0
Router(config)#interface fa0/0.2
Router(config-subif)#encapsulation dot1q
Router(config-subif)#ip address 192.168.20.1 255.255.255.0

Switch(config)#interface fa0/1
Switch(config-if)#switchport mode trunk
```

For this example, the following statements are true:

- The trunk link connects to Fa0/0 on the router and Fa0/1 on the switch.
- The physical interface F0/0 on the router has been divided into two subinterfaces, Fa0/0.1 and Fa0/0.2.
- The encapsulation type of 802.1q has been specified on the two subinterfaces of the router.
- The physical interface on the switch has been specified as a trunk link.
- The IP addresses 192.168.10.1 and 192.168.20.1 should be the default gateways of the computers located in VLANs 1 and 2, respectively.

The switchport mode nonegotiate command does not need to be executed because the switch is already configured for non-negotiation, as indicated by the output Negotiation of Trunking: Off. Trunk negotiation using the Dynamic Trunking Protocol (DTP) does not need to be enabled for a trunk to form.

The switchport trunk encapsulation 802.1q command does not need to be executed for a trunk to form. Also, the output Operational Trunking Encapsulation: dot1q indicates that 802.1q encapsulation is already configured.

The switchport trunk native VLAN command does not need to be executed. This command is used to change the native VLAN from its default of 1, but leaving it set to the default of 1 will not prevent the trunk from forming.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot inter-VLAN routing

References:

[Cisco > Catalyst 3560 Switch Software Configuration Guide, Rel. 12.2\(25\)SEE > Configuring VLANs > Configuring VLAN Trunks > Trunking Overview](#)

QUESTION 137

Which Cisco IOS command can be used to troubleshoot switch startup problems on a Cisco Catalyst 2950 switch?

- A. show test
- B. show diagnostic
- C. show post
- D. show switchstartup

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Cisco IOS command show post is used on the 2900/3500XL, 2950/2955, 3550, 2970, and 3750 series switches to view and troubleshoot issues related to the Power On Self-Test (POST) on the switch. This command will find the POST test that failed on startup.

The show test command is incorrect because it is a CatOS command, not a Cisco IOS command. The Cisco 2950 uses a Cisco IOS operating system and not the Catalyst operating system. The show test command is used on a switch to view any hardware errors that occurred at startup. It also provides information on the errors returned from the diagnostic tests. The following parameters can be used with this command:

- mod: An optional parameter used to specify the module number.
- diaglevel: Used to view the diagnostic level.
- diagfail-action: Used to view information on the action taken by the supervisor engine after the failure of a diagnostics test.

The following code is a sample output of this command for module 2:

```
Module 2 : 2-port 1000BaseX Supervisor
Network Management Processor (NMP) Status: (. = Pass, F = Fail, U = Unknown)
ROM: . Flash-EEPROM: . Ser-EEPROM: . NVRAM: . EOBC Comm: .
Line Card Firmware Status for Module 2 : PASS
Port Status :
Ports 1 2
-----

Line Card Diag Status for Module 2 (. = Pass, F = Fail, N = N/A)
Module 2
Cafe II Status :
NewLearnTest: .
IndexLearnTest: .
DontForwardTest: .
DontLearnTest: .
ConditionalLearnTest: .
BadBpduTest: .
TrapTest: .
Loopback Status [Reported by Module 2] :
Ports 1 2
-----

Channel Status :
Ports 1 2
-----
```

The show diagnostic command is incorrect because this command is used on the Catalyst 6000 series, not the 2950. A variant of the command, show diagnostics, is used for the Catalyst 4000 series. These commands can be used on the relevant switches to view any hardware errors that occurred on startup. This command displays the Power-On Self-Test (POST) results.

The show switchstartup command is not a valid Cisco IOS command.

Objective:

Infrastructure Management

Sub-Objective:

Use Cisco IOS tools to troubleshoot and resolve problems

References:

[Cisco>Home>Support>Product Support>End-of-Sale and End-of-Life Products>Cisco Catalyst 6000 Series Switches>Troubleshoot and Alerts> Troubleshooting TechNotes> Troubleshooting Switch Port and Interface Problems> Most Common Port and Interface Troubleshooting Commands for CatOS and Cisco IOS](#)

QUESTION 138

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.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot STP protocols

References:

[Cisco > Support > Configuring Spanning Tree-Protocol > How STP Works](#)

QUESTION 139

Which is the shortest possible notation of the following Internet Protocol version 6 (IPv6) address?

2001:0DB8:0000:0001:0000:0000:0000:F00D

- A. 2001:DB8::1::F00D
- B. 2001:DB8:0:1::F00D
- C. 2001:DB8:0:1:0:0:0:F00D
- D. 2001:0DB8:0:1::F00D

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The shortest possible notation of the IPv6 address 2001:0DB8:0000:0001:0000:0000:0000:F00D is 2001:DB8:0:1::F00D. The address is shortened according to the following rules:

- Remove leading zeros.
- Remove the consecutive fields of zeros with double colon (::).
- The double colon (::) can be used only once.

The option 2001:DB8::1::F00D is incorrect because the double colon (::) can be used only once in the process of shortening an IPv6 address.

The option 2001:DB8:0:1:0:0:0:F00D is incorrect because 2001:DB8:0:1:0:0:0:F00D can be further shortened to 2001:DB8:0:1::F00D.

The option 2001:0DB8:0:1::F00D is incorrect because 2001:0DB8:0:1::F00D can be further shortened to 2001:DB8:0:1::F00D.

Objective:

Network Fundamentals

Sub-Objective:

Compare and contrast IPv6 address types

References:

QUESTION 140

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

Objective:

WAN Technologies

Sub-Objective:

Describe WAN access connectivity options

References:

[Cisco > Support > Product Support > End-of-Sale and End-of-Life Products > Cisco IOS Software Releases 11.1 > Configure > Feature Guides > Clock Rate Command Enhancements Feature Module](#)

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

QUESTION 141

A device has an address of 192.168.144.21 and a mask of 255.255.255.240.

What will be the broadcast address for the subnet to which this device is attached?

- A. 192.168.144.23
- B. 192.168.144.28
- C. 192.168.144.31
- D. 192.168.144.32

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The broadcast address for the subnet to which this device is attached will be 192.168.144.31.

To determine the broadcast address of a network where a specific address resides, you must first determine the network ID of the subnetwork where the address resides. The network ID can be obtained by determining

the interval between subnet IDs. With a 28-bit mask, the decimal equivalent of the mask will be 255.255.255.240. The interval between subnets can be derived by subtracting the value of the last octet of the mask from 256. In this case, that operation would be 256 - 240. Therefore, the interval is 16.

The first network ID will always be the classful network you started with (in this case 192.168.144.0). Then each subnetwork ID in this network will fall at 16-bit intervals as follows:

192.168.144.0
192.168.144.16
192.168.144.32
192.168.144.48

At 192.168.144.48 we can stop, because the address that we are given as a guide is in the network with a subnet ID of 192.168.144.16. Therefore, since the broadcast address for this network will be 1 less than the next subnet ID (192.168.144.32), the broadcast address for the subnet to which this device is attached is 192.168.144.31.

All the other options are incorrect because none of these will be the broadcast address for the subnet to which this device is attached.

Objective:

Network Fundamentals

Sub-Objective:

Apply troubleshooting methodologies to resolve problems

References:

[Cisco > Technology Support > IP > IP Routing > Design TechNotes > IP Addressing and Subnetting for New Users > Document ID: 13788 > Understanding IP Addresses](#)

QUESTION 142

DRAG DROP

Click and drag the OSI layer, on the left, to the commands at which they test functionality. If a command can test more than one layer, choose the highest layer for which it can test. (It may be necessary to use an OSI layer multiple times.)

Select and Place:

Layers:

Layer 1
Layer 2
Layer 3
Layer 4
Layer 5
Layer 6
Layer 7

Command-Line Tool:

	ping
	show interface
	telnet
	show cdp neighbor
	ftp

Correct Answer:

Layers:

Layer 1
Layer 2
Layer 3
Layer 4
Layer 5
Layer 6
Layer 7

Command-Line Tool:

Layer 3	ping
Layer 2	show interface
Layer 7	telnet
Layer 2	show cdp neighbor
Layer 7	ftp

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Telnet operates at the application layer, which is Layer 7 of the OSI model. File transfer Protocol (FTP) is a generic command that is also used by some high-end Cisco routers but in a different format. FTP also operates at Layer 7. The ping command operates at the network layer, which is Layer 3 of OSI reference model. Therefore, it is used to test the connectivity up to Layer 3. The show interface command will display the status

of line protocol. If it displays the message interface up, line protocol up it means that Layer 2 is functioning correctly.

The show cdp neighbor command also operates at Layer 2, which is the data link layer.

Objective:

Infrastructure Management

Sub-Objective:

Perform device maintenance

References:

[Cisco Documentation > Internetworking Technology Handbook > Internetworking Basics > OSI model](#)

QUESTION 143

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 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.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot single area and multi-area OSPFv2 for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub, virtual-link, and LSAs)

References:

[Cisco > Support > IP > IP Routing > Technology Information > Technology White Paper > OSPF Design Guide > Document ID: 7039 > DR Election](#)

QUESTION 144

Which statement is TRUE of the CSMA/CD Ethernet media access method?

- A. It requires centralized monitoring and control.
- B. It is ideal for a switched network environment.
- C. It uses a back-off algorithm to calculate a random time value.
- D. Each station is allotted a time slot in which they can transmit data.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Carrier Sense Multiple Access - Collision Detection (CSMA/CD) Ethernet Media Access Control (MAC) method uses a back-off algorithm to calculate random times to transmit packets across a channel. When two stations start transmitting at same time, their signals will collide. The CSMA/CD method detects the collision and causes both stations to hold the retransmission for an amount of time determined by the back-off algorithm. This is done in an effort to ensure that the retransmitted frames do not collide.

CSMA/CD does not require centralized monitoring and control nor does it assign time slots to stations. Moreover, the CSMA/CD method is designed to work in non-switched environment. It is an alternative to a token-passing topology, in which each station waits in turn to receive a token that allows it to transmit data. With CSMA/CD, each station is capable of making the decision regarding when to transmit the data.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Describe and verify switching concepts

References:

[Cisco Documentation > Internetworking Technology Handbook > Ethernet Technologies](#)

QUESTION 145

Which of the following commands will enable a global IPv6 address based on the Modified EUI-64 format interface ID?

- A. ipv6 address 5000::2222:1/64
- B. ipv6 address autoconfig
- C. ipv6 address 2001:db8:2222:7272::72/64 link-local
- D. ipv6 enable

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

To configure the interface to create a global IPv6 address based on the Modified EUI-64 format interface ID, you must enable stateless autoconfiguration. In stateless autoconfiguration, the interface will receive the network prefix from the router advertisement (RA) and generate a full IPv6 address by spreading the 48-bit

MAC address of the interface across 64 bits to complete the address. This can all be done simply by executing the ipv6 address autoconfig command at the interface configuration prompt.

The command ipv6 address 5000::2222:1/64 is used to manually assign a full IPv6 address to the interface without using stateless autoconfiguration or the eui-64 keyword to manually specify the first 64 bits and allow the last 64 bits to be generated from the MAC address of the interface.

The command ipv6 address 2001:db8:2222:7272::72/64 link local is used to configure a link-local address manually without allowing the system to generate one from the MAC address, which is the default method.

The command ipv6 enable is used to allow the system to generate a link-local address from the MAC address. Because this is the default behavior, the command is not required if any other ipv6 commands have been issued. Regardless of how many manual IPv6 addresses you configure, a link local address is always generated by default.

Objective:

Network Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot IPv6 addressing

References:

[Cisco > Product Support > Security > Cisco ASA 5500-X Series Firewalls > Configure > Configuration Guides > Cisco Security Appliance Command Line Configuration Guide, Version 7.2 > Chapter: Configuring IPv6 > Enforcing the Use of Modified EUI-64 Interface IDs in IPv6 Addresses](#)
[Cisco > Support > Cisco IOS IPv6 Command Reference > ipv6 address](#)

QUESTION 146

Refer to the following partial output of the show interfaces command:

```
Serial 0 is down, 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 runts, 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 are the two troubleshooting steps that you should perform to resolve the problem depicted in the output? (Choose two.)

- A. Check the cable connections.
- B. Reset the equipment.
- C. Check the router configuration.
- D. Check the router configuration for the shutdown interface command.

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should check the cable connections and reset the equipment to troubleshoot the problem depicted in the output. The Serial 0 is down, line protocol is down message indicates that there is no carrier detect (CD) signal sensed by the router. This problem might be due to incorrect cabling or a possible hardware failure.

A complete list of the possible troubleshooting steps that should be performed to resolve this issue include:

- Checking the cable connections.
- Resetting the equipment.
- Checking the CD LED on the CSU/DSU.
- Reporting the issue to the leased-line provider.
- Replacing the faulty equipment.

The router configuration is not a possible issue in this scenario because both serial 0 and line protocol are down, indicating a problem in the physical layer. Configuration issues, such as an incorrect IP address, would be indicated in the second section of the output (line protocol is up/down). The second section, regardless of whether it says up or down is meaningless when the first section indicates a problem.

You should not check the router configuration for the shutdown interface command. When an interface has been manually shut down with this command, it will be indicated in the output as Serial 0 is administratively down, line protocol is down.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Troubleshoot interface and cable issues (collisions, errors, duplex, speed)

References:

[Cisco > Cisco IOS Interface and Hardware Component Command Reference > show interfaces summary](#)
[Cisco Documentation > Cisco 1700 Series Router Software Configuration Guide > Configuring a Leased Line > Troubleshooting Problems with Leased Lines](#)

QUESTION 147

When executed on a HSRP group member named Router 10, what effect does the following command have?

```
Router10(config-if)# standby group 1 track serial0 25
```

- A. It will cause the router to increase its HSRP priority by 25 if the Serial0 interface on the standby router goes down
- B. It will cause the router to shut down the Serial0 interface if 25 packets have been dropped
- C. It will cause the router to notify Router 25 if serial 0 goes down
- D. It will cause the router to decrement its HSRP priority by 25 if Serial 0 goes down

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

This command will cause the router to decrement its HSRP priority by 25 if Serial 0 goes down. Interface tracking can be configured in Hot Standby Routing Protocol (HSRP) groups to switch traffic to the standby router if an interface goes down on the active router. This is accomplished by having the active router track its interface. If that interface goes down, the router will decrement its HSRP priority by the value configured in the command. When properly configured, this will cause the standby router to have a higher HSRP priority, allowing it to become the active router and to begin serving traffic.

When the standby router in an HSRP group is not taking over the active role when the active router loses its

tracked interface, it is usually a misconfigured decrement value, such that the value does not lower the HSRP priority of the active router far enough for the standby to have a superior priority value.

The command will not cause the router to increase its HSRP priority by 25 if the Serial0 interface on the standby router goes down. HSRP routers track their own interfaces, not those of another router.

The command will not cause the router to shut down the Serial0 interface if 25 packets have been dropped. It will only do this if the link becomes unavailable.

The command will not cause the router to notify Router 25 is serial 0 goes down. The number 25 in the command is the decrement value, not the ID of another router.

Objective:

Infrastructure Services

Sub-Objective:

Configure, verify, and troubleshoot basic HSRP

References:

[Cisco > Home > Support > Technology Support > IP > IP Application Services > Design > Design Technotes > How to Use the standby preempt and standby track Commands](#)

[Cisco > Cisco IOS IP Application Services Command Reference > standby track](#)

QUESTION 148

You are a network administrator for your organization. Your organization has two Virtual LANs, 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, whereas switches B, D, and E have user machines connected to the Production VLAN only. (Click the Exhibit(s) button to view the network diagram.)

To meet a new requirement, Marketing VLAN users must communicate with Production VLAN users and vice versa. What changes would be required for the network in this scenario?

- A. Disable VTP pruning.
- B. Convert all switch ports into trunk ports.
- C. Create an access list with permit statements.
- D. Install a routing device or enable Layer 3 routing on a switch.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

In this scenario, either a Layer 3 device or Layer 3 routing on a switch would be required to implement inter-VLAN routing. Although you could use multiple physical interfaces for the VLAN traffic, using trunk links between the switches and an external router would make more efficient use of the physical interfaces that you have. Only trunk links can carry traffic from multiple VLANs. These data frames must be frame tagged over the trunk link to identify the VLAN that sourced the frame. The receiving switch sees the VLAN ID, and uses this information to forward the frame appropriately. Additionally, the cables used to connect the router to the switches must be a straight-through cable and not a crossover cable.

When trunks links do not appear to be operating, it is always a good idea to make sure the port used for the trunk link is set as a trunk link and not as an access link. For example, the output below of the show interface fastethernet 0/15 switchport command indicates that Switch2 will not trunk because the port is set as an access link. This is shown in line 5 of the output:

<<output omitted>>

```
Switch2#show Interface fastethernet 0/15 switchport
```

```
Name: Fa0/15
```

```
SwitchportEnabled
```

```
Administrative Mode: access
```


Operational Mode: access
<<output omitted>>

The VLAN Trunking Protocol (VTP) pruning feature restricts unnecessary broadcast traffic between multiple switches. It does not affect inter-VLAN traffic. Therefore, disabling VTP pruning will not permit inter-VLAN communication between the Marketing and Production VLANs.

Converting all switch ports into trunk ports will permit traffic from multiple VLANs to traverse over these links. However, traffic from one VLAN will be restricted to that VLAN only, and inter-VLAN communication will not be possible.

Access lists can permit or deny packets based on the packets' source/destination IP address, protocol, or port number. However, access lists can manipulate inter-VLAN traffic only when inter-VLAN traffic is enabled using a Layer 3 device or Layer 3 routing. Therefore, creating access lists will not enable inter-VLAN routing between the Marketing and Production VLANs.

Objective:

Network Fundamentals

Sub-Objective:

Describe the impact of infrastructure components in an enterprise network

References:

QUESTION 149

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.

Objective:

Infrastructure Services

Sub-Objective:

Configure, verify, and troubleshoot basic HSRP

References:

[Cisco IOS Master Command List, Release 12.4T>show ip route profile through sshow mpls atm-ldp summary>Cisco IOS IP Application Services Command Reference>show standby through show udp>show standby](#)

QUESTION 150

DRAG DROP

Match the Dynamic Trunking Protocol (DTP) configuration on the switch ports so that a trunk link can be established. (Click and drag the DTP modes on the left and place them with their corresponding port on the right.)

Select and Place:

Modes:
Nonegotiate
Trunk
Desirable
Auto

Ports:	
	Trunk or Desirable or Auto
	Trunk or Desirable or Auto
	Trunk or Desirable
	Nonegotiate

Correct Answer:

Modes:

Ports:	
Trunk	Trunk or Desirable or Auto
Desirable	Trunk or Desirable or Auto
Auto	Trunk or Desirable
Nonegotiate	Nonegotiate

Section: (none)

Explanation

Explanation/Reference:

Explanation:

There are five DTP modes:

- Trunk: Switch will establish trunk if other end port is configured as Trunk/Desirable/Auto.
- Dynamic Desirable: Switch will establish trunk if other end port is configured as Trunk/Desirable/Auto.
- Dynamic Auto: Switch will establish trunk if other end port is configured as Trunk/Desirable.

- Nonegotiate: Other end port should also be configured with Nonegotiate, or should be a device that does not support DTP.
- Access: No trunk establishment.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure and verify Layer 2 protocols

References:

QUESTION 151

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.

Objective:

Routing Fundamentals

Sub-Objective:

Troubleshoot basic Layer 3 end-to-end connectivity issues

References:

[Cisco > Technology Support > IP > IP Routing > Design Technotes > IP Addressing and Subnetting for New Users > Document ID: 13788 > Understanding IP Addresses](#)

[Cisco > Technology Support > IP > IP Addressing Services > Design Technotes > Subnet Zero and the All-Ones Subnet > Document ID: 13711](#)

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

QUESTION 152

Examine the following partial output of the show interfaces command.

```
Router# show interfaces ethernet 0/0
Ethernet0/0 is administratively down, line protocol is down
Hardware is AmdP2, address is 0003.e39b.9220 (bia 0003.e39b.9220)
Internet address is 10.1.0.254/16
MTU 1500 bytes, BW 10000 Kbit, DLY 1000 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
ARP type: ARPA, ARP Timeout 04:00:00
<<output omitted>>
```

Which of the following statements are true? (Choose all that apply.)

- A. the interface is functional
- B. the largest frame allowed through this connection is 1500 bytes
- C. the interface needs the no shutdown command executed to be functional
- D. the largest frame allowed through this connection is 10000 Kbs

Correct Answer: BC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

From this output, we can determine that the largest frame allowed through this connection is 1500 bytes and that the interface needs the no shutdown command executed to be functional. The portions of the output that tell us this are:

MTU 1500 bytes indicates that the Maximum Transmission Unit (MTU) is 1500 bytes. The MTU is the largest frame size allowed.

Ethernet0/0 is administratively down indicates that the interface has either been disabled or has never been enabled. The command no shutdown is used to enable an interface, and until enabled, it will not function.

The interface is not functional, as indicated by the Ethernet0/0 is administratively down portion of the output.

The largest frame allowed through this connection is not 10000 Kbs. It is 1500 bytes. It is interesting to note that the bandwidth of the connection is 10000 Kbs, as indicated by the section:

BW 10000 Kbit

Objective:

LAN Switching Fundamentals

Sub-Objective:

Troubleshoot interface and cable issues (collisions, errors, duplex, speed)

References:

QUESTION 153

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.

Objective:

Routing Fundamentals

Sub-Objective:

Compare and contrast distance vector and link-state routing protocols

References:

[Cisco > Articles > Network Technology > General Networking > Dynamic Routing Protocols](#)

QUESTION 154

Multiple routes to a destination already exist from various routing protocols.

Which of the following values is used FIRST to select the route that is inserted into the route table?

- A. composite metric
- B. administrative distance
- C. prefix length
- D. hop count

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

When multiple routes to a destination exist from various routing protocols, the first value to be evaluated is the administrative distance of the source of the route. The following are examples of default administrative distance values:

Connected	0
Static	1
eBGP	20
EIGRP (internal)	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
EIGRP (external)	170
iBGP	200
EIGRP summary route	5

The second value to be compared is the composite metric, or any metric value for that matter. It is only used when multiple routes exist that have the same administrative distance.

The prefix length is only used to compare two existing routes in the routing table that lead to the destination, yet have different mask or prefix lengths. In that case, the route with the longest prefix length will be chosen.

Hop count is ONLY used when comparing multiple RIP routes. It is not the first consideration when multiple routes from various routing protocols exist in a routing table.

Objective:

Routing Fundamentals

Sub-Objective:

Describe how a routing table is populated by different routing information sources

References:

[Support > Technology Support > IP > IP Routing > Troubleshoot and Alerts > Configuration Example and TechNotes > Route Selection in Cisco Routers](#)

QUESTION 155

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.

Objective:

Infrastructure Management

Sub-Objective:

Configure and verify device-monitoring protocols

References:

[Configuring SNMP Support > Understanding SNMP > SNMP Versions](#)

[Cisco IOS Network Management Command Reference > snmp-server engineID local through snmp trap link-status > snmp-server host](#)

QUESTION 156

You need to manually assign IPv6 addresses to the interfaces on an IPv6-enabled router. While assigning addresses, you need to ensure that the addresses participate in neighbor discovery and in stateless auto-configuration process on a physical link.

Which of the following addresses can be assigned to the interfaces?

- A. FEC0:0:0:1::1/64
- B. FE80::260:3EFF:FE11:6770/10
- C. 2001:0410:0:1:0:0:0:1/64
- D. 2002:500E:2301:1:20D:BDFF:FE99:F559/64

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The FE80::260:3EFF:FE11:6770/10 address can be assigned to an interface of the IPv6-enabled router. This address is a link-local address as it has the prefix FE80::/10. Link-local addresses can be configured for an interface either automatically or manually.

Link-local addresses are IPv6 unicast addresses that are configured on the interfaces of an IPv6-enabled router. With link-local addresses, the nodes can connect to a network (local link) and communicate with other nodes. In addition, these addresses participate in the neighbor discovery protocol and the stateless auto-configuration process.

The FEC0:0:0:1::1/64 address should not be used for the interfaces because this address is a site-local address. Site-local addresses are IPv6 equivalent addresses to IPv4's private address classes. These addresses are available only within a site or an intranet, which typically is made of several network links.

You should not use the 2001:0410:0:1:0:0:0:1/64 and 2002:500E:2301:1:20D:BDFF:FE99:F559 addresses for the interfaces. These two addresses are global unicast addresses as they fall in the range from 2000::/3 and to E000::/3. A global address is used on links that connect organizations to the Internet service providers (ISPs).

Objective:

Network Fundamentals

Sub-Objective:

Configure and verify IPv6 Stateless Address Auto Configuration

References:

[Cisco > Understanding IPv6 Link Local Address](#)

QUESTION 157

Which Cisco Internetwork Operating System (IOS) command is used to make the running configuration in Random Access Memory (RAM) to the configuration the router will use at startup?

- A. copy running-config startup-config
- B. copy flash running-config
- C. copy tftp flash
- D. copy running-config flash memory
- E. copy startup-config tftp
- F. copy tftp running-config
- G. copy running-config tftp

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The copy running-config startup-config command is used to make the running configuration in Random Access Memory (RAM) the configuration the router will use at startup. It saves the running configuration in RAM to the router's NVRAM. This command should always follow changes to the configuration; otherwise, the changes will be lost at the next router restart. The startup configuration loads into memory from NVRAM at boot and resides in memory. When the router restarts, memory information is lost.

The copy flash running-config command is incorrect because this would copy a configuration from the router's flash memory to the running configuration, causing it to be the active configuration. While this can be done, it is not a common practice. Configuration files are normally stored in NVRAM.

The copy tftp flash command is incorrect because this command is used to replace the IOS image with a backup IOS image stored on a TFTP server to the target router. A router can also act as a TFTP server for another router. When you execute this command, you will be prompted for the IP address or hostname of the TFTP server. This prompt will display as in this example:

```
router#enable
router#copy tftp flash
Address or name of remote host []? 192.168.1.5.2
```

Before performing an upgrade of the IOS version from a TFTP server, you should verify that the upgrade is necessary by verifying the current IOS version number. The IOS version number can be found in the output of the following commands:

- **show running-config**
- **show version**
- **show flash**

The copy running-config flash memory command is incorrect because this command would copy the running configuration to the router's flash memory. It is the opposite of the copy flash-running config command. While this can be done, it is not a common practice. Flash is typically used to store the Cisco IOS or operating system. Configuration files are normally stored in NVRAM.

The copy startup-config tftp command is incorrect because this command would be used to copy the current configuration stored in NVRAM to a TFTP server. When you execute this command, you will be prompted for the IP address or hostname of the TFTP server. This prompt will display as below:

```
router#copy start tftp
Address or name of remote host []? 192.168.1.5
```

Destination filename [router-config]?

The address 192.168.1.5 is the address of the TFTP server. If no file name is given, it will save the file as router-config.

The copy tftp running-config is incorrect. This command is used to merge a backup configuration located on a TFTP server with the configuration in RAM.

The copy running-config tftp command is incorrect. It is used to make a backup copy of the configuration residing in RAM to a TFTP server.

Objective:

Infrastructure Management

Sub-Objective:

Perform device maintenance

References:

[Cisco > Tech Notes > How To Copy a System Image from One Device to Another > Document ID: 15092](#)
[Cisco Documentation > Cisco IOS Release 12.4 Command References > Using Cisco IOS Software for Release 12.4 > Understanding Command Modes](#)

QUESTION 158

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.

Objective:

Network Fundamentals

Sub-Objective:

Describe the effects of cloud resources on enterprise network architecture

References:

[Cloud and Systems Management Benefits](#)

QUESTION 159

You recently implemented SNMPv3 to increase the security of your network management system. A partial output of the show run command displays the following output that relates to SNMP:

<output omitted>

```
snmp-server group TECHS v3 noauth read TECHS write TECHS
```

Which of the following statements is true of this configuration?

- A. It provides encryption, but it does not provide authentication
- B. It provides neither authentication nor encryption
- C. It provides authentication, but it does not provide encryption
- D. It provides both authentication and encryption

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

It provides neither authentication nor encryption. In SMNPv3, 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

In this case, the keyword noauth in the configuration indicates that no authentication and no encryption are provided. This makes the implementation no more secure than SNMPv1 or SNMPv2.

In SNMPv1 and SNMPv2, authentication is performed using a community string. When you implement SNMP using the noauth keyword, it does not use community strings for authentication. Instead it uses the configured user or group name (in this case TECHS). Regardless, it does not provide either authentication or encryption.

Objective:

Infrastructure Management

Sub-Objective:

Configure and verify device-monitoring protocols

References:

[SNMP Configuration Guide, Cisco IOS XE Release 3SE \(Catalyst 3850 Switches\) > SNMPv3](#)

QUESTION 160

You are the network administrator for your company. You want to upgrade the network, which is currently running on IPv4, to a fully functional IPv6 network. During the transition, you want to ensure that hosts capable only of IPv6 can communicate with hosts capable only of IPv4 on the network.

Which solution should you implement to accomplish the task in this scenario?

- A. IPv6 over IPv4 tunnels
- B. IPv6 over dedicated Wide Area Network (WAN) links
- C. Dual-Stack Backbones
- D. Protocol translation

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The protocol translation deployment model should be used to accomplish the task in this scenario. It is the only offered solution that does not require at least one end of the communication solution to support both IPv6 and IPv4.

The following deployment models are available for IPv4 to IPv6 migration:

- IPv6 over IPv4 tunnels: IPv6 traffic is encapsulated into IPv4 packets. Then these packets are transferred over IPv4 WAN. This model eliminates the need to create separate circuits to connect to the IPv6 networks. This model increases protocol overhead because of the IPv6 headers and requires the edge router at each end be capable of both protocols.
- Protocol translation: A method allowing an IPv6 host to communicate with an IPv4 host. This is accomplished with the help of Network Address Translation - Protocol Translation (NAT-PT) used to configure translation between IPv6 and IPv4 hosts. NAT-PT allows communication between IPv6 hosts and applications, and native IPv4 hosts and applications.
- IPv6 over dedicated WAN links: A new deployment of IPv6 is created. In this model, IPv6 hierarchy, addressing, and protocols are used by all nodes. However, this model involves cost for creating IPv6 WAN circuits. This solution is not designed for LAN translation but rather communication over WAN links.
- Dual-Stack Backbones: A hybrid model in which backbone routers have dual-stack functionality, which enables them to route both IPv4 and IPv6 packets. It is suitable for an enterprise that uses both IPv4 and IPv6 applications. Running IPv6 and IPv4 together in a network is known as dual-stack routing.

Objective:

Network Fundamentals

Sub-Objective:

Identify the appropriate IPv6 addressing scheme to satisfy addressing requirements in a LAN/WAN environment

References:

[Cisco > Products and Services > Security > Cisco IOS Network Address Translation \(NAT\) > Data Sheets and Literature > White Papers > Network Address Translator-Protocol Translator](#)

QUESTION 161

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.

Objective:

WAN Technologies

Sub-Objective:

Describe WAN access connectivity options

References:

[Cisco Documentation > Internetworking Technology Handbook > Introduction to WAN Technologies > Circuit Switching](#)

QUESTION 162

You are configuring the link between a Cisco 2950 series switch and a Cisco 2611 router. You have physically connected the router's Ethernet port to the switch using a straight-through cable. The switch has not been configured, except for a hostname. The router's hostname has also been configured, and the Ethernet port has been enabled. However, you forgot to assign an IP address to the Ethernet port.

You issue the show cdp neighbors command and get the following output:

```
RouterA#show cdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater
Device ID         Local Interface         Holdtime    Capability Platform Port
ID
SwitchA          Eth 0/0                   157         S          2950       Fas 0/0
```

If you did not configure IP addresses, how is this information being passed between the two devices?

- A. The devices established a connection using default IP addresses.
- B. The ip unnumbered command has been issued, which means the interface does not require an IP address to be configured.
- C. CDP is a Layer 2 protocol and does not require IP addresses to be configured.
- D. CDP uses its own IP addressing system.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

CDP is a Layer 2 protocol and does not require IP addresses to be configured. The structure of the OSI model requires that the upper-layer protocols rely on the lower-layer protocols for operation. Protocols at Layer 3 cannot be operational unless Layers 1 and 2 are operational. Conversely, lower-layer protocols do not rely on upper-layer protocols for their operation. Because CDP operates at Layer 2 of the OSI model, it does not require an IP address to be active, since IP addresses are a function of Layer 3.

The ip unnumbered command has not been issued in this scenario. This command can only be used on serial interfaces, not Ethernet interfaces. It allows a serial interface to use an address that is already applied to an Ethernet interface.

Information is not being passed between the devices through default IP addresses. There is no such thing as default IP addresses on Ethernet interfaces for Cisco routers.

Information is not being passed between the devices through CDP's IP addressing system. CDP does not have its own IP addressing system because it does not use IP addresses for its operation.

Objective:

Infrastructure Management

Sub-Objective:

Use Cisco IOS tools to troubleshoot and resolve problems

References:

[Cisco > Cisco IOS Network Management Command Reference > schema through show event manager session cli username > show cdp neighbors](#)

QUESTION 163

You are the network administrator for your company and have configured Cisco Discovery Protocol (CDP) in your network. You recently noticed that when devices send large numbers of CDP neighbor announcements, some devices are crashing. You decide to disable CDP on the router.

Which command should you use to achieve the objective?

- A. no cdp run
- B. set cdp disable
- C. no cdp enable
- D. no cdp advertise-v2

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should use the no cdp run command to disable CDP on the router. Due to a known vulnerability regarding the handling of CDP by Cisco routers and switches when devices send large numbers of CDP neighbor announcements, some devices can crash or cause abnormal system behavior. To overcome this problem, you can disable CDP for the entire router by using the no cdp run command.

You cannot use the set cdp disable command to disable CDP on the router. This command disables CDP on an entire Catalyst switch.

You cannot use the no cdp enable command to disable CDP on the router. This command disables CDP on a specific interface.

You cannot use the no cdp advertise-v2 command to disable CDP on the router. This command disables CDPv2 advertisements.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure and verify Layer 2 protocols

References:

[Cisco > Support > Using Cisco Discovery Protocol](#)

[Cisco > Support > Technology Support > Network Management > Cisco's Response to the CDP Issue > Document ID: 13621](#)

QUESTION 164

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.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot single area and multi-area OSPFv2 for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub, virtual-link, and LSAs)

References:

[Internetworking Technology Handbook > Open Shortest Path First \(OSPF\)](#)

QUESTION 165

Which Cisco Internetwork Operating System (IOS) command is used to assign a router a name for identification?

- A. description
- B. banner motd
- C. hostname
- D. banner exec

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The hostname command is used to assign the router a name for identification. This command is a global configuration mode command. The syntax of the command is as follows:

Router(config)# hostname [name]

The name parameter of the command specifies the new host name for the router.

The description command is incorrect because this command is used to set a description for an interface. The description command is an interface configuration mode command.

The banner motd command is used to specify a message of the day (MOTD) banner to users logging into the router. This is a global configuration mode command, but it does not assign a name to the router for identification.

The banner exec command enables a banner message to be displayed when an EXEC process is created; for example, if a line is activated or an incoming connection is made to a telnet line.

Objective:

Network Fundamentals

Sub-Objective:

Select the appropriate cabling type based on implementation requirements

References:

[Cisco > Cisco IOS Configuration Fundamentals Command Reference > F through K > hostname](#)

QUESTION 166

Which command is used to disable Cisco Discovery Protocol (CDP) on a Cisco router?

- A. disable cdp
- B. no cdp run
- C. no cdp enable
- D. no cdp advertise-v2

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The no cdp run command is used to disable CDP on a Cisco router globally. CDP is a Layer 2 (Data Link layer) protocol that discovers information about neighboring network devices. CDP does not use network layer protocols to transmit information because it operates at the Data Link layer. Therefore, it is useful to determine information about directly connected Cisco network devices, because it can operate when network protocols have not been configured or are misconfigured. The show cdp neighbors detail command is used to view the IP addresses of the directly connected Cisco devices.

The no cdp advertise-v2 command disables CDPv2 advertisements. It will not disable the protocol globally.

The no cdp enable command is used to disable CDP on an interface. In a situation where CDP needs to be disabled on a single interface only, such as the interface leading to the Internet, this command would be executed from interface configuration mode for that specific interface. It will not disable the protocol globally. For example, to disable CDP for only the serial0 interface, the command sequence would be:

Router#configure terminal

Router(config)#interface serial 0

Router(config-if)#no cdp enable

The disable cdp command is not a valid Cisco command.

Objective:
LAN Switching Fundamentals
Sub-Objective:
Configure and verify Layer 2 protocols

References:
[Cisco > Cisco IOS Network Management Command Reference > show cdp neighbors](#)

QUESTION 167

You wish to configure Secure Shell (SSH) support on your router so that incoming VTY connections are secure.

Which of the following commands must be configured? (Choose all that apply.)

- A. ip domain-name
- B. transport input ssh
- C. ip access-group
- D. crypto key generate rsa
- E. service config

Correct Answer: ABD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Secure Shell (SSH) provides a secure alternative to Telnet for remote management of a Cisco device.

Configuring Secure Shell (SSH) support on a Cisco router involves a minimum of three commands:

- ip domain-name [domain-name]: configures the DNS of the router (global configuration mode)
- crypto key generate rsa: generates a cryptographic key to be used with SSH (global configuration mode)
- transport input ssh: allows SSH connections on the router's VTY lines (VTY line configuration mode)

The transport input ssh command allows only SSH connectivity to the router, and prevents clear-text Telnet connections. To enable both SSH and Telnet, you would use the transport input ssh telnet command.

The ip access-group command is incorrect because this command is used to activate an access control list (ACL) on an interface, and does not pertain to SSH.

The service config command is incorrect because this command is used to automatically configure routers from a network server, and does not pertain to SSH.

Objective:
Infrastructure Security
Sub-Objective:
Configure, verify, and troubleshoot basic device hardening

References:
[Cisco > Support > Technology Support > Security and VPN > Secure Shell \(SSH\) > Design > Configuring Secure Shell on Routers and Switches Running Cisco IOS > Document ID: 4145](#)

QUESTION 168

Which command would be used to establish static translation between an inside local address and an inside global address?

- A. Router(config)# ip nat inside source static local-ip global-ip
- B. Router(config)# ip source nat inside static local-ip global-ip
- C. Router(config)# ip nat inside static source local-ip global-ip
- D. Router(config)# ip nat static inside source local-ip global-ip

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should use the following command:

Router(config)# ip nat inside source static local-ip global-ip

You can translate your own IP addresses into globally unique IP addresses when communicating outside of your network. You can configure static or dynamic inside source translation.

To establish static translation between an inside local address and an inside global address, you should use the ip nat inside source static local-ip global-ip command. This static configuration can be removed by entering the no ip nat inside source static global command.

The other options are incorrect as they are not valid Cisco IOS configuration commands.

Objective:

Infrastructure Services


Sub-Objective:

Configure, verify, and troubleshoot inside source NAT

References:

QUESTION 169

Assume that all ports on Layer 2 devices are in the same Virtual LAN (VLAN). View the given network topology. (Click the Exhibit(s) button.)



Which network device should be placed at the highlighted box to produce a total of two broadcast domains and seven collision domains in the network?

- A. Hub
- B. Bridge

- C. Switch
- D. Router

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A hub should be placed at the highlighted box to produce a total of two broadcast domains and seven collision domains in the network. Network devices segment collision domains and broadcast domains in the following manner:

- Hub: A Layer 1 device with all ports in same collision domain and broadcast domain.
- Bridge/Switch: Layer 2 devices on which all ports are in different collision domains, but in the same broadcast domain (assuming that all ports are in the same VLAN or no VLAN is configured).
- Routers: A Layer 3 device on which every port is a separate collision as well as broadcast domain.

The bridge shown in the graphic has three ports populated by active links, resulting in three collision domains.

The switch shown in the exhibit has four ports populated with the links, resulting in four collision domains.

Together these two devices create seven collision domains.

Because the scenario requires that there be no more than seven collision domains, the device in the highlighted box must not create any further collision domains. A hub is a device that has all its ports in the same collision domain and will not create any further collision domains in the topology.

A bridge or switch cannot be the correct option because these will also add collision domains.

In the exhibit, the router has two ports with active links, which will result into two broadcast domains. Because the scenario states there are no more than two broadcast domains, the device in the highlighted box must not be a router. Routers are used to segment broadcast domains.

Objective:

Network Fundamentals

Sub-Objective:

Describe the impact of infrastructure components in an enterprise network

References:

QUESTION 170

Which Cisco IOS command is used on a Catalyst 2950 series switch to verify the port security configuration of a switch port?

- A. show interfaces port-security
- B. show port-security interface
- C. show ip interface
- D. show interfaces switchport

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The show port-security interface command displays the current port security and status of a switch port, as in this sample output:

```
Switch# show port-security interface fastethernet0/1
Port Security: Enabled
Port status: SecureUp
Violation mode: Shutdown
Maximum MAC Addresses: 2
Total MAC Addresses: 2
Configured MAC Addresses: 2
Aging Time: 30 mins
Aging Type: Inactivity
SecureStatic address aging: Enabled
Security Violation count: 0
```

The sample output indicates that port security has been enabled on interface FastEthernet0/1, and that a maximum of two MAC addresses has been configured. A violation policy of Shutdown indicates that if a third MAC address attempts to make a connection, the switch port will be disabled.

The violation mode setting has three possible values that take the following actions when a violation occurs:

- protect Drops packets with unknown source addresses until you remove a sufficient number of secure MAC addresses to drop below the maximum value.
- restrict Drops packets with unknown source addresses until you remove a sufficient number of secure MAC addresses to drop below the maximum value and causes the SecurityViolation counter to increment. It will send a Syslog message and an SNMP trap as well.
- shutdown Puts the interface into the error-disabled state immediately and sends an SNMP trap notification

The show ip interface command is incorrect because it displays protocol-related information about an interface, and nothing pertaining to switch port security.

The show interfaces switchport command is incorrect because it displays non-security related switch port information, such as administrative and operational status and trunking.

The show interfaces port-security command is incorrect because this is not a valid Cisco command.

Objective:

Infrastructure Security

Sub-Objective:

Configure, verify, and troubleshoot port security

References:

QUESTION 171

Which of the following statements are TRUE regarding the following output? (Choose all that apply.)

```
Router# show ip route
```

```
Gateway of last resort is 192.168.15.1 to network 0.0.0.0
```

```
<<output omitted>>
```

```
D 192.168.10.0 [90/2172416] via 192.168.15.254, 0:01:42, Serial0/1/0
C 192.168.14.0 is directly connected, Serial0/0/0
D 192.168.52.0 [90/2172416] via 192.168.15.254, 0:00:35, Serial0/1/0
[90/2172416] via 192.168.15.5, 0:02:05, Serial0/0/0
C 192.168.15.0 is directly connected, Serial0/1/0
C 192.168.20.0 is directly connected, Serial0/0/1
S 192.168.50.0 [1/0] via 192.168.53.1
C 192.168.33.0 is directly connected, Loopback1
D 192.168.25.0 [90/2196545] via 192.168.20.254, 0:01:20, Serial0/0/1
```

- A. There are four default routes on this router.
- B. There are four physically connected interfaces on this router.
- C. This router is running EIGRP.
- D. The metric for the routes learned via a routing protocol is 90.
- E. A packet for the 192.168.52.0 network will be load-balanced across two paths.

Correct Answer: CE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

This router is running EIGRP and a packet for the 192.168.52.0 network will be load-balanced across two paths.

EIGRP routes display with a D code in the leftmost column of the show ip route command. The D stands for Diffusing Update Algorithm (DUAL), which is the algorithm used by EIGRP to determine the best and potential backup paths to each remote network. There are four EIGRP-learned routes in this exhibit.

When two routes with equal metrics exist in the routing table, EIGRP will send packets using both paths. In the output there are two routes listed for the 192.168.52.0 network. Both have the same metric value (2172416). Therefore, packets will be sent to that network via the Serial 0/1/0 interface to the neighbor at 192.168.15.254 and via the Serial 0/0/0 interface to the neighbor at 192.168.15.5. Both paths, either directly or indirectly, lead to the 192.168.52.0 network, and both paths have the same cost.

There are not four default routes on this router. The D represents EIGRP-learned routes, not default routes. There is one default route, as indicated by the line of output that says Gateway of last resort is 192.168.15.1 to network 0.0.0.0. Because Serial0/1/0 is directly connected to the 192.168.15.0 network, packets that are destined for networks not found in the routing table will be sent out on that interface.

The C in the leftmost column of the show ip route command represents directly connected networks, of which there are four in the exhibit. Closer examination, however, reveals that one of these entries (for network 192.168.33.0) is connected to a loopback interface (Loopback1), as opposed to a physical interface:

```
C 192.168.33.0 is directly connected, Loopback1
```

Loopback interfaces are virtual, software interfaces that appear in the routing table, but do not represent a physical interface on the router. Therefore, there are three physically connected interfaces on this router, not four.

The metric for the routes learned via a routing protocol is not 90. The 90 in the scenario output is the administrative distance (AD) of the route, and the 2196545 is the metric value (see below):

D 192.168.25.0 [90/2196545] via 192.168.20.254, 0:01:20, Serial0/0/1

Objective:

Routing Fundamentals

Sub-Objective:

Interpret the components of routing table

References:

QUESTION 172

You are purchasing a device to upgrade your network. You need to determine the type of device required, as well as the number and type of required interfaces. The device will host three LAN subnets and a T1 Internet connection.

Which of the following device and interface combinations will support this requirement without providing any unnecessary interfaces or using subinterfaces?

- A. a switch with one Ethernet interface and three serial interfaces
- B. a router with one serial interface and three Ethernet interfaces
- C. a router with one serial interface and one Ethernet interface
- D. a switch with one modem and three serial interfaces

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

This deployment will require a router with one serial interface and three Ethernet interfaces. When LAN subnets and the Internet must be connected, you must deploy a device that can make decisions based on IP addresses. This is the function of a router. Each LAN subnet will require a separate Ethernet interface, and the T1 connection requires a serial interface, so the router must have one serial interface and three Ethernet interfaces.

A switch cannot be used to connect separate subnets and the Internet. This requires a router. Switches make forwarding decisions based on MAC addresses. In this deployment, decisions must be made on the basis of IP addresses. Moreover, switches only have Ethernet interfaces, so a switch could not handle the T1 connection.

A router with one serial and one Ethernet interface will not be sufficient. Each LAN subnet will require a separate Ethernet interface.

Objective:

Network Fundamentals

Sub-Objective:

Describe the impact of infrastructure components in an enterprise network

References:

[Cisco > Home > Internetworking Technology Handbook > Internetworking Basics > Bridging and Switching Basics](#)

QUESTION 173

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

- A. show vtp status
- B. show vtp domain
- C. show vtp statistics

D. show vtp counters

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

show vtp {counters | status}

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

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

```
Router#show vtp counters

VTP statistics:
Summary advertisements received: 7
Subset advertisements received: 6
Request advertisements received: 0
Summary advertisements transmitted: 894
Subset advertisements transmitted: 13
Request advertisements transmitted: 3
Number of config revision errors: 0
Number of config digest errors: 0
Number of V1 summary errors: 0
VTP pruning statistics:

Trunk Join Transmitted Join Received Summary advts received
from on-pruning-capable device
-----
Fa0/2 43450 42691 6
```

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

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

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot inter-VLAN routing

References:

QUESTION 174

You are the network administrator for your company. The Chief Technical Officer of the company is looking for a routing solution that satisfies the following requirements:

- No routing protocol advertisements
- Increased network security
- No routing protocol overhead
- Not concerned about fault tolerance

Which of the following routing techniques matches the criteria?

- A. Dynamic routing
- B. Hybrid routing
- C. Static routing
- D. Public routing

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The static routing technique matches the criteria given in this scenario. Static routing is a process of manually entering routes into a routing table. Static routes are not recommended for large networks because static routes are manually configured on the router. However, if a single link is used to connect an enterprise to an Internet Service Provider (ISP), then static routing is the best option.

The following are characteristics of static routing:

- Configuring static routes does not create any network traffic.
- Manually configured static routes do not generate routing updates and therefore do not consume any network bandwidth.
- Router resources are used more efficiently.
- Static routes are not recommended for large networks because they are manually configured on the router and maintaining the routes can become problematic.
- Static route configuration is not fault tolerant, because static routes do not automatically adapt to changes in the network.

The dynamic routing option is incorrect because route updates consume bandwidth and overhead. While the scenario is not concerned with routing protocol overhead, it states that there should be no bandwidth consumption by route advertisements.

Hybrid routing and public routing are not valid routing techniques in Cisco terminology.

Objective:

Routing Fundamentals

Sub-Objective:

Compare and contrast static routing and dynamic routing

References:

[Cisco Documentation > Internetworking Technology Handbook > Routing Basics > Algorithm Types](#)

QUESTION 175

How many IP addresses are available for hosts in the 192.168.16.64 /26 subnet?

- A. 14
- B. 30
- C. 62
- D. 126

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

There are 62 IP addresses available for hosts in the 192.168.16.64 /26 subnet.

The number of host addresses is calculated as $2^n - 2$, where n is the number of host bits and 2 is subtracted to exclude the network address and the broadcast address.

An IP address has 32 available bits divided into four octets. In the 192.168.16.66 /26 address, the /26 indicates that there are 26 masking bits, or that 26 bits are reserved for the network portion of the address. This leaves 6 bits for the host addresses ($32 - 26 = 6$).

The following formula is used to calculate the number of IP addresses available for hosts:

Network address: 192.168.16.0

Subnet mask in decimal: 255.255.255.192

Subnet mask in binary: 11111111.11111111.11111111.11000000

Number of bits used for masking = 2^6

Number of hosts bits in the address = 6

Using the formula for calculating the number of hosts per subnet, we find:

Hosts formula: $2^{\text{number-of-host-bits}} - 2$

Hosts: $2^6 - 2 = 62$

For subnet 192.168.16.64, the valid host range starts from 192.168.16.65 and runs to 192.168.16.126. For subnet 192.168.16.128, the valid host range starts from 192.168.16.129 and runs to 192.168.16.190.

The options 14, 30, and 126 are incorrect because 62 IP addresses are available for hosts in the 192.168.16.64/26 subnet.

The correct mask for the size network desired is critical to proper network function. For example, assume a router has an interface Fa0/0 hosting a LAN with 20 computers configured as shown in the following output of show interfaces command:

```
Router# show interfaces
FastEthernet0 is up, line protocol is up
Hardware address is 000b.12bb.4587
Internet address 192.168.10.30/30
```

In this example, the computers will not be able to access anything beyond the LAN because the mask /30 only allows for 2 addresses when 21 (including the router interface) are required.

Objective:

Network Fundamentals

Sub-Objective:

Apply troubleshooting methodologies to resolve problems

References:

[Cisco > Technology Support > IP > IP Routing > Design TechNotes > IP Addressing and Subnetting for New Users > Document ID: 13788 > Understanding IP Addresses](#)

QUESTION 176

Refer to the following sample output:

```
GigabitEthernet0/2 is up, line protocol is up
Internet address is 11.1.1.1/16
Broadcast address is 255.255.255.255
Address determined by setup command
MTU is 1500 bytes
Helper address is not set
Directed broadcast forwarding is disabled
Outgoing access list is not set
Inbound access list is not set
Proxy ARP is enabled
Local Proxy ARP is disabled
Security level is default
Split horizon is enabled
ICMP redirects are always sent
ICMP unreachable are always sent
ICMP mask replies are never sent
IP fast switching is enabled
IP fast switching on the same interface is disabled
IP Flow switching is disabled
IP CEF switching is enabled
IP Feature Fast switching turbo vector
IP VPN Flow CEF switching turbo vector
IP multicast fast switching is enabled
IP multicast distributed fast switching is disabled
IP route-cache flags are Fast, CEF
Router Discovery is disabled
IP output packet accounting is disabled
IP access violation accounting is disabled
TCP/IP header compression is disabled
RTP/IP header compression is disabled
Policy routing is enabled, using route map PBR
Network address translation is disabled
BGP Policy Mapping is disabled
IP Multi-Processor Forwarding is enabled <===== MPF information
IP Input features, "PBR",
are not supported by MPF and are IGNORED
IP Output features, "NetFlow",
are not supported by MPF and are IGNORED
```

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

- A. show interfaces
- B. show interfaces summary
- C. show ip interface
- D. show interfaces serial

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The show ip interface command will produce the displayed output. The show ip interface command is used to view the usability status of Internet Protocol (IP) interfaces. The complete syntax of this command is:

show ip interface [type number] [brief]

Following is a brief description of the parameters used in this command:

type: An optional parameter that refers to the type of interface.

number: An optional parameter that refers to the interface number.

brief: An optional parameter used to view a summarized display of the usability status information for every interface

The show interfaces command does not generate the displayed output. This command is used to view information regarding statistics for specific interfaces.

The show interfaces summary command does not generate the displayed output. This command provides a summarized view of all interfaces configured on a device.

The show interfaces serial command does not generate the displayed output. This command is used to view information for a serial interface.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Troubleshoot interface and cable issues (collisions, errors, duplex, speed)

References:

QUESTION 177

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.



Objective:

Routing Fundamentals

Sub-Objective:

Compare and contrast distance vector and link-state routing protocols

References:

[Cisco > Articles > Cisco Networking Academy > CCNP 1: Advanced IP Addressing Management](#)

[Cisco > Internetworking Technology Handbook > Routing Information Protocol \(RIP\)](#)

QUESTION 178

You have the following configuration on your router:

```
ip dhcp pool POOLNAME
network 10.1.0.0 255.255.255.0
default-router 10.1.0.254
dns-server 10.1.0.200
```

What command would you run to prevent the last available IP address in the scope from being allocated to a host via DHCP?

- A. ip dhcp restrict 10.1.0.254
- B. ip dhcp excluded-address 10.1.0.253
- C. ip dhcp excluded-address 10.1.0.254
- D. ip dhcp 10.1.0.253 excluded-address

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

In this scenario, you would run the ip dhcp excluded-address 10.1.0.253 command in global configuration mode to prevent DHCP allocation of the last available IP address in the scope. The ip dhcp excluded-address command is used to prevent DHCP from handing out IP addresses that are already statically configured on your network. The command can include a single IP address to exclude, or an entire range, such as:

```
Router(config)# ip dhcp excluded-address 10.1.0.100 10.1.0.125
```

The command above would block the entire range of 10.1.0.100 through 10.1.0.125 from being allocated by

DHCP. If the next IP address in sequence to be assigned would have been 10.1.0.100, DHCP will skip the range and assign 10.1.0.126 as the next host address.

You would not execute `ip dhcp excluded-address 10.1.0.254`. This is the address of the router and it will automatically be excluded.

The other commands are incorrect because they are not valid Cisco IOS commands.

Objective:

Infrastructure Services

Sub-Objective:

Configure and verify DHCP on a router (excluding static reservations)

References:

[Cisco > Support > Cisco IOS Software > Configuring the Cisco IOS DHCP Server > Excluding IP Addresses](#)

QUESTION 179

You have executed the following commands on switch55:

```
switchA(config)# dot1x system-auth-control
switchA(config)# aaa new-model
switchA(config)# radius-server host 192.168.105.67 key firstKey111
switchA(config)# aaa authentication dot1x default group radius
switchA(config)# interface range Fa 0/1 - 11
switchA(config-if)# switchport mode access
switchA(config-if)# dot1x port-control auto
```

What is the result of executing the given commands? (Choose two.)

- A. Only the listed RADIUS server is used for authentication
- B. 802.1X authentication is enabled on the Fa0/1 interface only
- C. The key for the RADIUS server is firstKey111
- D. AAA is not enabled on the switch

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

As a result of executing these commands, the default list is used for the RADIUS server for authentication, and the key for the RADIUS server is firstKey111.

A RADIUS server combines the authentication and authorization processes. Before you configure the RADIUS server, you should enable AAA by using the `aaa new-model` command in global configuration mode. Then, you can specify the location of the RADIUS server and the key using the `radius-server host` command. In this case, the RADIUS server is located at the IP address 192.168.105.67 and requires the key firstKey111 as the encryption key. This key must be mutually agreed upon by the server and the clients.

The `aaa authentication dot1x default group radius` command creates a method list for 802.1X authentication. The default group radius keywords specify that the default method will be to use all listed RADIUS servers to authenticate clients. Since only one is listed, it will be the only one used.

It is incorrect to state that 802.1X authentication is enabled only on the Fa0/1 interface. The interface range Fa 0/1 - 11 and the `dot1x port-control auto` commands specify that 802.1X authentication is enabled on the interfaces Fa0/1 to Fa0/11.

It is incorrect to state that AAA is not enabled on the switch. The `aaa new-model` command enables AAA

globally on the switch.

Objective:

Infrastructure Security

Sub-Objective:

Describe device security using AAA with TACACS+ and RADIUS

References:

[Cisco > Support > Cisco IOS Security Command Reference: Commands A to C > aaa new-model](#)

[Cisco > Support > Cisco IOS Security Command Reference: Commands D to L > dot1x port-control](#)

[Cisco > Support > Cisco IOS Security Command Reference: Commands M to R > radius-server host](#)

QUESTION 180

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.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot STP-related optional features

References:

[Cisco > Technology Support > LAN Switching > Spanning Tree Protocol > Technology White Paper > Understanding Rapid Spanning Tree Protocol \(802.1w\)](#)

QUESTION 181

Which Cisco IOS command can be issued on a router to test the connectivity of one interface from another interface on the same router?

- A. ping (with no address specified)
- B. ping (with an address specified)
- C. tracert
- D. traceroute

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The extended ping Cisco IOS utility, which is issued with no address specified, can be issued on a router to test connectivity between two remote routers. The ping utility uses Internet Control Messaging Protocol (ICMP) packets. An ICMP echo request is sent to the destination host. Upon its receipt, the destination host responds to the sending host with an ICMP echo reply. When the echo reply is received, the connectivity is verified.

Below is sample output of the extended ping command:

```
Router R#ping
Protocol [ip]:
Target IP address: 10.10.10.1
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 12.1.10.2
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 ping command with an address specified is incorrect because you when you issue this command you will either receive a reply from the destination or a destination unreachable message. It will not prompt for additional information as shown which is what allows you to specify the endpoints for the ping.

The traceroute command is not correct for this scenario because this command traces the path between the host issuing the command and the target network.

The tracert command is not a Cisco IOS command, but a Microsoft command.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot IPv4 and IPv6 static routing

References:

[Cisco > Tech Notes > Using the Extended ping and Extended traceroute Commands > Document ID: 13730 > The Extended ping Command](#)

QUESTION 182

Which of the following statements best describes the result of issuing the command `standby 44 timers 3 1` on an HSRP router?

- A. The holdtime will be set to a value of 3, and the hellotime will be set to a value of 1.
- B. The status of the standby router will be displayed as unknown expired.
- C. The role of active router will be passed repeatedly from one router to another.
- D. The router will be configured to reassume the role of active router in the event that the router fails and is subsequently restarted.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

When the command `standby 44 timers 3 1` is issued on a Hot Standby Routing Protocol (HSRP) router, the role of active router will be passed repeatedly from one router to another. This behavior occurs when the timers are set incorrectly. The syntax for the standby timers command is `standby [group-number] timers [hellotime holdtime]`.

The hellotime variable is the number of seconds between hello messages and is set to a value of 3 by default.

The holdtime variable is the number of seconds that the HSRP standby router will wait before assuming that the active router is down; if the standby router believes the active router to be down, it will assume the role of active router.

The holdtime is set to a value of 10 by default. The holdtime should be set to a value at least three times the value of the hellotime. Otherwise, the active router might not be able to respond before the standby router assumes that the active router is down and becomes the new active router.

Because the command `standby 44 timers 3 1` sets the hellotime to a value of 3 and the holdtime to a value of 1, the role of active router will be passed from one standby router to the next. To set the holdtime to a value of 3 and the hellotime to a value of 1, the command `standby 44 timers 1 3` should be issued. To reset the timer values to their default values, the command `no standby group-number timers` should be issued.

The status of the standby router will be displayed as unknown expired if a Physical layer problem exists. The unknown expired status can also be displayed if only one HSRP router is configured for the subnet.

To configure an HSRP router to reassume the role of active router in the event that the router fails and is subsequently restarted, the command `standby group-number preempt` should be issued. When the HSRP active router fails or is shut down, the standby router assumes the role of active router. By default, when the original HSRP active router is restarted, it does not take the role of active router away from the original standby router, even if the original active router has a higher priority value. The command `standby group-number preempt` changes this default behavior.

The holdtime will not be set to a value of 3, and the hellotime will not be set to a value of 1. On the contrary, the hellotime will be set to a value of 3 and the holdtime will be set to a value of 1.

Objective:

Infrastructure Services

Sub-Objective:

Configure, verify, and troubleshoot basic HSRP

References:

[Cisco IOS IP Application Services Command Reference > show vrrp through synguard \(virtual server\) > standby timers](#)

QUESTION 183

Which Cisco Internetwork Operating System (IOS) command is used to assign a router a name for identification?

- A. description
- B. banner motd
- C. hostname
- D. banner exec

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The hostname command is used to assign the router a name for identification. This command is a global configuration mode command. The syntax of the command is as follows:

Router(config)# hostname [name]

The name parameter of the command specifies the new host name for the router.

The description command is incorrect because this command is used to set a description for an interface. The description command is an interface configuration mode command.

The banner motd command is used to specify a message of the day (MOTD) banner to users logging into the router. This is a global configuration mode command, but it does not assign a name to the router for identification.

The banner exec command enables a banner message to be displayed when an EXEC process is created; for example, if a line is activated or an incoming connection is made to a telnet line.

Objective:

Infrastructure Management

Sub-Objective:

Configure and verify initial device configuration

References:

[Cisco > Cisco IOS Configuration Fundamentals Command Reference > F through K > hostname](#)

QUESTION 184

A new trainee is setting up a router in a test lab, and he asks you to describe the use of the connector marked BRI on the router.

Which is a correct use for this connector?

- A. A WAN interface for a T1 connection
- B. A LAN interface to connect to a switch
- C. An interface to connect a console cable
- D. A WAN interface for an ISDN connection

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:


The connector marked BRI is used for an Integrated Services Digital Network (ISDN) connection, specifically a basic rate interface (BRI). An ISDN basic rate interface provides three channels: a D channel for control signaling, and two B or bearer channels for data, resulting in 128 bits of bandwidth.

A WAN interface for a T1 connection would be connected to a serial port on the router, not the BRI interface. It would not accept a basic rate ISDN connection.

A LAN interface to connect to a switch would be an Ethernet connection that used either an RJ-45 connector or a legacy AUI connector. It would not accept a basic rate ISDN connection.

An interface to a console connector will look like an RJ-45 Ethernet connector but will only accept a console or rollover cable, and is used to manage the router. It would not accept a basic rate ISDN connection.

These various ports can be seen on the backplane of a router as shown below:



Objective:

WAN Technologies

Sub-Objective:

Describe WAN access connectivity options

References:

<http://www.tutorialsworld.com/networking/routers/cisco-routers-ios.htm#Hardware%20Components>:
[Cisco>Support>Product Support>End-of-Sale and End-of-Life Products>Cisco 3600 Series Multiservice Platforms>Troubleshoot and Alerts> Troubleshooting TechNotes> Understanding the 1-Port ISDN BRI \(S/T\) WAN Interface Card \(WIC-1B-S/T or WIC36-1B-S/T\)](#)

QUESTION 185

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

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

Correct Answer: AB

Section: (none)

Explanation

Explanation/Reference:

Explanation:

OSPF will perform an election for a designated router (DR) and backup designated router (BDR) on every multi-access network segment. Multi-access segments are defined as segments where more than two hosts can reach each other directly, such as a shared Ethernet segment (broadcast multi-access) or Frame Relay (non-broadcast multi-access, or NBMA).

DR and BDR elections do not occur on point-to-point or point-to-multipoint segments. Point-to-point and point-to-multipoint segments are not considered multi-access segments. OSPF routers on these network types will

establish an adjacency without a DR/BDR election.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot single area and multi-area OSPFv2 for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub, virtual-link, and LSAs)

References:

[Cisco > Support > IP > IP Routing > Technology Information > Technology White Paper > OSPF Design Guide > Document ID: 7039 > DR Election](#)

QUESTION 186

You have a class C address range and are planning a network that has an average of 50 hosts per subnet.

How many host bits will have to be borrowed for subnetting so that the maximum number of subnets can be implemented?

- A. 1
- B. 2
- C. 3
- D. 4
- E. 6

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A class C address has 8 bits in host space. By using 2 bits from the host space for subnetting, leaving 6 host bits, you can create subnets that can accommodate up to 62 hosts each ($2^6 - 2 = 62$). This will ensure that the requirement of 50 hosts per subnet is met and the maximum number of subnets is provided.

The formulas to calculate the number of subnets and hosts are:

Number of subnets = $2^{\text{number-of-subnet-bits}}$

Number of hosts per subnet = $2^{\text{number-of-host-bits}} - 2$

If you take 1 bit for subnetting:

Number of subnets = $2^1 = 2$

Number of hosts per subnet = $2^7 - 2 = 126$

This results in a mask of 255.255.255.128 or /25. Since each subnet need not be bigger than 50, this solution would not maximize the number of subnets.

If you take 2 bits for subnetting:

Number of subnets = $2^2 = 4$

Number of hosts per subnet = $2^6 - 2 = 62$

This results in a mask of 255.255.255.192 or /26. This solution would create more subnets, but the subnets are smaller than the requirement.

If you take 3 bits for subnetting:

Number of subnets = $2^3 = 8$

Number of hosts per subnet = $2^5 - 2 = 30$

This results in a mask of 255.255.255.224 or /27. This would create more subnets, but the subnets are smaller than the requirement.

If you take 4 bits for subnetting:

Number of subnets = $2^4 = 16$

Number of hosts per subnet = $2^4 - 2 = 14$

This results in a mask of 255.255.255.240 or /28. This solution would create more subnets, but the subnets are smaller than the requirement.

If you take 6 bits for subnetting:

Number of subnets = $2^6 = 64$

Number of hosts per subnet = $2^2 - 2 = 2$

This mask, 255.255.255.252 or /30, yields only 2 IP addresses, but is quite commonly used on a point-to-point link, such as between two routers. This solution would create more subnets, but the subnets are smaller than the requirement.

You will always subtract 2 from the number of hosts (the formula of $2^{\text{number-of-host-bits}} - 2$) because the all-zeros bit address is reserved for the network address and the all-ones bit address is reserved for the broadcast address.

Prior to Cisco IOS Software Release 12.0, it was common practice to subtract 2 from the networks formula ($2^{\text{number-of-subnet-bits}}$) to exclude addresses of all 1s and all 0s (called the all-ones subnet and subnet zero). Today that range is usable, except with some legacy systems. On certain networks with legacy software, you may need to use the previous formula ($2^n - 2$) to calculate the number of subnets.

Objective:

Network Fundamentals

Sub-Objective:

Apply troubleshooting methodologies to resolve problems

References:

[Cisco > Technology Support > IP > IP Routing > Design TechNotes > IP Addressing and Subnetting for New Users > Document ID: 13788 > Understanding IP Addresses](#)

[Cisco > Technology Support > IP > IP Addressing Services > Design TechNotes > Subnet Zero and the All-Ones Subnet > Document ID: 13711](#)

QUESTION 187

What are the three types of Internet Protocol version 6 (IPv6) addresses? (Choose three.)

- A. Unicast
- B. Broadcast
- C. Dual-cast
- D. Anycast
- E. Multicast

Correct Answer: ADE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Unicast, multicast, and anycast are types of IPv6 addresses.

The following are the IPv6 address types:

- Unicast address: These types of addresses are used to define a single destination interface. A packet sent to a unicast address is delivered to the specific interface.
- Multicast address: These types of addresses are used to define a group of hosts. When a packet is sent to a multicast address, it is delivered to all the hosts identified by that address. Multicast addresses begin with the prefix FF00::/8 and the second octet identifies the range over which the multicast address is propagated. Some special case IPv6 multicast addresses:
 - FF01:0:0:0:0:0:0:1: Indicates all-nodes address for interface-local scope.
 - FF02:0:0:0:0:0:0:2: Indicates all-routers address for link-local.
- Anycast address: These types of addresses are used to identify a set of devices. These addresses are also assigned to more than one interface belonging to different nodes. A packet sent to an anycast address is delivered to just one of the interfaces, based on which one is closest. For example, if an anycast address is assigned to a set of routers, one in India and another in the U.S., the users in the U.S. will be routed to U.S. routers and the users in India will be routed to a server located in India.

The broadcast option is incorrect because these types of addresses are not supported by IPv6. Broadcast functionality is provided by multicast addressing.

The dual-cast option is incorrect because this is not a valid Cisco address type.

Objective:

Network Fundamentals

Sub-Objective:

Compare and contrast IPv6 address types

References:

QUESTION 188

Which media access control method is used by Ethernet technology to minimize collisions in the network?

- A. CSMA/CD
- B. token passing
- C. back-on algorithm
- D. full-duplex

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Carrier Sense Multiple Access - Collision Detection (CSMA/CD) is used by Ethernet technology to minimize collisions in the network. The CSMA/CD method uses a back-off algorithm to calculate random time for retransmission after a collision. When two stations start transmitting at the same time, their signals will collide. The CSMA/CD method detects the collision, and both stations hold the retransmission for a certain amount of time that is determined by the back-off algorithm. This is an effort to help ensure that the retransmitted frames do not collide.

Token passing is used by the token-ring network topology to control communication on the network.

Full-duplex is the Ethernet communication mode that allows workstation to send and receive simultaneously. With the use of full-duplex, the bandwidth of the station can effectively be doubled. Hubs are not capable of handling full-duplex communication. You need dedicated switch ports to allow full-duplex communication.

The back-on algorithm is an invalid option. There is no such contention method.

Objective:

LAN Switching Fundamentals

Sub-Objective:
Describe and verify switching concepts

References:
[Cisco Documentation > Internetworking Technology Handbook > Ethernet Technologies](#)

QUESTION 189

Which of the following IP addresses are valid Class B host addresses if a default Class B mask is in use?
(Choose all that apply.)

- A. 10.6.8.35
- B. 133.6.5.4
- C. 192.168.5.9
- D. 127.0.0.1
- E. 190.6.5.4

Correct Answer: BE

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The IP addresses 133.6.5.4 and 190.6.5.4 are both valid Class B addresses when a default mask is in use. The Class B default mask is 255.255.0.0 and the range of valid addresses is 128.0.0.0-191.255.255.255.

The IP address 10.6.8.35 is a Class A address. The Class A default mask is 255.0.0.0 and the range of valid addresses is 1.0.0.0 - 127.255.255.255, with the exception of the range 127.0.0.1 - 127.255.255.255, which is reserved and cannot be assigned.

The IP address 192.168.5.9 is a Class C address. The Class C default mask is 255.255.255.0 and the range of valid addresses is 192.0.0.0 - 223.255.255.255.

The IP address 127.0.0.1 is a Class A address, but it comes from a reserved portion that cannot be assigned. The range 127.0.0.1 - 127.255.255.255 is used for diagnostics, and although any address in the range will work as a diagnostic address, 127.0.0.1 is known as the loopback address. If you can ping this address, or any address in the 127.0.0.1 - 127.255.255.255 range, then the NIC is working and TCP/IP is installed. The Class A default mask is 255.0.0.0 and the range of valid addresses is 1.0.0.0 - 127.255.255.255, with the exception of the range 127.0.0.1 - 127.255.255.255, which is reserved and cannot be assigned.

Objective:
Network Fundamentals
Sub-Objective:
Compare and contrast IPv4 address types

References:
[Cisco > Design Tech Notes > IP Routing > IP Addressing and Subnetting for New Users > Understanding IP Addresses > Document ID: 13788](#)

QUESTION 190

What is the purpose of using the show arp command?

- A. To view the ARP statistics only for a particular interface
- B. To view details regarding neighboring devices discovered by ARP
- C. To view global ARP information such as timer and hold time
- D. To view the Address Resolution Protocol (ARP) cache

Correct Answer: D

Section: (none)
Explanation

Explanation/Reference:

Explanation:

The show arp command is used to view the Address Resolution Protocol (ARP) cache. ARP is used by the Internet Protocol (IP) to find the Media Access Control (MAC) address or the hardware address of a host. The main function of ARP is to translate IP addresses to MAC addresses. The process of obtaining the address of a computer in the network is known as address resolution. This process is accomplished by sending an ARP packet from a source to a destination host. The destination host responds to the ARP packet by replying back to the source and including its own MAC address. Once the source host receives the reply, it will update its ARP cache with the new MAC address.

The complete syntax of the show arp command is:

show arp [ip-address [locationnode-id] | hardware-address [locationnode-id] | traffic [locationnode-id | interface-instance] | trace [error [locationnode-id] | dev [locationnode-id] | events [locationnode-id] | table [locationnode-id] packets [locationnode-id] | [locationnode-id]] | type instance| [locationnode-id]

The following is a brief description of the parameters used with this command:

ip-address: An optional parameter that displays specific ARP entries.

locationnode-id: An optional parameter that displays the ARP entry for a specific location. The method for entering the node-id argument is rack/slot/module notation.

hardware-address: An optional parameter that displays ARP entries that match the 48-bit MAC address.

traffic: An optional parameter that displays ARP traffic statistics.

interface instance: Either a physical interface instance or a virtual interface instance:

Physical interface instance: the naming notation is rack/slot/module/port and a slash mark between values is required as part of the notation where:

rack refers to the chassis number of the rack.

slot refers to the physical slot number of the line card.

module refers to the module number. A physical layer interface module (PLIM) is always 0.

port refers to the physical port number of the interface.

Virtual interface instance: the number range is variable depending on the type of interface.

trace: An optional parameter that displays the ARP entries in the buffer.

error: An optional parameter that displays the ARP error logs.

dev: An optional parameter that displays the ARP internal logs.

events: An optional parameter that displays the ARP events logs.


table: An optional parameter that displays the ARP cache logs.

packets: An optional parameter that displays the ARP packet receive and reply logs.

type instance: An optional parameter that specifies the interface for which you want to view the ARP cache.

An example of the output of the show arp command is shown below along with a diagram of the network in which the router resides.

```
R1#show arp
Protocol Address Age (min) Hardware Addr Type Interface
Internet 192.0.5.1 120 0000.a710.4baf ARPA FastEthernet 0/1
Internet 192.0.5.6 105 0000.a710.859b ARPA FastEthernet 0/1
Internet 58.59.6.3 42 0000.a710.68cd ARPA FastEthernet 0/0
Internet 58.59.6.4 59 0000.0c01.7bbd ARPA FastEthernet 0/0
```



From the information above, we can make the following conclusions about the actions R1 will take when it receives data from PC1 destined for the Web server:

- The data frames will be forwarded out the Fa0/0 interface of R1
- R1 will place the MAC address of the Web Server (0000.0c01.7bbd) in the destination MAC address of the frames
- R1 will put the MAC address of the forwarding Fa0/0 interface (0000.a710.68cd) in the place of the source MAC address

The option stating that the show arp command is used to view the ARP statistics only for a particular interface is incorrect because this command is used to view the ARP cache. You can also view the information for a particular interface with the help of the interface instance parameter.

The options stating that the show arp command is used to view the details of neighboring devices discovered by the ARP or to view global ARP information, such as hold time and timer, are both incorrect because these are both Cisco Discovery protocol (CDP) functions, not ARP functions. The show cdp neighbors detail command is used to display details regarding the neighboring devices that are discovered by CDP, and the show cdp command displays global CDP information, such as timer and hold-time.

Objective:

Network Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot IPv4 addressing and subnetting

References:

QUESTION 191

Which of the following statements are NOT true, based on the output below?

```

Access1#show spanning-tree vlan 10

VLAN0010

Spanning tree enabled protocol rstp
Root ID Priority 24586
Address 0015.63f6.b700
Cost 19
Port 107 (FastEthernet3/0/1)
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

Bridge ID Priority 32778 (priority 32768 sys-id-ext 10)
Address 000f.f794.3d00
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
Aging Time 300

Interface Role Sts Cost Prio.Nbr Type
-----
Fa3/0/1 Root FWD 19 128.107 P2p
Fa3/0/2 Altn BLK 19 128.108 P2p
-----

```

- A. This switch is the root bridge.
- B. This switch has a priority of 32778.
- C. This switch has a MAC address of 0015.63f6.b700.
- D. All ports will be in a state of discarding, learning, or forwarding.
- E. All designated ports are in a forwarding state.
- F. This switch is using the default priority for STP

Correct Answer: AC

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The upper half of the output provides information about the root bridge. It indicates that the root bridge has a bridge priority of 24586 and a MAC address of 0015.63f6.b700. The bottom half of the output pertains to the current switch, and indicates that this switch has a bridge priority of 32778 and a MAC address of 000f.f794.3d00.

The value of the switch bridge priority is arrived at by adding the configured priority of 32768, which is indicated by the line priority 32768 sys-id-ext 10, to the VLAN ID of 10. Because 32768 is the default bridge priority for STP, this switch is set to the default priority for STP.

The priority of this switch is 32778. The bridge priority is arrived at by adding the configured priority of 32768 to the VLAN ID of 10.

This switch is not the root bridge, as indicated by the differences in priorities and MAC addresses between the root ID and the bridge ID output. If this were the root bridge, the MAC addresses and priority values would be the same in both the Root ID and the Bridge ID sections.

Finally, when a switch is using RSTP, as indicated by the output Spanning tree enabled protocol rstp, all ports will be in a state of discarding, learning, or forwarding, with all designated ports in a forwarding state. When RSTP has converged, all ports will be in either the discarding or forwarding states.

Objective:

LAN Switching Fundamentals

Sub-Objective:
Configure, verify, and troubleshoot STP protocols

References:
[Cisco > Cisco IOS Bridging Command Reference > show spanning-tree](#)

QUESTION 192

Which of the following values will be used by a router to make a routing decision when two routes have been learned from OSPF?

- A. cost
- B. administrative distance
- C. composite metric
- D. hop count

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

When two routes have been learned by OSPF to same network, the best route will be chosen based on lowest cost. Cost is the metric used in OSPF to choose the best route from all candidate routes learned through OSPF.

Administrative distance is a measure of the trustworthiness of the routing information source. It is a value used by a router to choose between multiple known routes that have been learned from different routing sources, such as different routing protocols. When routes are learned from the same routing protocol, their administrative distance will be equal, and the router will then choose the route with the lowest metric value of the routing protocol. In this case, that metric is the OSPF cost.

The composite metric is the metric used by EIGRP to choose a route when multiple routes have been learned by EIGRP.

Hop count is the metric used by RIP to choose a route when multiple routes have been learned by RIP.

Objective:

Routing Fundamentals


Sub-Objective:

Describe how a routing table is populated by different routing information sources

References:
[Support > Technology Support > IP > IP Routing > Troubleshoot and Alerts > Troubleshooting TechNotes > Route Selection in Cisco Routers](#)

QUESTION 193

Your company's network must make the most efficient use of the IP address space. In the following diagram, the circles define separate network segments. The requirements of each network segment are given in the diagram. (Click the Exhibit(s) button.)



Users complain of connectivity issues. You need to discover the problems with the network configuration.

What are the three problems with the network diagram? (Choose three.)

- A. The 172.16.1.0/30 segment requires more user address space.
- B. The 172.16.2.0/26 segment requires more user address space.
- C. The 172.16.3.0/25 segment requires more user address space.
- D. The 172.16.2.64/26 segment requires more user address space.
- E. Interface fa0/2 has an IP address that belongs to the 172.16.2.64/26 segment.
- F. Interface fa0/4 has an IP address that belongs to the 172.16.2.0/26 segment.
- G. Interface fa0/3 has an IP address outside the 172.16.3.0/25 segment.

Correct Answer: AFG

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The given exhibit has three problems:

- The 172.16.1.0/30 segment requires more user address space.
- Interface Fa0/4 has an IP address that belongs to the 172.16.2.0/26 segment.
- Interface Fa0/3 has an IP address outside the 172.16.3.0/25 segment.

The 172.16.1.0/30 segment, as configured, will only support two hosts. This segment needs to support three hosts, the two servers, and the Fa0/1 interface. The number of hosts that a subnet is capable of supporting is a function of the number of host bits in the subnet mask. When that has been determined, the following formula can be used to determine the number of hosts yielded by the mask:

$$2^n - 2 = X$$

(where n = the number of host bits in the mask and X = the number of hosts supported)

In this example with a 30-bit mask, 2 host bits are left in the mask. When that is plugged into the formula, it yields only two usable addresses. The -2 in the formula represents the two addresses in each subnet that cannot be assigned to hosts, the network ID and the broadcast address. Therefore, the segment should be configured with the 172.16.1.0/29 address range, which supports up to six hosts.

Interface fa0/4, as configured, has an IP address that belongs to the 172.16.2.0/26 segment. With a 26-bit mask and the chosen class B address, the following network IDs are created:

```
172.16.0.0
172.16.0.64
172.16.1.128
172.16.1.192
172.16.2.0
172.16.2.64
172.16.2.128
172.16.2.192
172.16.2.0
172.16.2.64
172.16.2.128
172.16.2.192
...and so on, incrementing each time by 64 in the last octet
```

The 172.16.2.0/26 segment is allocated host addresses in the 172.16.2.1 through 172.16.2.62 range (the last address, 172.16.2.63, is the broadcast address and cannot be assigned). Interface fa0/4 should be assigned an IP address in the 172.16.2.64/26 range, which includes host addresses in the 172.16.2.65 through 172.16.2.126 range.

Interface Fa0/3, as configured, has an IP address outside the 172.16.3.0/25 segment. With a 25-bit mask and the chosen class B address, the following network IDs are created:

```
172.16.0.0
172.16.0.128
172.16.1.0
172.16.1.128
172.16.2.0
172.16.2.128
172.16.3.0
172.16.3.128
...and so on, incrementing each time by 128 in the last octet
```


Interface Fa0/3 should be allocated an IP address in the 172.16.3.1 through 172.16.3.126 range.

The 172.16.2.0/26 segment does not require more user address space. With a 26-bit mask, 6 bits are left for hosts, and by using the above formula it can be determined that it will yield 62 hosts. It requires 32.

The 172.16.2.64/26 segment does not require more user address space. With a 26-bit mask, 6 bits are left for hosts, and by using the above formula it can be determined that it will yield 62 hosts. It requires 32.

Interface Fa0/2 does not have an IP address that belongs to the 172.16.2.64/26 segment. The 172.16.2.64/26 segment includes addresses 172.16.2.65-172.16.2.126. Because its address is 172.16.2.1, it belongs in the 172.16.2.0/26 network (from 172.16.2.1-172.16.2.62), so it is correctly configured.

The network should be configured as shown in the following image:



Objective:
 Network Fundamentals
 Sub-Objective:
 Configure, verify, and troubleshoot IPv4 addressing and subnetting

References:
[IP Addressing and Subnetting for New Users](#)

QUESTION 194

What is the possible IP range that can be assigned to hosts on a subnet that includes the address 192.168.144.34/29?

- A. 192.168.144.32 - 192.168.144.63
- B. 192.168.144.33 - 192.168.144.38
- C. 192.168.144.33 - 192.168.144.48
- D. 192.168.144.28 - 192.168.144.40

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Range 192.168.144.33 - 192.168.144.38 is the correct answer. To determine the range of addresses that can be assigned in a subnet, you must first determine the network ID of the subnetwork and the broadcast address of the subnetwork. All addresses that can be assigned to hosts will lie between these endpoints. The network ID can be obtained by determining the interval between subnet IDs. With a 29-bit mask, the decimal equivalent

of the mask will be 255.255.255.248. The interval between subnets can be derived by subtracting the value of the last octet of the mask from 256. In this case, that operation would be $256 - 248 = 8$. Therefore, the interval is 8.

The first network ID will always be the classful network you started with (in this case 192.168.144.0). Each subnetwork ID will fall at 8-bit intervals as follows:

192.168.144.0
192.168.144.8
192.168.144.16
192.168.144.24
192.168.144.32
192.168.144.40

We can stop at the 192.168.144.40 address because the address given in the scenario, 192.168.144.34, is in the network with a subnet ID of 192.168.144.32. Therefore, since the broadcast address for this network will be 1 less than the next subnet ID (192.168.144.39), the valid range of IP addresses is 192.168.144.33 - 192.168.144.38. 192.168.144.39 will be the broadcast address for the next subnet, and 192.168.144.40 will be the first valid address in the next subnet.

None of the other answers is the correct range.

Objective:

Network Fundamentals

Sub-Objective:

Apply troubleshooting methodologies to resolve problems

References:

[Cisco > Technology Support > IP > IP Routing > Design TechNotes > IP Addressing and Subnetting for New Users > Document ID: 13788 > Understanding IP Addresses](#)

QUESTION 195

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 that messages are authenticated using SHA with a key of CISCO.

Objective:

Infrastructure Management

Sub-Objective:

Configure and verify device-monitoring protocols

References:

[SNMP Configuration Guide, Cisco IOS XE Release 3SE \(Catalyst 3850 Switches\) > SNMPv3](#)

QUESTION 196

Based on the command output below, which of the interfaces on Router1 are trunk ports?

```
Router1# show mac-address-table

Dynamic Addresses Count: 14
Secure Addresses (User-defined) Count: 0
Static Addresses (User-defined) Count: 0
System Self Addresses Count: 23
Total MAC addresses: 33
Non-static Address Table:
Destination Address Address Type VLAN Destination Port
```

```
-----
0010.0de0.e289 Dynamic 1 FastEthernet0/1
0010.7b00.1540 Dynamic 1 FastEthernet0/5
0010.7b00.1545 Dynamic 1 FastEthernet0/5
0060.5cf4.0076 Dynamic 3 FastEthernet0/1
0060.5cf4.0077 Dynamic 3 FastEthernet0/1
0060.5cf4.1315 Dynamic 2 FastEthernet0/1
0060.70cb.f301 Dynamic 1 FastEthernet0/2
00e0.1e42.9978 Dynamic 1 FastEthernet0/3
```

<output omitted>

- A. Fa0/1
- B. Fa0/2
- C. Fa0/3
- D. Fa0/5

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Interface Fa0/1 is a trunk port. The output shows that it has MAC addresses that belong to VLANs 1, 2 and 3. Only trunk ports can carry traffic from multiple VLANs.

Fa0/2 is not a trunk port. It only carries traffic from VLAN 1.

Fa0/3 is not a trunk port. It only carries traffic from VLAN 1.

Fa0/5 is not a trunk port. It only carries traffic from VLAN 1.

Objective:

Infrastructure Management

Sub-Objective:

Use Cisco IOS tools to troubleshoot and resolve problems

References:

QUESTION 197

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

```
%192.168.16.0 overlaps with FastEthernet0/0
```

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

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

Correct Answer: CD

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

It's also a valuable skill to be able to recognize these problems before the router tells you about them. In the diagram below, you should be able to spot the problem with the two planned addresses on the router as being in the same subnet before you receive the error message.

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

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

Objective:

Routing Fundamentals

Sub-Objective:

Troubleshoot basic Layer 3 end-to-end connectivity issues

References:

[Cisco Documentation > Configuring IP Addressing](#)

QUESTION 198

You are the network administrator for your company. Your company has opened a new site in London. The Chief Technical Officer (CTO) of the company wants to implement a routing protocol that can provide the following features:

- Supports multiple large networks
- Does not require a hierarchical physical topology
- Supports VLSM
- Provides loop prevention and fast convergence
- Provides load balancing over un-equal cost links

Which routing protocol should be implemented in the new site?

- A. Enhanced Interior Gateway Routing Protocol (EIGRP)
- B. Open Shortest Path First (OSPF)
- C. Interior Gateway Routing Protocol (IGRP)

D. Routing Information Protocol version 2 (RIPv2)

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Enhanced Interior Gateway Routing Protocol (EIGRP) is the routing protocol that should be implemented for this scenario. EIGRP is a classless protocol that allows the use of variable length subnet masks (VLSM) and classless interdomain routing (CIDR) for the allocation of IP addresses. The following are characteristics of EIGRP:

- Supports large networks due to high scalability.
- Does not require a hierarchical physical topology.
- Provides loop prevention and fast convergence by using Diffusing Update Algorithm (DUAL).
- Performs equal cost load balancing by default.
- Can be configured to perform unequal-cost load balancing.
- Supports VLSM and CIDR.
- Is a hybrid routing protocol (a distance-vector protocol that also provides link-state protocol characteristics).
- Is a classless protocol.
- Sends partial route updates only when there are changes.
- Supports Message-Digest algorithm 5 (MD5) authentication.
- Has an administrative distance is 90 for EIGRP internal routes, 170 for EIGRP external routes, and 5 for EIGRP summary routes.
- Is only used with Cisco platforms.

All the other options are incorrect because they would not provide the features required in this scenario.

OSPF requires a hierarchical physical topology.

IGRP does not support VLSM.

RIPv2 is not designed for multiple large networks.

Objective:

Routing Fundamentals

Sub-Objective:

Compare and contrast distance vector and link-state routing protocols

References:

[Cisco > Articles > Cisco Certification > CCNP > Shooting Trouble with IP](#)

[Cisco > Support > IP > IP Routing > Design > Design TechNotes > Introduction to EIGRP > Document ID: 13669](#)

QUESTION 199

Refer to the following configuration on a Cisco router to allow Telnet access to remote users:

Router(config)#line vty 0 2

Router(config-line)#login

Router(config-line)#password guest

How many users can Telnet into this router at the same time?

- A. 0
- B. 1
- C. 2
- D. 3
- E. 5

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The given configuration will allow three users to Telnet into the router at the same time. The line vty 0 2 command specifies a range from 0 to 2; therefore, three simultaneous Telnet sessions are allowed on this Cisco router. The commands in the exhibit can be explained as follows:

Router(config)#line vty 0 2 (determines which of the five possible terminal lines are being configured. In this case, they are lines 0 through 2. It also determines the number of lines available, in that any line with no password configured will be unusable.)

Router(config-line)#login (specifies that a password will be required)

Router(config-line)#password guest (specifies the password)

The default configuration allows five simultaneous Telnet sessions on the Cisco router. For the default configuration, you would issue the vty 0 4 command in global configuration mode.

You must configure a password when enabling a router for Telnet access. Without a password, the login access to the router will be disabled and you will receive the following error message if you try to Telnet to the router:

```
router# telnet 10.10.10.1
Trying 10.10.10.1 ... Open
Password required, but none set
[Connection to 10.10.10.1 closed by foreign host]
```

Objective:

Infrastructure Management

Sub-Objective:

Configure and verify device management

References:

QUESTION 200

Which of the following are characteristics of Enhanced Interior Gateway Routing Protocol (EIGRP)? (Choose all that apply.)

- A. Requires a hierarchical physical topology
- B. Does not require a hierarchical physical topology
- C. Uses Diffusing Update Algorithm (DUAL) to provide loop prevention
- D. Uses Bellman-Ford algorithm to provide loop prevention
- E. Supports Message-Digest Algorithm 5 (MD5) authentication
- F. Does not support Message-Digest Algorithm 5 (MD5) authentication
- G. Can differentiate between internal and external routes
- H. Uses a 32-bit metric

Correct Answer: BCEGH

Section: (none)

Explanation

Explanation/Reference:

Explanation:

EIGRP does not require a hierarchical physical topology. It uses Diffusing Update Algorithm (DUAL) to provide loop prevention, and it supports Message-Digest Algorithm 5 (MD5) authentication. It can differentiate between internal and external routes, and uses a 32-bit metric.

EIGRP is a classless protocol that allows the use of variable length subnet masks (VLSM) and supports classless interdomain routing (CIDR) for allocation of IP addresses. The following are characteristics of EIGRP:

- Supports large networks due to high scalability
- Provides fast convergence using the Diffusing Update Algorithm (DUAL)
- Performs equal and unequal load balancing by default
- Supports variable length subnet masks (VLSM) and classless interdomain routing (CIDR)
- Is a hybrid routing protocol (distance-vector protocol) that also provides link-state protocol characteristics
- Is a classless protocol
- Sends partial route updates only when there are changes, reducing bandwidth usage for routing updates
- Has an administrative distance of 90 for EIGRP internal routes, 170 for EIGRP external routes, and 5 for EIGRP summary routes
- Is used only with Cisco platforms
- Provides support for IP IPX and AppleTalk protocols
- Can differentiate between internal and external routes
- Uses a 32-bit metric

EIGRP can load-balance up to four unequal cost paths. To do so, use the variance n command to instruct the router to include routes with a metric of less than n times the minimum metric route for that destination. The variable n can take a value between 1 and 128. The default is 1, which means equal cost load balancing.

The option stating that EIGRP requires a hierarchical physical topology is incorrect because EIGRP does not require or support a hierarchical routing topology.

The option stating that EIGRP uses Bellman-Ford algorithm to provide loop prevention is incorrect. EIGRP uses DUAL to provide loop prevention.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot EIGRP for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub)

References:

[Cisco > Internetworking Technology Handbook > Enhanced Interior Gateway Routing Protocol \(EIGRP\)](#)

QUESTION 201

Which prompt indicates the configuration mode at which Cisco IOS debug commands can be issued?

- A. router>
- B. router#
- C. router(config)#
- D. router(config-if)#

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You would use privileged EXEC mode, as indicated by the router# prompt, to issue Cisco IOS show and debug commands. All debug commands are entered in privileged EXEC mode. A brief description of all the debugging commands can be displayed by entering the following command in privileged EXEC mode at the command line:

debug?

Debugging output consumes high CPU processing power and can leave the system unusable. The debug commands should be reserved to troubleshoot specific problems, preferably with the help of Cisco technical support staff.

The prompt router> indicates user exec mode, which provides limited access to the router.

The prompt router(config)# indicates global configuration mode, which allows configuration settings affecting the entire router. Passing through this mode is also required to access configuration mode for specific interfaces as well.

The prompt router(config-if)# indicates interface configuration mode, which allows configuration of the interface specified when entering this mode.

Objective:

Infrastructure Management

Sub-Objective:

Configure and verify device management

References:

[Cisco > Support > Cisco IOS Software > Using the Command-Line Interface in Cisco IOS Software](#)

QUESTION 202

Which Cisco Internetwork Operating System (IOS) command can be used to configure the location of the configuration file?

- A. boot buffersize
- B. configure
- C. boot config
- D. service config

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The boot config command will configure the location of the configuration file. It must be followed by the copy run start command to be effective at next reboot. The syntax of the command is as follows:

boot config device:filename

The parameters of the command are as follows:

- Device : Specifies the device that contains the configuration file.
- Filename : Specifies the name of the configuration file.

The boot buffersize command is incorrect because this command is used to modify the buffer size used to load the IOS image. Moreover, this command no longer functions in IOS 12.4.

The configure command is incorrect because this command is used to enter the global configuration mode.

The service config command is incorrect because this command is used to enable autoloading of configuration files from a network server.

Objective:

Infrastructure Management

Sub-Objective:

Perform device maintenance

References:

[Cisco > Cisco IOS Configuration Fundamentals Command Reference > A through B > boot config](#)

QUESTION 203

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.

Objective:

WAN Technologies

Sub-Objective:

Describe WAN access connectivity options

References:

[Cisco > Articles > Network Technology > General Networking > IPsec Overview Part Two: Modes and Transforms](#)


[Cisco > The Internet Protocol Journal > The Internet Protocol Journal - Volume 3, No. 1, March 2000 > IP Security](#)

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

QUESTION 204

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 192.168.3.0/24 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.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot EIGRP for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub)

References:

[Cisco > Home > Support > Technology Support > IP > IP Routing > Technology Information > Technology White](#)

QUESTION 205

You have a router that is not syncing with its configured time source.

Which of the following is NOT a potential reason for this problem?

- A. The reported stratum of the time source is 12
- B. The IP address configured for the time source is incorrect
- C. NTP authentication is failing
- D. There is an access list that blocks port 123

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A reported stratum of 12 will not cause a router's inability to synchronize with its configured time source. The stratum value describes the device's distance from the clock source, measured in NTP server hops. When a router reports a stratum value over 15, it is considered unsynchronized. Therefore, a report of 12 could be normal.

The other options describe potential reasons for a lack of synchronization.

When you are configuring the local router with a time source, if the IP address configured for the time source is incorrect, then no synchronization will occur.

If NTP authentication is configured between the local router and its time source, and that process is failing (for example, due to a non-matching key or hashing algorithm), then synchronization will not occur.

If there were an access list applied to any interface in the path between the local router and its time source that blocks port 123 (the port used for NTP), then synchronization will not occur.

Objective:

Infrastructure Services

Sub-Objective:

Configure and verify NTP operating in a client/server mode

References:

[Cisco > Support > Product Support > Switches > Cisco Nexus 6000 Series Switches > Configure > Configuration Guides > Cisco Nexus 6000 Series NX-OS System Management Configuration Guide, Release 7.x > Chapter: Configuring NTP](#)

QUESTION 206

Which Cisco IOS command allows you to change the setting of the configuration register?

- A. boot config
- B. configuration-register edit
- C. config-register
- D. edit configuration-register

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The config-register command is used to change the setting of the configuration register. The configuration register has the boot field setting, which specifies the order in which the router should look for bootstrap information. The router contains a 16-bit software register, which is stored in the non-volatile random access memory (NVRAM). The config-register command is used to modify the default configuration register. The most common use of changing this register is to instruct the router to ignore the stored configuration file and boot as a new router with no configuration. This process is normally used when a router has a password that is not known and must be reset. For security purposes, this procedure can only be performed from the console connection, which means it requires physical access to the router.

Normally the setting of this register is 0x2102, which tells the router to look for a configuration file. If the file exists, it will use it. If none exists, the router will boot into ROM and present the user with a menu-based setup. This would be the default behavior for a new router as well.

To view the value of the configuration register, use the show version command as displayed below. The register setting can be seen at the bottom of the output in bold.

```
Cisco IOS Software, 3600 Software (C3660-I-M), Version 12.3(4)T
TAC Support: http://www.cisco.com/tac
Copyright (c) 1986-2003 by Cisco Systems, Inc.
Compiled Thu 18-Sep-03 15:37 by ccai
ROM: System Bootstrap, Version 12.0(6r)T, RELEASE SOFTWARE (fc1)
ROM:
C3660-1 uptime is 1 week, 3 days, 6 hours, 41 minutes
System returned to ROM by power-on
System image file is "slot0:tftpboot/c3660-i-mz.123-4.T"

Cisco 3660 (R527x) processor (revision 1.0) with 57344K/8192K bytes of memory.
Processor board ID JAB055180FF
R527x CPU at 225Mhz, Implementation 40, Rev 10.0, 2048KB L2 Cache

3660 Chassis type: ENTERPRISE
2 FastEthernet interfaces
4 Serial interfaces
DRAM configuration is 64 bits wide with parity disabled.
125K bytes of NVRAM.
16384K bytes of processor board System flash (Read/Write)

Flash card inserted. Reading filesystem...done.
20480K bytes of processor board PCMCIA Slot0 flash (Read/Write)
```

Configuration register is 0x2102

To change this setting would require issuing these commands, followed by a restart:

```
Router#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)#config
Router(config)#config-register 0x2142
```

By setting register to 0x2142, the router will ignore a configuration file at reboot if it exists. The router will then enter setup mode and prompt for you to enter initial system configuration information, as would happen with a new router. This enables the user to bypass an unknown password, since the password is contained in the file.

The boot config command is incorrect because this command is used to set the device where the configuration file is located (flash, slot, etc.) and file name for the configuration file, which helps the router to configure itself during startup.

The configuration-register edit command and the edit configuration-register commands are incorrect because

they are not valid Cisco IOS commands.

Objective:

Infrastructure Management

Sub-Objective:

Perform device maintenance

References:

[Cisco > Support > Routers > Cisco 10000 Series Routers > Troubleshoot and Alerts > Troubleshooting TechNotes > Use of the Configuration Register on All Cisco Routers > Document ID: 50421](#)

QUESTION 207

Examine the partial output from two adjacent routers:

```
RTR78# show ip ospf
Routing Process 201 with ID 192.0.2.1 VRF default
  Stateful High Availability enabled
  Graceful-restart is configured
  Grace period: 60 state: Inactive
  Last graceful restart exit status: None
  Supports only single TOS(TOS0) routes
  Supports opaque LSA
  This router is an autonomous system boundary
Administrative distance 110
Reference Bandwidth is 40000 Mbps
Initial SPF schedule delay 3000.000 msec,
minimum inter SPF delay of 2000.000 msec,
maximum inter SPF delay of 4000.000 msec
Initial LSA generation delay 3000.000 msec,
```

```
RTR79# show ip ospf
Routing Process 202 with ID 192.0.2.1 VRF default
  Stateful High Availability enabled
  Graceful-restart is configured
  Grace period: 60 state: Inactive
  Last graceful restart exit status: None
  Supports only single TOS(TOS0) routes
  Supports opaque LSA
  This router is an autonomous system boundary
Administrative distance 110
Reference Bandwidth is 30000 Mbps
Initial SPF schedule delay 3000.000 msec,
minimum inter SPF delay of 2000.000 msec,
maximum inter SPF delay of 4000.000 msec
Initial LSA generation delay 3000.000 msec,
```

Which of the following statements describes why the two routers are NOT forming an OSPF neighbor adjacency?

- A. The process IDs do not match
- B. The router IDs are misconfigured
- C. The distance is misconfigured
- D. The reference bandwidth does not match

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The output shows that the router IDs for RTR78 and RTR79 are the same value, which should not be the case. One of the two routers has been misconfigured with the other router's ID. This will prevent an OSPF neighbor adjacency from forming.

Other issues that can prevent an adjacency are:

- Mismatched OSPF area number
- Mismatched OSPF area type
- Mismatched subnet and subnet mask
- Mismatched OSPF HELLO and dead timer values

The process IDs do not have to match. It does not matter whether they match or do not match because the process ID is only locally significant on the device.

The administrative distance is not misconfigured in the output. Both routers are using the default OSPF administrative distance of 110.

If the reference bandwidths do not match, it will affect the calculation of the path cost, but it will not prevent an adjacency from forming.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot single area and multi-area OSPFv2 for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub, virtual-link, and LSAs)

References:

[Support > Technology Support > IP > IP Routing > Troubleshoot and Alerts > Troubleshooting TechNotes > OSPF Neighbor Problems Explained](#)

QUESTION 208

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.

Objective:

Routing Fundamentals

Sub-Objective:

Compare and contrast interior and exterior routing protocols

References:

[Cisco > Articles > Cisco Certification > CCNP > Shooting Trouble with IP](#)

[Cisco > Internetworking Technology Handbook > Open Shortest Path First \(OSPF\)](#)

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

QUESTION 209

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 DHCP OFFER, DHCP ACK, and DHCP NAK, 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.

Objective:

Infrastructure Security

Sub-Objective:

Describe common access layer threat mitigation techniques

References:

[Home > Support > Product Support > Switches > Cisco Catalyst 4500 Series Switches > Configure > Configuration Guides > Chapter: Configuring DHCP Snooping and IP Source Guard > Configuring DHCP Snooping on the Switch](#)

QUESTION 210

You execute the ping command from a host, but the router does not have a path to its destination.

Which of the following ICMP message types will a client receive from the router?

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

destination message:

1. Network unreachable
2. Host unreachable
3. Protocol unreachable
4. Port unreachable
5. Fragmentation needed and Don't Fragment (DF) bit set
6. Source route failed

An ICMP redirect message would not be received. This type of response is received when the router is configured to direct clients to a different router for better routing.

An ICMP time-exceeded message would not be received. This type of response occurs when the router successfully sent the packet but did not receive an answer within the allotted time; in other words, the time-to-live of the ICMP packet has been exceeded.

An ICMP echo-reply message would not be received. This would be the response received if the destination received the ping command and responded successfully.

Objective:

Routing Fundamentals

Sub-Objective:

Troubleshoot basic Layer 3 end-to-end connectivity issues

References:

[Cisco > Internetworking Technology Handbook > Internet Protocols \(IP\) > Internet Control Message Protocol \(ICMP\)](#)

QUESTION 211

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].

Objective:

WAN Technologies

Sub-Objective:

Configure and verify PPP and MLPPP on WAN interfaces using local authentication

References:

[Cisco > Support > WAN > Point-to-Point Protocol \(PPP\) > Design Technotes > Understanding and Configuring PPP CHAP Authentication > Document ID: 25647](#)

QUESTION 212

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 the 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.

Objective:

Infrastructure Management

Sub-Objective:

Troubleshoot network connectivity issues using ICMP echo-based IP SLA

References:

[Home > Support > Technology support > IP > IP application services > Technology information > Technology white paper > Cisco IOS IP Service Level Agreements User Guide](#)
[Cisco IOS IP SLAs Command Reference > icmp-echo through probe-packet priority > ip sla](#)
[Cisco IOS IP SLAs Command Reference > icmp-echo through probe-packet priority > ip sla schedule](#)
[Cisco > Cisco IOS IP SLAs Command Reference > icmp-echo](#)

QUESTION 213

Which of the following statements is NOT true regarding flow control?

- A. It determines the rate at which the data is transmitted between the sender and receiver.
- B. It can help avoid network congestion.
- C. It manages the data transmission between devices.
- D. It uses a cyclic redundancy check (CRC) to identify and remove corrupted data.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

It is NOT true that flow control uses a cyclic redundancy check (CRC) to identify and remove corrupted data. CRC is an error-checking schema that checks and removes corrupted data. It is a calculation that is performed at the source. Flow control uses CRC to identify corrupted data for the purpose of requesting re-transmission, but it does not use CRC to remove the corrupted data from the packet. If corruption is detected, the entire packet will be dropped.

Flow control is a function that ensures that a sending device does not overwhelm a receiving device. The following statements are TRUE regarding flow control:

- Flow control controls the amount of data that the sender can send to the receiver.
- Flow control determines the rate at which the data is transmitted between the sender and receiver.

- Flow control of certain types can aid in routing data around network congestion

Types of flow control include windowing, buffering, and congestion avoidance:

- Windowing- a process whereby the sender and receiver agree to increase or decrease the number of packets received before an acknowledgment is required based on network conditions. This packet number is called a window. When conditions are favorable, the window size will be increased. During unfavorable network conditions, it will be decreased.
- Buffering- the ability of a network card to store data received but not yet processed in a buffer (memory). This enhances its ability to handle spikes in traffic without dropping any data.
- Congestion avoidance - a process that some routing protocols can perform by adding information in each frame that indicates the existence of congestion on the network, allowing the router to choose a different routing path based on this information.

Objective:

Network Fundamentals

Sub-Objective:

Compare and contrast OSI and TCP/IP models

References:

[Cisco Documentation > Internetworking Technology Handbook > Routing Basics > Internet Protocols > TCP Packet Format](#)

QUESTION 214

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

```
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 shows 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.

Objective:

Infrastructure Services

Sub-Objective:

Configure, verify, and troubleshoot basic HSRP

References:

[Cisco > Cisco IOS IP Application Services Command Reference > show ip sockets through standby name > show standby](#)

[Cisco > Cisco IOS IP Application Services Configuration Guide, Release 12.4 > Part 1: First Hop Redundancy Protocols > Configuring HSRP](#)

QUESTION 215

Which of the following is NOT a mode of Dynamic Trunking Protocol (DTP)?

- A. dynamic auto
- B. dynamic trunk
- C. dynamic desirable
- D. nonegotiate

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Dynamic trunk is not a DTP mode. DTP is a Cisco proprietary trunk negotiation protocol and is used to determine if two interfaces on connected devices can become a trunk. There are five modes of DTP:

- Trunk: Puts the interface into permanent trunking mode and negotiates to convert the link into a trunk link. The interface becomes a trunk interface even if the neighboring interface does not agree to the change.
- Access: Puts the interface into permanent nontrunking mode and negotiates to convert the link into a nontrunk link. The interface becomes a nontrunk interface even if the neighboring interface does not agree to the change.
- Dynamic desirable: Makes the interface actively attempt to convert the link to a trunk link. The interface becomes a trunk interface if the neighboring interface is set to trunk, desirable, or auto mode.
- Dynamic auto: Makes the interface willing to convert the link to a trunk link. The interface becomes a trunk interface if the neighboring interface is set to trunk or desirable mode. This is the default mode for all Ethernet interfaces in Cisco IOS.
- Nonegotiate: Puts the interface into permanent trunking mode but prevents the interface from generating DTP frames. You must configure the neighboring interface manually as a trunk interface to establish a trunk link. Use this mode when connecting to a device that does not support DTP.

If one side's mode of link is in trunk mode, dynamic desirable mode, or dynamic auto mode, and the other side is trunk or dynamic desirable, a trunk will form. Nonegotiate mode enables trunking but disables DTP.

Objective:

LAN Switching Fundamentals

Sub-Objective:

Configure and verify Layer 2 protocols

References:

QUESTION 216

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

- A. Authentication Header (AH) in tunnel mode
- B. Authentication Header (AH) in transport mode
- C. Encapsulating Security Payload (ESP) in tunnel mode
- D. Encapsulating Security Payload (ESP) in transport mode

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

- Authentication Header (AH) uses protocol 51.
- ESP uses protocol 50.

AH is defined in Request for Comments (RFC) 1826 and 2402. AH does not perform data encryption and therefore, information is passed as clear text. The purpose of AH is to provide data integrity and authentication, and anti-reply service (optional). It ensures that a packet that crosses the tunnel is the same packet that left the peer device and no changes have been made. It uses a keyed hash to accomplish this.

ESP is defined in RFC 2406. ESP can provide data integrity and authentication, but its primary purpose is to encrypt data crossing the tunnel. There are two reasons why ESP is the preferred building block of IPSec tunnels:

- The authentication component of ESP does not include any Layer 3 information. Therefore, this component can work in conjunction with a network using Network Address Translation (NAT).
- On Cisco devices, ESP supports encryption using Advanced Encryption Standard (AES), Data Encryption Standard (DES), or Triple DES (3DES).

Tunnel mode is used between Virtual Private Network (VPN) gateways such as routers, firewalls, and VPN concentrators.

Transport mode is used between end-stations or between an end-station and a VPN gateway.

The options AH in tunnel mode and AH in transport mode are incorrect because AH does not provide encryption.

The option ESP in transport mode is incorrect because transport mode is used between end-stations or between an end-stations and a VPN gateway.

Objective:

WAN Technologies

Sub-Objective:

Describe WAN access connectivity options

References:

[Cisco > Articles > Network Technology > General Networking > IPSec Overview Part Two: Modes and Transforms](#)

[Cisco > The Internet Protocol Journal > The Internet Protocol Journal - Volume 3, No. 1, March 2000 > IP Security](#)

QUESTION 217

If a routing table contains multiple routes for the same destination, which were inserted by the following

methods, which route will the router use to reach the destination network?

- A. The route inserted by RIP
- B. The route inserted by OSPF
- C. The route inserted by BGP
- D. The route configured as a static route

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A static route will be preferred because it has the lowest administrative distance. Routing protocols are dynamic routing methods. With the default configuration, static routes are preferred over dynamic routes.

The default administrative distance for the offered options is:

- RIP 120
- OSPF 110
- eBGP 20
- Static 1

When Routing Information Protocol (RIP), Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), and static routing is enabled on a router, the router will prefer the static route.

Objective:

Routing Fundamentals

Sub-Objective:

Interpret the components of routing table

References:

[Cisco Documentation > Internetworking Technology Handbook > Routing Basics](#)

QUESTION 218

Which Cisco IOS command is used to view the information about the interfaces on which Cisco Discovery Protocol (CDP) is enabled?

- A. show cdp interface
- B. show interfaces
- C. show cdp
- D. show cdp interfaces

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The show cdp interface command is used to view the information about the interfaces on which Cisco Discovery Protocol (CDP) is enabled.

The syntax of the command is as follows:

Router# show cdp interface [type number]

The parameters of the command are as follows:

type: specifies the type of interface for which information is required

number: specifies the number of interfaces for which information is required

The output of the show cdp interface command is as follows:

```
Router#show cdp interface
Serial0 is up, line protocol is up, encapsulation is SMDS
Sending CDP packets every 100 seconds
Holdtime is 300 seconds
Serial1 is up, line protocol is up, encapsulation is SMDS
Sending CDP packets every 60 seconds
Holdtime is 180 seconds
Ethernet0 is up, line protocol is up, encapsulation is ARPA
Sending CDP packets every 120 seconds
Holdtime is 360 seconds
```

The show interfaces command is incorrect because this command is used to view configured interfaces on the router. The output of this command can be very useful, especially when troubleshooting a connection with no connectivity. Consider the output of the command on the following two routers that are connected with a serial interface:

```
NewYork#show interfaces s0
Serial0 is up, line protocol is up
Hardware is HD64570
Internet Address is 192.168.10.1/24
MTU 1500 bytes,BW 1544 Kbit
Reliability 255/255
Encapsulation HDLC, loopback not set
Keepalive set (10 sec)
```

```
LosAngeles#show interfaces s1
Serial0 is up, line protocol is up
Hardware is HD64570
Internet Address is 192.168.11.2/24
MTU 1500 bytes,BW 56000 Kbit
Reliability 255/255
Encapsulation HDLC, loopback not set
Keepalive set (10 sec)
```

Notice that the following settings are correct:

- The encapsulation matches (HDLC)
- The physical connection is good (indicated by Serial0 is up)

Notice, however, that the IP addresses 192.168.10.1 and 192.168.11.2 are NOT in the same subnet when using a 24-bit mask. With a 24-bit mask, the two addresses should agree through the first three octets, and these do not. Problems such as this can be located through inspection of the output produced by the show interfaces command.

The show cdp command is incorrect because this command is used to view the global CDP information.

The show cdp interfaces command is incorrect because this command does not exist in the Cisco command reference. There is a show cdp interface command, which displays CDP activity on a per-interface basis.

Objective:

LAN Switching Fundamentals

Sub-Objective:


Configure and verify Layer 2 protocols

References:

[Cisco > Cisco IOS Network Management Command Reference > show cdp interface](#)

QUESTION 219

You have three EIGRP routers that are connected as shown in the diagram below.



Router A and Router C do not seem to be exchanging information. You execute commands on all three routers, and receive as output the information shown below:

```
Router A# show ip eigrp neighbors
Address Interface holdtime uptime Q Seq SRTT RTO
192.168.10.2 S1 13 1:20:10 100 458 0 30

routerA# show run
<output omitted>
router eigrp 56
network 192.168.10.0
network 192.168.9.0
no auto-summary

routerC# show run
<output omitted>
router eigrp 56
network 192.168.11.0
no auto-summary
```

What needs to be done to make Routers A and C start exchanging information?

- A. Execute the auto-summary command on Router A
- B. Execute the network 192.168.9.0 command under EIGRP 56 on Router C
- C. Correct the IP address on the S1 interface of Router C
- D. Recreate the EIGRP configuration on Router C as EIGRP 55

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Router C is not displayed in the neighbor table of Router A, which indicates that Router C and Router A are not forming a neighbor relationship or exchanging information. This is because Router C does not have EIGRP configured for its S1 interface. You can see this is missing from its configuration in the output of the show run command for RouterC. To solve the issue, you should execute the network 192.168.9.0 command under the

EIGRP 56 configuration on Router C. Then Router C will start sending hellos on that interface and the two routers will become neighbors.

The `show ip eigrp neighbors` command displays the following information for each EIGRP neighbor. In parentheses is the value of each found in the output of router A for Router B:

```
IP address (192.168.10.2)
Local interface (S1)
Retransmit interval (13)
Queue count (100)
```

There is no need to execute the `auto-summary` command on Router A. It will not affect the establishment of a neighbor relationship between Routers A and C.

There is no need to correct the IP address on the S1 interface of Router C. The address 192.168.9.1 is correctly located in the same subnet as the address on S0 of Router A.

Finally, changing the EIGRP configuration on Router C to EIGRP 55 will not help. Router C will not start sending hellos on its S1 interface until EIGRP is enabled on the S1 interface. Until then, the Routers A and C will not form a neighbor relationship and will not share information.

Objective:

Routing Fundamentals

Sub-Objective:

Configure, verify, and troubleshoot EIGRP for IPv4 (excluding authentication, filtering, manual summarization, redistribution, stub)

References:

[Cisco > Cisco IOS IP Routing Configuration Guide, Release 12.4 > Configuring EIGRP > Enabling EIGRP](#)

QUESTION 220

You are the network administrator for your company. You recently configured Cisco Discovery Protocol (CDP) in the network. You want to view output regarding all of the neighboring devices discovered by CDP. This information should include network address, enabled protocols, and hold time.

Which Cisco Internetwork Operating System (IOS) command would allow you to accomplish this task?

- A. `show cdp`
- B. `show cdp entry`
- C. `show cdp neighbor entries`
- D. `show cdp neighbors detail`

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

In this scenario, you should use the `show cdp neighbors detail` command to view the details of the neighboring devices that were discovered by CDP. CDP is a Layer 2 (data link layer) protocol used to find information about neighboring network devices. The `show cdp neighbors detail` command is used to view details such as network address, enabled protocols, and hold time. The complete syntax of this command is:

`show cdp neighbors [type number] [detail]`

The command parameters are defined in this way:

type: An optional parameter which specifies the type of interface used to connect to the neighbors for which you require information.

number: An optional parameter used to specify the interface number connected to the neighbors for which you want information.

detail: An optional parameter used to get detailed information about neighboring devices, such as network address, enabled protocols, software version and hold time.

The following code is a sample partial output of the show cdp neighbors detail command:

```
Device ID: RTR2511
Entry address(es):
IP address: 178.10.20.1
Platform: cisco 2511, Capabilities: Router
Interface Serial 0
Holdtime : 123 sec
<output omitted>
```

```
-----
Device ID: RTR2611-Edge
Entry address(es):
IP address: 10.10.1.2
Platform: cisco 2611, Capabilities: Router
Interface Ethernet 0
Holdtime : 123 sec
<output omitted>
```

The show cdp command is incorrect because this command is used to view global CDP information such as the timer and hold time.

The show cdp entry command is incorrect because this command is used to view information about a specific neighboring device.

The show cdp neighbor entries command is incorrect because this is not a valid Cisco IOS command.

Objective:

LAN Switching Fundamentals

Sub-Objective:


Configure and verify Layer 2 protocols

References:

[Cisco > Cisco IOS Network Command Reference, Release 12.4 > show cdp neighbors](#)

QUESTION 221

How many collision domains are in a LAN with four hubs and two bridges that are connected directly to each other, as shown in the following figure? (Click the Exhibit(s) button.)



- A. four
- B. five
- C. six
- D. fourteen

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

A bridge segments the LAN into separate collision domains. The figure in this scenario has five segments created between these two bridges. Therefore, there will be five collision domains (segments) on the LAN if the two bridges are directly connected as shown in the exhibit. Hubs do not create LAN segments; they act as port aggregators and signal amplifiers.

It is also worth noting that with no router in the diagram, the entire network is a single broadcast domain. If a router were present, each of its interfaces could host a different subnet and each of those same interfaces would be a separate broadcast domain.

Objective:

LAN Switching Fundamentals

Sub-Objective:


Describe and verify switching concepts

References:

[Internetwork Design Guide -- Designing Switched LAN Internetworks > Comparison of LAN Switches and Routers](#)

QUESTION 222

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.

Objective:

Routing Fundamentals

Sub-Objective:

Compare and contrast distance vector and link-state routing protocols

References:

QUESTION 223

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.

Objective:

WAN Technologies

Sub-Objective:

Configure, verify, and troubleshoot PPPoE client-side interfaces using local authentication

References:

QUESTION 224

With which type of service is bandwidth and latency the biggest consideration?

- A. streaming video
- B. telnet sessions
- C. FTP transfers
- D. authentication traffic

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Streaming video places the largest demand on both bandwidth and latency. Video traffic is real-time and benefits from dedicated bandwidth with QoS implementation to ensure quality. Moreover, this service can tolerate very little latency.

Telnet and FTP sessions are both low bandwidth users and can tolerate a high degree of latency since the data can be reassembled when all pieces arrive, which is not possible when data is coming in real-time, and waiting for retransmissions and reassembly is not feasible.

Authentication traffic is not sensitive to latency and does not require much bandwidth either.

Objective:

WAN Technologies

Sub-Objective:

Describe basic QoS concepts

References:

[Cisco Documentation > Internetworking Technology Handbook > Voice/Data Integration Technologies](#)

QUESTION 225

You are configuring a serial link between a Cisco router and a router produced by another vendor. What would be the advantages of using Point to Point Protocol (PPP) over High Level Data Link Control (HDLC) in this scenario?

- A. HDLC has a proprietary "type" field that may be incompatible with equipment from other vendors.
- B. HDLC is not available on non-Cisco routers.
- C. PPP is faster.
- D. PPP performs error checking.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

High Level Data Link Control (HDLC) has a proprietary "type" field that may be incompatible with equipment from other vendors. It is recommended that PPP always be used when combining equipment from multiple vendors because this Data Link layer WAN protocol is an industry standard. PPP is implemented in the same manner on all PPP-capable equipment.

HDLC is available on non-Cisco routers. However, the Cisco implementation has a "type" field that may prevent

the connection from working.

PPP is not faster than HDLC.

PPP performs error checking, but so does HDLC.

Objective:

WAN Technologies

Sub-Objective:

Describe WAN access connectivity options

References:

[Cisco > Internetworking Technology Handbook > Point to Point Protocol \(PPP\)](#)

QUESTION 226

What Cisco IOS command produced the following as a part of its output?

```
Port Security: Enabled
Port status: SecureUp
Violation mode: Shutdown
Maximum MAC Addresses: 2
Total MAC Addresses: 2
Configured MAC Addresses: 2
Aging Time: 30 mins
Aging Type: Inactivity
SecureStatic address aging: Enabled
Security Violation count: 0
```

- A. show interfaces port-security
- B. show port-security interface
- C. show ip interface
- D. show interfaces switchport

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The output is the result of executing the show port-security interface command. The sample output indicates that port security has been enabled on the interface, and that a maximum of two MAC addresses has been configured. A violation mode of Shutdown indicates that if a third MAC address attempts to make a connection, the switch port will be disabled. It is useful to note that you must specify a port number when you execute the command. In this case, the command was Switch# show port-security interface fastethernet0/1.

The output was not produced by the show interfaces port-security command. This is not a valid Cisco command.

The output was not produced by the show ip interface command. It displays protocol-related information about an interface, and nothing pertaining to switch port security. An example of its output follows:

```
Router# show ip interface fastethernet0/1
fastethernet0/1 is up, line protocol is up
Internet address is 10.1.1.1/16
Broadcast address is 255.255.255.255
Address determined by setup command
MTU is 1500 bytes
Helper address is not set
Directed broadcast forwarding is disabled
Outgoing access list is not set
Inbound access list is not set
Proxy ARP is enabled
Local Proxy ARP is disabled
Security level is default
Split horizon is enabled
ICMP redirects are always sent
ICMP unreachable are always sent
ICMP mask replies are never sent
IP fast switching is enabled
IP fast switching on the same interface is disabled
IP Flow switching is disabled
IP CEF switching is enabled
IP Feature Fast switching turbo vector
IP VPN Flow CEF switching turbo vector
IP multicast fast switching is enabled
IP multicast distributed fast switching is disabled
IP route-cache flags are Fast, CEF
Router Discovery is disabled
IP output packet accounting is disabled
IP access violation accounting is disabled
TCP/IP header compression is disabled
RTP/IP header compression is disabled
Policy routing is enabled, using route map PBR
Network address translation is disabled
BGP Policy Mapping is disabled
IP Multi-Processor Forwarding is enabled <===== MPF information
IP Input features, "PBR",
are not supported by MPF and are IGNORED
IP Output features, "NetFlow",
are not supported by MPF and are IGNORED
```

The output was not produced by the show interfaces switchport command. This command displays non-security related switch port information, such as administrative and operational status and trunking:


```
Cat2950# show interfaces fastethernet0/1 switchport
Name: Po1
Switchport: Enabled
Administrative Mode: trunk
Operational Mode: trunk
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 10 (VLAN0010)
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001

Protected: false
Voice VLAN: none (Inactive)
Appliance trust: none
```

Objective:

Infrastructure Security

Sub-Objective:

Configure, verify, and troubleshoot port security

References:

[Cisco > Support > Cisco IOS Security Command Reference: Commands S to Z > show port-security](#)

QUESTION 227

Which of the following commands helps you determine the Layer 1 and Layer 2 up/down status of a Cisco interface?

- A. show controllers
- B. show running-config
- C. show interfaces trunk
- D. show interfaces

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The show interfaces command displays the Layer 1 and Layer 2 operational status of an interface, along with other information.

```

Router# show interfaces
Ethernet 1 is up, line protocol is up
Hardware is MCI Ethernet, address is 0000.0c00.750c (via 0000.0c00.750c)
Internet address is 205.108.28.8, subnet mask is 255.255.0.0
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 10:09*: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, 586 runts, 705 giants
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
Input queue: 0/75/0 (size/max/drops); Total output drops: 0
Output queue: 7/64/0 (size/threshold/drops)
Conversations 2/9 (active/max active)

```

Of interest in this output is the information contained on line 14 (also shown below). The figures for Runts and Giants (packets that are either too large or too small) indicate that collisions are occurring or that the NIC is malfunctioning:

```
Received 354125 broadcasts, 586 runts, 705 giants
```

As a part of troubleshooting this increase in collisions, you can also identify the speed of the interface by reading line 4, which says the bandwidth is 10000 Kbit, indicating a FastEthernet interface.

```
MTU 1500 bytes, BW 10000 Kbit, DLY 100000 usec, rely 255/255, load 1/255
```

You can also specify a particular interface for which information should be displayed, as shown below:

```

Router# show interfaces ethernet 0/0
Ethernet0/0 is administratively down, line protocol is down
Hardware is AmdP2, address is 0003.e39b.9220 (via 0003.e39b.9220)
Internet address is 10.1.0.254/16
MTU 1500 bytes, BW 10000 Kbit, DLY 1000 usec,
reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA, loopback not set
Keepalive set (10 sec)
ARP type: ARPA, ARP Timeout 04:00:00
<<output omitted>>

```

The sample output indicates that interface Ethernet 0/0 is in an administratively down, line protocol is down state. The first statement indicates the Layer 1 (Physical) status, while the second statement indicates the Layer 2 (Data Link) status of the interface. A status of administratively down always indicates that the interface is in a shutdown state; the interface can be activated by executing the command `no shutdown`. This command also indicates the configured bandwidth of the interface (10000 Kbit in this case).

The following lines of information concern the Physical layer:

```
Hardware is AmdP2, address is 0003.e39b.9220 (via 0003.e39b.9220)
Ethernet0/0 is administratively down, line protocol is down
```

This output indicates that the link has not been enabled. There are other combinations of up and down states that can indicate other conditions. For example, the following indicates the link is functioning:

Ethernet0/0 is up, line protocol is up

The output below indicates a problem at the other end of the link, perhaps meaning that the interface on the other end has not been enabled or that the port to which it is connected has been disabled.

Ethernet0/0 is up, line protocol is down (not connect)

The following lines of information concern the Data Link layer:

```
Encapsulation ARPA
line protocol is down
```

The show controllers command provides Layer 1 information only, including the type of cable detected (DTE/DCE) on a serial interface.

The show running-config command displays the current active configuration of the router, but does not indicate the operational status of its interfaces.

The show interfaces trunk command will not show the Layer 1 and Layer 2 up/down status of a Cisco interface. It will show all interfaces configured to be trunks. This command is very useful when you need to locate trunk interfaces on a switch with which you are not familiar. In the output of the command below, the three trunking interfaces are the Fa0/3, Fa0/9, and Fa0/12 interfaces.

```
Switch# show int trunk
Port Mode          Encapsulation    native vlan
Fa0/3              on                802.1q          1
Fa0/9              desirable        802.1q          1
Fa0/12             desirable        802.1q          1
```

Objective:

LAN Switching Fundamentals

Sub-Objective:

Troubleshoot interface and cable issues (collisions, errors, duplex, speed)

References:

[Cisco > Catalyst 6500 Series Cisco IOS Command Reference, 12.1 E > show bootvar to show ip cache > show interfaces](#)

QUESTION 228

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).

Objective:

Routing Fundamentals

Sub-Objective:

Interpret the components of routing table

References:

[Cisco > Support > IP > IP Routing](#)

QUESTION 229

Examine the output shown below:

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

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

What is true of this configuration?

- A. The link-local address of R1 is FE80::2
- B. The link-local address of R1 is FE80::1
- C. The area ID is 1
- D. No adjacency has formed

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The output shows that the link-local address of R1 is FE80::1. R1's link-local address appears in the output of R2 because the show ipv6 eigrp neighbors command displays information about the neighbor, not the local router.

The link-local address of R1 is not FE80::2. That is the link-local address of R2.

Because the area ID is not displayed in the output, we do not know its value. The only 1 in the output is the value representing the process ID of both routers, IPv6-EIGRP neighbors for process 1.

It is not true that no adjacency has formed. There is an adjacency present; if there were not, the two routers would not appear in each other's output of the show ipv6 eigrp neighbors command.

Objective:

Routing Fundamentals

Sub-Objective:

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

References:

[Home > Support > Cisco IOS IPv6 Command Reference > show ipv6 eigrp neighbors](#)

QUESTION 230

DRAG DROP

Click and drag the network components and functions to their corresponding descriptions on the right.

Select and Place:

Components:
TCP/IP
Router
Layer 2 switching
Hierarchical model

Descriptions:	
	Performed using a destination MAC address within a frame
	Provides a framework for designing internetworks in layers
	A suite of protocols used to transmit data
	Separates broadcast domains and connects different networks

Correct Answer:

Components:

Descriptions:	
Layer 2 switching	Performed using a destination MAC address within a frame
Hierarchical model	Provides a framework for designing internetworks in layers
TCP/IP	A suite of protocols used to transmit data
Router	Separates broadcast domains and connects different networks

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The following network components and functions should be matched to these corresponding descriptions:

- Transmission Control Protocol (TCP)/Internet Protocol (IP): TCP/IP is a suite of data communication protocols.
- Router: Separates broadcast domains while connecting different networks. Routers also provide a medium for connecting Local Area Network (LAN) and Wide Area Network segments.
- Layer 2 switching: Performed using a destination MAC address within a frame. In Layer 2 switching, switching is based on Media Access Control (MAC) addresses.
- Hierarchical model: Enables the designing of internetworks into layers. There are three layers in the hierarchical network design:
 - Core layer: Provides high-speed data transfer between sites.
 - Distribution layer: Includes LAN-based routers and Layer 3 switches and enables routing between Virtual Local Area Networks (VLANs).
 - Access layer: Provides workgroup and end-user access, and is also referred to as the desktop layer.

Objective:

Network Fundamentals

Sub-Objective:

Describe the impact of infrastructure components in an enterprise network

References:

[Cisco wiki > Main Page > Internet Protocols](#)

[Cisco wiki > Internetwork Design Guide -- Designing Switched LAN Internetworks > General Network Design Principles > 5.6.1 Figure: Hierarchical network design model](#)