Exam : 642-812

Title : Building Converged Cisco Multilayer Switched Networks (BCMSN)

Ver : 10.05.07

QUESTION 1
Certkiller uses layer 3 switches in the Core of their network. Which method of Layer 3 switching uses a forwarding information base (FIB)?

A. Topology-based switching
B. Demand-based switching
C. Route caching
D. Flow-based switching
E. None of the above

Answer: A

Explanation:
The Layer 3 engine (essentially a router) maintains routing information, whether from static routes or dynamic routing protocols. Basically, the routing table is reformatted into an ordered list with the most specific route first, for each IP destination subnet in the table. The new format is called a Forwarding Information Base (FIB) and contains routing or forwarding information that the network prefix can reference.
In other words, a route to 10.1.0.0/16 might be contained in the FIB, along with routes to 10.1.1.0/24 and 10.1.1.128/25, if those exist. Notice that these examples are increasingly more specific subnets. In the FIB, these would be ordered with the most specific, or longest match, first, followed by less specific subnets. When the switch receives a packet, it can easily examine the destination address and find the longest match entry in the FIB. The FIB also contains the next-hop address for each entry. When a longest match entry is found in the FIB, the Layer 3 next-hop address is found, too.

QUESTION 2
You need to design the VLAN scheme for the Certkiller network. Which two statements are true about best practices in VLAN design? (Select two)

A. Routing should occur at the access layer if voice VLANs are utilized. Otherwise, routing should occur at the distribution layer.
B. Routing should always be performed at the distribution layer.
C. VLANs should be localized to a switch.
D. VLANs should be localized to a single switch unless voice VLANs are being utilized.
E. Routing should not be performed between VLANs located on separate switches.

Answer: B, C

Explanation:
In the distribution layer, uplinks from all access layer devices are aggregated, or come together. The distribution layer switches must be capable of processing the total volume of traffic from all the connected devices. These switches should have a port density of high-speed links to support the collection of access layer switches.
VLANs and broadcast domains converge at the distribution layer, requiring routing,

filtering, and security. The switches at this layer must be capable of performing multilayer switching with high throughput. Only certain Catalyst switch models can provide multilayer switching; be sure to understand which ones can do this.
A switched environment offers the technology to overcome flat network limitations. Switched networks can be subdivided into VLANs. By definition, a VLAN is a single broadcast domain. All devices connected to the VLAN receive broadcasts from other VLAN members. However, devices connected to a different VLAN will not receive those same broadcasts. (Naturally, VLAN members also receive unicast packets directed toward them from other VLAN members.)
A VLAN consists of defined members communicating as a logical network segment. In contrast, a physical segment consists of devices that must be connected to a physical cable segment. A VLAN can have connected members located anywhere in the campus network, as long as VLAN connectivity is provided between all members. Layer 2 switches are configured with a VLAN mapping and provide the logical connectivity between the VLAN members.

QUESTION 3
The Certkiller la network is shown below:

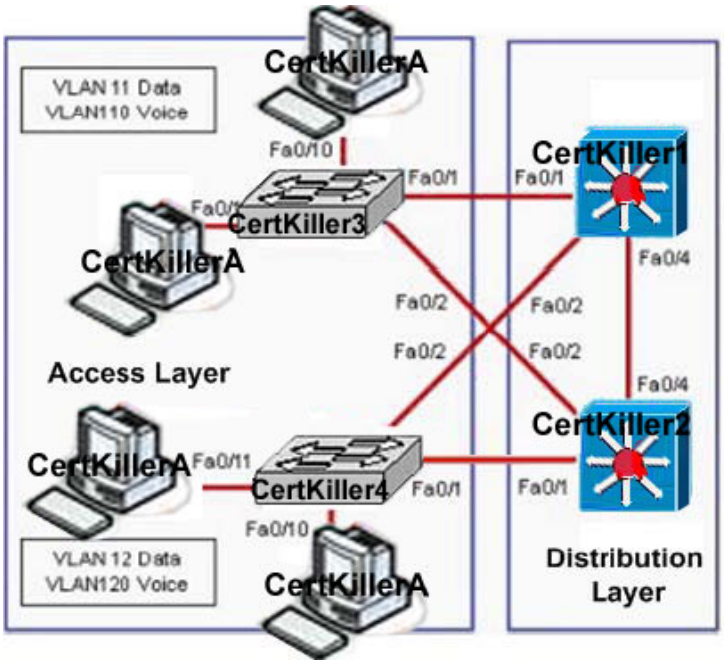

On the basis of the information provided in the exhibit above, which two sets of procedures are best practices for Layer 2 and 3 failover alignment? (Select two)

A. Configure the Certkiller 1 switch as the standby HSRP router and the STP root for VLANs 11 and 110. Configure the Certkiller 2 switch as the standby HSRP router and the STP root for VLANs 12 and 120.
B. Configure the Certkiller 1 switch as the active HSRP router and the backup STP root for VLANs 11 and 110. Configure the Certkiller 2 switch as the active HSRP router and the backup STP root for VLANs 12 and 120.
C. Configure the Certkiller 2 switch as the active HSRP router and the STP root for all VLANs. Configure the Certkiller 1 switch as the standby HSRP router and backup STP

root for all VLANs.
D. Configure the Certkiller 1 switch as the active HSRP router and the STP root for all VLANs. Configure the Certkiller 2 switch as the standby HSRP router and backup STP root for all VLANs.
E. Configure the Certkiller 1 switch as the standby HSRP router and the backup STP root for VLANs 12 and 120. Configure the Certkiller 2 switch as the standby HSRP router and the backup STP root for VLANs 11 and 110.
F. Configure the Certkiller 1 switch as the active HSRP router and the STP root for VLANs 11 and 110. Configure the Certkiller 2 switch as the active HSRP router and the STP root for VLANs 12 and 120.

Answer: E, F

Explanation:
Basically, each of the routers that provides redundancy for a given gateway address is assigned to a common HSRP group. One router is elected as the primary, or active, HSRP router, another is elected as the standby HSRP router, and all the others remain in the listen HSRP state. The routers exchange HSRP hello messages at regular intervals, so they can remain aware of each other's existence, as well as that of the active router.
HSRP election is based on a priority value (0 to 255) that is configured on each router in the group. By default, the priority is 100. The router with the highest priority value (255 is highest) becomes the active router for the group. If all router priorities are equal or set to the default value, the router with the highest IP address on the HSRP interface becomes the active router. To set the priority, use the following interface configuration command:
Switch(config-if)# standby group priority priority
When HSRP is configured on an interface, the router progresses through a series of states before becoming active. This forces a router to listen for others in a group and see where it fits into the pecking order. The HSRP state sequence is Disabled, Init, Listen, Speak, Standby, and, finally, Active.
You can configure a router to preempt or immediately take over the active role if its priority is the highest at any time. Use the following interface configuration command to allow preemption:
Switch(config-if)# standby group preempt [delay seconds]

QUESTION 4
The Certkiller LAN switches are being configured to support the use of Dynamic VLANs. Which of the following are true of dynamic VLAN membership? (Select all that apply)

A. VLAN membership of a user always remains the same even when he/she is moved to another location.
B. VLAN membership of a user always changes when he/she is moved to another location.
C. Membership can be static or dynamic.
D. Membership can be static only.

E. None of the above.

Answer: A, C

Explanation:
Dynamic VLAN memberships are based on the users MAC address connected to the port. If you have VTP server, a VTP database file, a VTP client switch, and a dynamic port; regardless of where your physical location is, you can still remain in the same VLAN.
Incorrect Answers:
B: This was true before the use of Dynamic VLAN membership, as VLANs were assigned to ports, not users.
D: VLAN memberships can be either static or dynamic.

QUESTION 5
The Certkiller LAN switches are being configured to support the use of Dynamic VLANs. What should be considered when implementing a dynamic VLAN solution? (Select two)

A. Each switch port is assigned to a specific VLAN.
B. Dynamic VLANs require a VLAN Membership Policy Server.
C. Devices are in the same VLAN regardless of which port they attach to.
D. Dynamic VLAN assignments are made through the command line interface.

Answer: B, C

Explanation:
With VLAN Membership Policy Server (VMPS), you can assign switch ports to VLANs dynamically, based on the source Media Access Control (MAC) address of the device connected to the port. When you move a host from a port on one switch in the network to a port on another switch in the network, the switch assigns the new port to the proper VLAN for that host dynamically.
Note: There are two types of VLAN port configurations: static and dynamic.
Incorrect Answers
A: In a static VLAN, the administrator assigns switch ports to the VLAN, and the association does not change until the administrator changes the port assignment. However, this is not the case of dynamic VLANs.
D: The Command Line Interface is not used for dynamic VLAN assignments.
Reference: Cisco Online, Configuring Dynamic Port VLAN Membership with VMPS

QUESTION 6
What is the preferred method of filtering bridged traffic within a VLAN?

A. Ethernet maps
B. Router ACLs
C. VLAN access maps

D. IP ACLs

Answer: C

Explanation:
VLAN ACLs or VLAN maps access-control all packets (bridged and routed). You can use VLAN maps to filter traffic between devices in the same VLAN.
Each VLAN access map can consist of one or more map sequences, each sequence with a match clause and an action clause. The match clause specifies IP, IPX, or MAC ACLs for traffic filtering and the action clause specifies the action to be taken when a match occurs. When a flow matches a permit ACL entry, the associated action is taken and the flow is not checked against the remaining sequences. When a flow matches a deny ACL entry, it will be checked against the next ACL in the same sequence or the next sequence. If a flow does not match any ACL entry and at least one ACL is configured for that packet type, the packet is denied.
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a0080160

QUESTION 7
You are assigning VLANs to the ports of switch CK1 . What VLAN number value is an assigned to the default VLAN?

A. VLAN 1003
B. VLAN 1
C. VLAN ON
D. VLAN A
E. VLAN 0

Answer: B

Explanation: The default VLAN is VLAN 1. Although this VLAN can be modified, it can not be deleted from the switch. The following VLANs are on by default for all Cisco Catalyst switches:
VLAN 1 - Default VLAN
VLAN 1002 - Default FDDI VLAN
VLAN 1003 - Default Token Ring VLAN
VLAN 1004 - Default FDDI Net VLAN
VLAN 1005 - Default Token Ring Net VLAN
Incorrect Answers:
A: This is the default Token Ring VLAN that is installed in the switch IOS. It is seldom used.
C: ON is a VTP configuration mode, but is not a normal VLAN name.
D: Although any VLAN can be named VLAN A, it is not created by default.
E: Although in Cisco IOS the number 0 has significance (i.e. ethernet 0, console port 0, serial 0) in VLANs 1 is the default. VLAN 0 is an invalid VLAN and can not be used.

QUESTION 8
The VLANs in switch CK1 are being modified. Which of the following are updated in CK1 every time a VLAN is modified? (Select all that apply)

A. Configuration revision number
B. Configuration revision flag field
C. Configuration revision reset switch
D. Configuration revision database
E. None of the above.

Answer: A, D

Explanation:
For accountability reasons, every time a VLAN is modified the revision number changes, as does the information in the configuration revision database (as that is where the VLAN information is stored).
Incorrect Answers:
B, C: The configuration revision flag field, and the configuration revision reset switch don't exist in this context.

QUESTION 9
If you needed to transport traffic coming from multiple VLANs (connected between switches), and your CTO was insistent on using an open standard, which protocol would you use?

A. 802.11B
B. spanning-tree
C. 802.1Q
D. ISL
E. VTP
F. Q.921

Answer: C

Explanation:
The act involved in the above question is trunking. The two trunking protocols in the answer choices are: 802.1Q and ISL. ISL is Cisco proprietary and IEEE 802.1Q is based on an open standard. When non-Cisco switches are used along with Cisco switches and trunking is required, it is best to use the 802.1Q encapsulation.
Incorrect Answers:
A: This standard is used in wireless networking and has nothing to do with VLAN switching.
B: The Spanning Tree Protocol (STP) is used to prevent loops within a bridged network. Each VLAN runs a separate instance of the STP and this is enabled by default.
D: This is the alternative Cisco proprietary method of trunking.

E: VLAN Trunking Protocol (VTP) is a Layer 2 messaging protocol that manages the addition, deletion, and renaming of VLANs on a network-wide basis. It is not used to actually transport VLAN traffic.
F: This is an ISDN signalling standard and is not related with VLAN switching.

QUESTION 10
Under what circumstances should an administrator prefer local VLANs over end-to-end VLANs?

A. Eighty percent of traffic on the network is destined for Internet sites.
B. There are common sets of traffic filtering requirements for workgroups located in multiple buildings.
C. Eighty percent of a workgroup's traffic is to the workgroup's own local server.
D. Users are grouped into VLANs independent of physical location.
E. None of the above

Answer: A

Explanation:
This geographic location can be as large as an entire building or as small as a single switch inside a wiring closet. In a geographic VLAN structure, it is typical to find 80 percent of the traffic remote to the user (server farms and so on) and 20 percent of the traffic local to the user (local server, printers, and so on).
Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 93

QUESTION 11
What are some virtues of implementing end-to-end VLANs? (Choose two)

A. End-to-end VLANs are easy to manage.
B. Users are grouped into VLANs independent of a physical location.
C. Each VLAN has a common set of security and resource requirements for all members.
D. Resources are restricted to a single location.

Answer: B, C

Explanation:
In an end-to-end VLAN, users are grouped into VLANs independent of physical location and dependent on group or job function.
Each VLAN has a common set of security requirements for all members.
Incorrect Answers:
A: End to end VLANs are more difficult to manage than local VLANs, due to the physical distances that they can span.
D: In an end-to-end VLAN, network resources are generally distributed across the entire enterprise wide area network.

QUESTION 12
Which of the following statements is true about the 80/20 rule (Select all that apply)?

A. 20 percent of the traffic on a network segment should be local
B. no more than 20 percent of the network traffic should be able to move across a backbone.
C. no more than 80 percent of the network traffic should be able to move across a backbone.
D. 80 percent of the traffic on a network segment should be local

Answer: B, D

Explanation:
The 80/20 rule in network design originated from Pareto's Principle. The Italian economist Vilfredo Paretto came up with the discovery that 20% of the people controlled 80% of the wealth and applied the principle of how inputs don't match outputs in real life. Keeping this number in mind, 80% of network traffic should be local to a segment and 20% should move across the backbone.
Note: With the availability of inexpensive bandwidth and centralized data centers, this rule appears to have become obsolete. In fact, most networks have taken on the 20/80 rules, as opposed to the legacy 80/20 rule.

QUESTION 13
Which two factors give merit to the 20/80 LAN traffic model? (Select two)

A. The Internet
B. Local servers
C. Server farms
D. Localized applications
E. More powerful desktop PC's

Answer: A, C

Explanation:
Remote services (server farms, Internet, etc.) are factors which contributed to increased backbone traffic.
Also consider:
This geographic location can be as large as an entire building or as small as a single switch inside a wiring closet. In a geographic VLAN structure, it is typical to find 80 percent of the traffic remote to the user (server farms and so on) and 20 percent of the traffic local to the user (local server, printers, and so on).
Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 93

QUESTION 14
The Certkiller LAN is becoming saturated with broadcasts and multicast traffic.

What could you do to help a network with many multicasts and broadcasts?

A. Creating smaller broadcast domains by implementing VLANs.
B. Separate nodes into different hubs.
C. Creating larger broadcast domains by implementing VLANs.
D. Separate nodes into different switches.
E. All of the above.

Answer: A

Explanation:
Controlling broadcast propagation throughout the network is important to reduce the amount of overhead associated with these frames. Routers, which operate at Layer 3 of the OSI model, provide broadcast domain segmentation for each interface. Switches can also provide broadcast domain segmentation using virtual LANs (VLANs). A VLAN is a group of switch ports, within a single or multiple switches, that is defined by the switch hardware and/or software as a single broadcast domain. A VLAN's goal is to group devices connected to a switch into logical broadcast domains to control the effect that broadcasts have on other connected devices. A VLAN can be characterized as a logical network.
Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 8

QUESTION 15
In the three-layer hierarchical network design model; what's associated with the access layer? (Select two)

A. optimized transport structure
B. high port density
C. boundary definition
D. data encryption
E. local VLANs
F. route summaries

Answer: B, E

Explanation:
The access layer is the outermost layer, and it is composed of the least sophisticated network equipment. The most important function of the access layer is high port density, since these devices connect the individual end users. The access layers are also where VLANs are implemented, since VLANs are assigned on a per-port basis.

QUESTION 16
In the three-layer hierarchical network design model, what's associated with the core layer? (Select two)

A. Access control list

B. Data encryption
C. Optimized transport
D. Address aggregation
E. Packet switching
F. Boundary definition

Answer: C, E

Explanation:
A hierarchical network design includes the following three layers:
-The backbone (core) layer that provides optimal transport between sites
-The distribution layer that provides policy-based connectivity
-The local-access layer that provides workgroup/user access to the network
The distribution layer of the network is the demarcation point between the access and core layers and helps to define and differentiate the core. The purpose of this layer is to provide boundary definition and is the place at which packet manipulation can take place.
In the campus environment, the distribution layer can include several functions, such as the following:
Address or area aggregation
1. Departmental or workgroup access
2. Broadcast/multicast domain definition
3. Virtual LAN (VLAN) routing
4. Any media transitions that need to occur
5. Security
The distribution layer can be summarized as the layer that provides policy-based connectivity
Reference: www.alteridem.net/networking/idg4/idgbasic.htm

QUESTION 17
Two Certkiller switches are connected as shown below:

Configuration exhibit

```
Hostname CertKillerA
!
Interface fastethernet 0/10
spanningtree vlan 1-5 port priority 10
Switchport mode trunk
!
Interface fastethernet 0/12
spanningtree vlan 6-10 port priority 10
switchport mode trunk
```

Please refer to the exhibit above. Given the partial configuration of the two Cisco Certkiller switches, which two statements are true about VLAN traffic? (Select two)

A. VLANs 1-5 will be blocked if fa0/10 goes down.
B. VLANs 6-10 will use fa0/10 as a backup only.

C. VLANs 1-5 will use fa0/10 as a backup only.
D. VLANs 6-10 have a port priority of 128 on fa0/10.
E. VLANs 1-10 are configured to load share between fa0/10 and fa0/12.

Answer: B, E

Explanation:
Spanning-Tree Protocol (STP) is a Layer 2 protocol that utilizes a special-purpose algorithm to discover physical loops in a network and effect a logical loop-free topology. STP creates a loop-free tree structure consisting of leaves and branches that span the entire Layer 2 network. The actual mechanics of how bridges communicate and how the STP algorithm works will be discussed at length in the following topics. Note that the terms bridge and switch are used interchangeably when discussing STP. In addition, unless otherwise indicated, connections between switches are assumed to be trunks.
Load sharing can be accomplished using a couple of methods. The most common method of load sharing is through root bridge placement on a per-VLAN basis. This will distribute traffic for separate VLANs across separate paths to different root bridges. A separate method divides the bandwidth supplied by parallel trunks connecting switches. To avoid loops, STP normally blocks all but one parallel link between switches. Using load sharing, traffic can be divided between the links according to which VLAN the traffic belongs.
Load sharing can be configured on trunk ports by using STP port priorities or STP path costs. For load sharing using STP port priorities, both load-sharing links must be connected to the same switch. For load sharing using STP path costs, each load-sharing link can be connected to the same switch or to two different switches.
Load Sharing Using STP Port Priorities
When two ports on the same switch form a loop, the STP port priority setting determines which port is enabled and which port is in a blocking state. The priorities on a parallel trunk port can be set so that the port carries all the traffic for a given VLAN. The trunk port with the higher priority (lower values) for a VLAN is forwarding traffic for that VLAN. The trunk port with the lower priority (higher values) for the same VLAN remains in a Blocking state for that VLAN. One trunk port sends or receives all traffic for the VLAN.

QUESTION 18
SIMULATION
Certkiller topology exhibit:

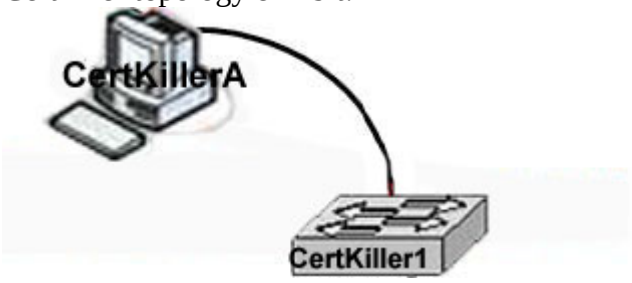

You work as a network engineer at Certkiller .com. Certkiller .com is a large international company with offices on all continents. You work at the Chicago

Office. Your boss at Certkiller .com, Miss Certkiller, has asked you to install a Catalyst 3500 to provide for another 24 users. Your instructions are as follows:
* install it in an IDF
* extremely import to have the proper configuration of the Catalyst before it is placed into the product information
* it should not participate in VTP
* it should forward VTP advertisements that are received on trunk ports
* all nontrunking interfaces, Fa0/1 to Fa0/24) should immediately to the forwarding state of the Spanning tree.
* configure all FastEthernet ports (the user ports) such as they are nontrunking.
* the fastEthernet interfaces 0/12 through 0/24 should be placed in VLAN 20

Answer:

Explanation:
Certkiller 1#conf t
Certkiller 1(config)#vtp mode transparent
Certkiller 1(config)#interface range fa0/1 - 24
Certkiller 1(config-if-range)#switchport mode access
Certkiller 1(config-if-range)#spanning-tree portfast
Certkiller 1(config)#interface range fa0/12 - 24
Certkiller 1(config-if-range)#switchport access vlan 20
Certkiller 1(config-if-range)#end
Certkiller 1# copy run start
Spanning tree PortFast is a Catalyst feature that causes a switch or trunk port to enter the spanning tree Forwarding state immediately, bypassing the Listening and Learning states. IOS-based switches only use PortFast on access ports connected to end stations.
When a device is connected to a port, the port normally enters the spanning tree Listening state. When the Forward Delay timer expires, the port enters the Learning state. When the Forward Delay timer expires a second time, the port is transitioned to the Forwarding or Blocking state. When PortFast is enabled on a switch or trunk port, the port is immediately transitioned to the Forwarding state. As soon as the switch detects the link, the port is transitioned to the Forwarding state (less than 2 seconds after the cable is plugged in).
Certkiller 1(Config-if-range)#switchport mode access : Brings the interfaces into access mode
Certkiller 1(Config-if-range)#spanning-tree portfast : Enables the PortFast on interface.
Certkiller 1(Config-if-range)#switchport access vlan 20 : Makes the members of vlan 20

QUESTION 19
The following output was seen on a Certkiller switch:

```
CertKiller1# show interfaces fastethernet 0/1 switchport
Name: Fa0/1
Switchport: Enabled
Administrative Mode: static access
Operational Mode: down
Administrative Trunking Encapsulation: negotiate
Negotiation of Trunking: Off
Access Mode VLAN: 10 (DATA)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk Native VLAN tagging: enabled
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk private VLANs: none
Switchport private-vlan: none
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL

Protected: false
Unknown unicast blocked: disabled
Unknown multicast blocked: disabled
Appliance trust: none
CertKiller1#
```

Study the exhibit carefully. The Certkiller user who is connected to interface FastEthernet 0/1 is on VLAN 10 and cannot access network resources. On the basis of the information in the exhibit, which command sequence would correct the problem?

A. Certkiller 1(config)# vlan 10
Certkiller 1(config-vlan)# no shut
B. Certkiller 1(config)# interface fastethernet 0/1
Certkiller 1(config-if)# switchport mode access
Certkiller 1(config-if)# switchport access vlan 10
C. Certkiller 1(config)# interface fastethernet 0/1
Certkiller 1(config-if)# switchport mode access
D. Certkiller 1(config)# interface fastethernet 0/1
Certkiller 1(config-if)# no shut
E. Certkiller 1(config)# vlan 10
Certkiller 1(config-vlan)# state active

Answer: D

Explanation:
In Exhibit Operation Mode is down, it means interface is in down state. Just bring into up state using no shutdown command

---

QUESTION 20
You need make configuration changes to an existing layer 3 switch in the Certkiller network. On a multilayer Catalyst switch, which interface command is used to convert a Layer 3 interface to a Layer 2 interface?

A. swithport access vlan vlan-id
B. switchport

C. switchport mode access
D. no switchport
E. None of the above

Answer: B

Explanation:
The switchport command puts the port in Layer 2 mode. Then, you can use other switchport command keywords to configure trunking, access VLANs, and so on.

QUESTION 21
What is a characteristic of assigning a static VLAN membership?

A. VMPS server lookup is required
B. Easy to configure
C. Easy of adds, moves, and changes
D. Based on MAC address of the connected device

Answer: B

Explanation:
Static port VLAN membership on the switch is assigned manually by the administrator on a port-by-port basis.
Characteristics of static VLAN configurations include the following:
1. Secure
2. Easy to configure
3. Straight forward to monitor
4. Works well in networks where moves, adds, and changes are rare.
Incorrect Answers:
A: VMPS server lookups are a function of dynamic VLANs and are not used with statically assigned VLANs.
C: Moves, adds, and changes, would require a network administrator to change the configuration of the switch every time a change is required.
D: This would describe a function of dynamic VLAN configurations, where the MAC address of the end user determines the VLAN that it belongs in, instead of the physical port.

QUESTION 22
Static VLANs are being used on the Certkiller network. What is true about static VLAN's?

A. Devices use DHCP to request their VLAN.
B. Attached devices are unaware of any VLANs.
C. Devices are assigned to VLANs based on their MAC addresses.
D. Devices are in the same VLAN regardless of which port they attach to.

Answer: B

Explanation:
LAN port VLAN membership can be assigned manually on a port-by-port basis. When you assign LAN ports to VLANs using this method, it is known as port-based, or static, VLAN membership.
Attached devices will be unaware of any VLANs.
Incorrect Answers:
A: The DHCP service is not involved in VLAN assignment.
C: Devices are not assigned to VLAN based on their MAC addresses. This is a function of dynamic VLANs.
D: Static VLANs are configured on a port by port basis.
Reference: Configuring VLANs
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/121_8aex/swconfig/vlans.htm

---

QUESTION 23
You are the network administrator at Certkiller and have justapplied this VLAN access map on one of your switches:
Router(config)#vlan access-map thor 10
Router(config-access-map)#match ip address net_10
Router(config-access-map)#action forward
Router(config-access-map)#exit
Router(config)#vlan filter thor vlan-list 12-15
What will this configuration result in?

A. All VLAN 12 through 15 IP traffic matching net_10 is forwarded and all other IP packets are dropped.
B. IP traffic matching net_10 is dropped an all other IP packets are forwarded to VLANs 12 through 15.
C. IP traffic matching vlan-list 12-15 is forwarded on all other IP packets are dropped.
D. All VLAN 12 through 15 IP traffic is forwarded, other VLAN IP traffic matching net_10 is dropped.

Answer: A

Explanation:
* vlan access-map thor 10 Defines the VLAN access map. Optionally, you can specify the VLAN access map sequence number.
* match ip address net_10 Configures the match clause in a VLAN access map sequence.
* action forward Configures the action clause in a VLAN access map sequence.
* vlan filter thor vlan-list 12-15 Applies the VLAN access map to the specified VLANs.
VLAN access maps can be applied to VLANs.
Each VLAN access map can consist of one or more map sequences, each sequence with a match clause and an action clause. The match clause specifies IP, IPX, or MAC ACLs for traffic filtering and the action clause specifies the action to be taken when a match

occurs. When a flow matches a permit ACL entry, the associated action is taken and the flow is not checked against the remaining sequences. When a flow matches a deny ACL entry, it will be checked against the next ACL in the same sequence or the next sequence. If a flow does not match any ACL entry and at least one ACL is configured for that packet type, the packet is denied.
To use access-control for both bridged and routed traffic, you can use VACLs alone or a combination of VACLs and ACLs. You can define ACLs on the VLAN interfaces to use access-control for both the input and output routed traffic. You can define a VACL to use access-control for the bridged traffic.
Reference:
http://www.cisco.com/en/US/products/hw/routers/ps368/products_configuration_guide_chapter09186a00801611

QUESTION 24
When configuring a VLAN for dynamic membership; which of the following guidelines NOT required?

A. Configure a VMPS server
B. Turn off trunking on the port
C. Turn off 802.1x port security
D. Configure the spanning-tree PortFast feature
E. All of the above are required for dynamic VLAN configuration

Answer: C
Reference:

Explanation:
Turning port security on or off is not necessary for enabling dynamic VLANs.
These guidelines and restrictions apply to dynamic port VLAN membership:
1. You must configure VMPS before you configure ports as dynamic.
2. When you configure a port as dynamic, spanning-tree PortFast is enabled automatically for that port. Automatic enabling of spanning-tree PortFast prevents applications on the host from timing out and entering loops caused by incorrect configurations. You can disable spanning-tree PortFast mode on a dynamic port.
3. If you reconfigure a port from a static port to a dynamic port on the same VLAN, the port connects immediately to that VLAN. However, VMPS checks the legality of the specific host on the dynamic port after a certain period.
4. Static secure ports cannot become dynamic ports. You must turn off security on the static secure port before it can become dynamic.
5. Static ports that are trunking cannot become dynamic ports. You must turn off trunking on the trunk port before changing it from static to dynamic.

QUESTION 25
Switch Certkiller 1 needs to have a port assigned to an existing VLAN. Which IOS command could you use to assign a switch port to a VLAN?

A. switchport mode access
B. switchport trunk access
C. switchport access vlan
D. switchport vlan

Answer: C

Explanation:
To assign a switchport to the VLAN, you would use the switchport access vlan interface configuration command.
Reference: CCNP Switching Exam Certification Guide: page 104, David Hucaby & Tim Boyles, Cisco Press 2001, ISBN 1-58720 000-7

QUESTION 26
What's true with VLAN port associations? (Select all that apply)

A. ASIC enhances the performance of the association
B. VLAN membership is based on Port through port-to-VLAN association.
C. Routing table enhances the performance of the association
D. VLAN membership is based on Port through port-to-WAN ID association.

Answer: A, B

Explanation:
ASIC (Application Specific Integrated Circuits), layer 2 switches have ASIC chips to help them with wire speed hardware switching. With ASIC, the performance of this association is very high, and is more desirable than the complex routing table lookup type of operation.
VLAN membership is based on Port through port-to-VLAN association.
Incorrect Answers:
C: Routing tables are not consulted when transferring VLAN traffic since VLANs are handled at layer 2 and routing occurs at layer 3.
D: VLAN associations deal with layer two VLANs, not layer 3 WAN IDs.

QUESTION 27
Which Cisco switch command would you use to map VLANs 10 to 20 to MST instance 1?

A. Switch(config)#vlan 10-20 instance 1
B. Switch(config)#instance 1 vlan 10-20
C. Switch(config-mst)#vlan 10-20 instance 1
D. Switch(config-mst)#instance 1 vlan 10-20
E. None of the above

Answer: D

Explanation:
Beginning in privileged EXEC mode, follow these steps to specify the MST region configuration and enable MSTP. This procedure is required.
Command Purpose
Step1 configure terminal Enter global configuration mode.
Step2 spanning-tree mst configuration Enter MST configuration mode.
Step3 instance instance-id vlan vlan-range Map VLANs to an MST instance.
For instance-id, the range is 1 to 15.
For vlan vlan-range, the range is 1 to 4094.
When you map VLANs to an MST instance, the mapping is incremental, and the VLANs specified in the command are added to or removed from the VLANs that were previously mapped.
To specify a VLAN range, use a hyphen; for example, instance 1 vlan 1-63maps VLANs 1 through 63 to MST instance 1.
To specify a VLAN series, use a comma; for example, instance 1 vlan 10, 20, 30maps VLANs 10, 20, and 30 to MST instance 1.
Step4 name name Specify the configuration name. The name string has a maximum length of 32 characters and is case sensitive.
Step5 revision version Specify the configuration revision number. The range is 0 to 65535.
Step6 show pending Verify your configuration by displaying the pending configuration.
Step7 exit Apply all changes, and return to global configuration mode.
Step8 spanning-tree mode mst Enable MSTP. RSTP is also enabled.
Caution
Changing spanning-tree modes can disrupt traffic because all spanning-tree instances are stopped for the previous mode and restarted in the new mode.
You cannot run both MSTP and PVST+ or both MSTP and rapid PVST+ at the same time.
Step9 end Return to privileged EXEC mode.
Step10 show running-config Verify your entries.
Step11 copy running-config startup-config (Optional) Save your entries in the configuration file.
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps5528/products_configuration_guide_chapter09186a00802

QUESTION 28
Which Cisco IOS command assigns a Catalyst switch port to VLAN 10?

A. switchport mode vlan 10
B. switchport trunk vlan 10
C. switchport access vlan 10
D. switchport mode access vlan 10

Answer: C

Explanation:

Switchport access:
Use the switchport access interface configuration command to configure a port as a static-access port. The port operates as a member of the configured VLAN.
Use the no form of this command to reset the access mode to the default VLAN for the switch.
Syntax
switchport access vlan vlan-id
no switchport access vlan vlan-id
Syntax Description
vlan vlan-id
ID of the VLAN. Valid IDs are from 1 to 1005. Do not enter leading zeroes.
Defaults
All ports are in static-access mode in VLAN 1.
Command Modes
Interface configuration.
Usage Guidelines
An access port can be assigned to only one VLAN.
When the no switchport access vlan form is used, the access mode is reset to static access on VLAN 1.
Example
The following example shows how to assign a port to VLAN 2 (instead of the default VLAN 1):
Switch(config-if)# switchport access vlan 2
You can verify the previous command by entering the show interface interface-id switchport command in privileged EXEC mode and examining information in the Administrative Mode and Operational Mode rows.

QUESTION 29
SIMULATION
You are connected to a Catalyst switch via a console cable as shown below:

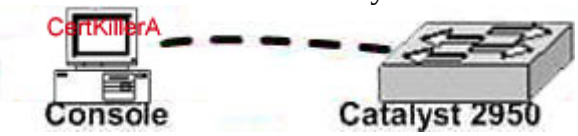

You work as a systems administrator at the Certkiller .com main office in the greater Toronto area. The number of employees on your floor has exceeded the infrastructure of your current network equipment. Your CTO has ordered a new switch chassis, but it's going to be another 6-8 working days until it arrives. In the meantime you can to connect 24 new workstations to an old Cisco Catalyst 2950, which your junior administrator has just finished erasing, and rebooting (to purge old VLAN information).
Your tasks are to:
* disable VTP
* Ensure that all non-trunking interfaces do not participate in Spanning Tree by default by globally configuring PortFast.
For the following two tasks, you are required to use global commands to configure the ports:

1. Ensure all FastEthernet interfaces are in permanent non-trunking mode.
2. Place FastEthernet interfaces 0/12 through 0/24 in VLAN 20.
Start by clicking on host CertK iA.

Answer:

Explanation:
enable
configure terminal
Switch(config)#vtp mode transparent (disable vtp)
Switch(config)#spanning-tree portfast default (Globally, enable portfast on all ports)
Switch(config)#interface range fa0/1 - 24 (select interfaces)
Switch(config-if)#switchport mode access (set ports for access mode, NOT Trunking)
switch(config)#interface range fa0/12 - 24 (The 4th task is to "Place FastEthernet interfaces 0/12 through 0/24 in VLAN20")
switch(config-if-range)#switchport access vlan 20
switch(config-if-range)#end
exit
Switch(config-if)#interface range fa0/12 - 24 (select interfaces)
Switch(config-if)#switchport access vlan 20 (assign ports to vlan 20)
end
copy running-config startup-config (save configuration)
The role of the VLAN Trunking Protocol (VTP) is to maintain VLAN configuration consistency across the entire network. VTP is a messaging protocol that uses Layer 2 trunk frames to manage the addition, deletion, and renaming of VLANs on a network-wide basis from a centralized switch that is in the VTP server mode. VTP is responsible for synchronizing VLAN information within a VTP domain. This reduces the need to configure the same VLAN information on each switch.
Transparent
VTP transparent switches do not participate in VTP. A VTP transparent switch does not advertise its VLAN configuration, and does not synchronize its VLAN configuration based on received advertisements. However, in VTP Version 2, transparent switches do forward VTP advertisements that the switches receive out their trunk ports. VLANs can be configured on a switch in the VTP transparent mode, but the information is local to the switch (VLAN information is not propagated to other switches) and is stored in NVRAM.
Spanning tree PortFast is a Catalyst feature that causes a switch or trunk port to enter the spanning tree Forwarding state immediately, bypassing the Listening and Learning states. IOS-based switches only use PortFast on access ports connected to end stations.
When a device is connected to a port, the port normally enters the spanning tree Listening state. When the Forward Delay timer expires, the port enters the Learning state. When the Forward Delay timer expires a second time, the port is transitioned to the Forwarding or Blocking state. When PortFast is enabled on a switch or trunk port, the port is immediately transitioned to the Forwarding state. As soon as the switch detects the link, the port is transitioned to the Forwarding state (less than 2 seconds after the cable is plugged in).

Certkiller 1(Config-if-range)#switchport mode access : Brings the interfaces into access mode
Certkiller 1(Config-if-range)#spanning-tree portfast : Enables the PortFast on interface.
Certkiller 1(Config-if-range)#switchport access vlan 20 : Makes the members of vlan 20

QUESTION 30
You are the network administrator of a network with the active VLANs: 1, 2, 3, 4, 10, 20, and 50. However, you only need to carry VLANs 1,2,10 and 20 on a trunk. Which of the following commands should you use to fulfil this requirement? (Select all that apply.)

A. switchport trunk allowed vlan remove 3,4,50
B. switchport trunk allowed vlan except 3,4,50
C. switchport trunk allowed vlan except 1,2,10,20
D. switchport trunk allowed vlan add 1,2,10,20
E. switchport trunk disallowed vlan remove 3,4,50
F. switchport trunk disallowed vlan add 3,4,50

Answer: A, B, D

Explanation:
switchporttrunk allowed vlan vlan_list
The vlan_list format is all | none | [add | remove | except] vlan_atom[,vlan_atom...], where:
* all specifies all VLANs from 1 to 4094. This keyword is not supported on commands that do not permit all VLANs in the list to be set at the same time.
* none indicates an empty list. This keyword is not supported on commands that require certain VLANs to be set or at least one VLAN to be set.
* add adds the defined list of VLANs to those currently set, instead of replacing the list.
* remove removes the defined list of VLANs from those currently set, instead of replacing the list.
* except lists the VLANs that should be calculated by inverting the defined list of VLANs.
* vlan_atom is either a single VLAN number from 1 to 4094 or a continuous range of VLANs described by two VLAN numbers, the lesser one first, separated by a hyphen.
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps663/products_command_reference_chapter09186a008014

QUESTION 31
When you're building up a VLAN, which of the following are part of the sequence? (Select all that apply)

A. Assign ports.
B. Create VLAN.
C. Create VLAN naming scheme
D. Configure ports for trunking.

E. Remove the trunk when the trunk is no longer needed.
F. Set the baud rate for ports
G. Verify configuration.

Answer: A, B, D, E, G

Explanation:
D, E: To create VLANs on a Catalyst switch, you must first enable the VLAN Trunking Protocol (VTP). The switch must be in VTP server or transparent mode to do this. VTP clients can not create VLANs.
B: The next step is to create the VLAN.
A: Once the VLAN is created, the final step is to assign individual ports to the VLAN.
G: After everything is configured, it should be verified.
Incorrect Answers:
C: A VLAN naming scheme isn't necessary because VLANs are numbered by default when they're created. Note that A VTP domain name must be created, but the VLANs themselves are not required to be named.
F: The baud rate doesn't have to be set for ports. Setting baud rate is for out-of-band console connections.

QUESTION 32
Two Certkiller switches are connected together as shown below:

Certkiller 1 configuration exhibit:

```
hostname CertKiller1

Interface fastethernet 0/12
switchport mode dynamic auto
switchport trunk encapsulation dot1q
switch out trunk native clan 5
!
!
```

Certkiller 2 configuration exhibit:

```
hostname CertKiller2

Interface fastethernet 0/21
switchport mode dynamic desirable
switchport trunk encapsulation dot1q
!
!
!
```

Based on the information shown above, which three statements are true? (Select three)

A. DTP packets are sent from Switch Certkiller 2.
B. DTP is not running on Switch Certkiller 1.
C. Only VLANs 1-1001 will travel across the trunk link.
D. A trunk link will be formed.
E. The native VLAN for Switch Certkiller 2 is vlan 1.

Answer: A, D, E

Explanation:
You can manually configure trunk links on Catalyst switches for either ISL or 802.1Q mode. In addition, Cisco has implemented a proprietary, point-to-point protocol called Dynamic Trunking Protocol (DTP) that negotiates a common trunking mode between two switches. The negotiation covers the encapsulation (ISL or 802.1Q) as well as whether the link becomes a trunk at all.
You can configure the trunk encapsulation with the switchport trunk encapsulation command, as one of the following:
1. isl-VLANs are tagged by encapsulating each frame using the Cisco ISL protocol.
2. dot1q-VLANs are tagged in each frame using the IEEE 802.1Q standard protocol. The only exception is the native VLAN, which is sent normally and not tagged at all.
1. negotiate(the default)-The encapsulation is negotiated to select either ISL or IEEE 802.1Q, whichever is supported by both ends of the trunk. If both ends support both types, ISL is favored. (The Catalyst 2950 switch does not support ISL encapsulation.)
In the switchport mode command, you can set the trunking mode to any of the following:
1. trunk-This setting places the port in permanent trunking mode. The corresponding switch port at the other end of the trunk should be similarly configured because negotiation is not allowed. You should also manually configure the encapsulation mode.
2. dynamicdesirable (the default)-The port actively attempts to convert the link into trunking mode. If the far-end switch port is configured to trunk, dynamic desirable, or dynamic auto mode, trunking is successfully negotiated.
3. dynamicauto-The port converts the link into trunking mode. If the far-end switch port is configured to trunk or dynamic desirable, trunking is negotiated.
Because of the passive negotiation behavior, the link never becomes a trunk if both ends of the link are left to the dynamic auto default.

QUESTION 33
Two Certkiller switches are connected via a trunk link. In this network, the original frame is encapsulated and an additional header is added before the frame is carried over a trunk link. At the receiving end, the header is removed and the frame is forwarded to the assigned VLAN. This describes which technology?

A. DISL
B. ISL
C. DTP
D. IEEE 802.1Q
E. MPLS
F. None of the above

Answer: B

Explanation:
Inter-Switch Link Protocol

The Inter-Switch Link (ISL) protocol is a Cisco proprietary method for preserving the source VLAN identification of frames passing over a trunk link. ISL performs frame identification in Layer 2 by encapsulating each frame between a header and trailer. Any Cisco switch or router device configured for ISL can process and understand the ISL VLAN information. ISL is primarily used for Ethernet media, although Cisco has included provisions to carry Token Ring, FDDI, and ATM frames over Ethernet ISL. (A Frame-Type field in the ISL header indicates the source frame type.)
When a frame is destined out a trunk link to another switch or router, ISL adds a 26-byte header and a 4-byte trailer to the frame. The source VLAN is identified with a 10-bit VLAN ID field in the header. The trailer contains a cyclic redundancy check (CRC) value to ensure the data integrity of the new encapsulated frame. Figure 6-3 shows how Ethernet frames are encapsulated and forwarded out a trunk link. Because tagging information is added at the beginning and end of each frame, ISL is sometimes referred to as double tagging.

QUESTION 34
The Certkiller core switches use 802.1Q trunks to connect to each other. How does 802.1Q trunking keep track of multiple VLAN's?

A. It tags the data frame with VLAN information and recalculates the CRC value
B. It encapsulates the data frame with a new header and frame check sequence
C. It modifies the port index of a data frame to indicate the VLAN
D. It adds a new header containing the VLAN ID to the data frame
E. None of the above

Answer: A

Explanation:
The IEEE 802.1Q protocol can also carry VLAN associations over trunk links. However, this frame identification method is standardized, allowing VLAN trunks to exist and operate between equipment from multiple vendors.
In particular, the IEEE 802.1Q standard defines an architecture for VLAN use, services provided with VLANs, and protocols and algorithms used to provide VLAN services.
Like Cisco ISL, IEEE 802.1Q can be used for VLAN identification with Ethernet trunks. Instead of encapsulating each frame with a VLAN ID header and trailer, 802.1Q embeds its tagging information within the Layer 2 frame. This method is referred to as single-tagging or internal tagging.
802.1Q also introduces the concept of a native VLAN on a trunk. Frames belonging to this VLAN are not encapsulated with any tagging information. In the event that an end station is connected to an 802.1Q trunk link, the end station can receive and understand only the native VLAN frames. This provides a simple way to offer full trunk encapsulation to the devices that can understand it, while giving normal access stations some inherent connectivity over the trunk.

QUESTION 35
Refer to the following exhibits:

Exhibit #1:

```
CertKiller1# show interface fa0/13 switchport
Name: Fa0/13
Switchport: Enabled
Administrative Mode: dynamic auto
Operational Mode: down
Administrative Trunking Encapsulation: negotiate
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk Native VLAN tagging: enabled
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk private VLANs: none
Operational private-vlan: none
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL

<omitted output>
```

Exhibit #2:

```
CertKiller1# show interface fa0/13 switchport
Name: Fa0/13
Switchport: Enabled
Administrative Mode: dynamic desirable
Operational Mode: down
Administrative Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk Native VLAN tagging: enabled
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk private VLANs: none
Operational private-vlan: none
Trunking VLANs Enabled: 1,10,20
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL

<omitted output>
```

Study the exhibits carefully. The switchport output in Exhibit #1 displays the default settings of interface FastEthernet 0/13 on switch Certkiller 1. Figure 2 displays the desired interface settings. Which command sequence would configure interface FastEthernet 0/13 as displayed in Exhibit #2?

A. Certkiller 1(config-if)# switchport trunk encapsulation dot1q
Certkiller 1(config-if)# switchport mode dynamic auto
Certkiller 1(config-if)# switchport trunk native DATA
Certkiller 1(config-if)# witchport trunk allowed vlan add 1,10,20
B. Certkiller 1(config-if)# switchport trunk encapsulation dot1q
Certkiller 1(config-if)# switchport mode dynamic desirable

Certkiller 1(config-if)# switchport trunk native vlan DATA
Certkiller 1(config-if)# switchport trunk allowed vlan 1,10,20
C. Certkiller 1(config-if)# switchport trunk encapsulation dot1q
Certkiller 1(config-if)# switchport mode trunk
Certkiller 1(config-if)# switchport trunk native DATA
Certkiller 1(config-if)# switchport trunk allowed vlan 1,10,20
D. Certkiller 1(config-if)# switchport trunk encapsulation dot1q
Certkiller 1(config-if)# switchport mode dynamic desirable
Certkiller 1(config-if)# switchport trunk native vlan 10
E. Certkiller 1(config-if)# switchport trunk encapsulation dot1q
Certkiller 1(config-if)# switchport mode dynamic desirable
Certkiller 1(config-if)# switchport trunk native vlan 10
Certkiller 1(config-if)# switchport trunk allowed vlan 1,10,20

Answer: E

Explanation:
The IEEE 802.1Q protocol can also carry VLAN associations over trunk links. However, this frame identification method is standardized, allowing VLAN trunks to exist and operate between equipment from multiple vendors.
In particular, the IEEE 802.1Q standard defines an architecture for VLAN use, services provided with VLANs, and protocols and algorithms used to provide VLAN services.
Like Cisco ISL, IEEE 802.1Q can be used for VLAN identification with Ethernet trunks. Instead of encapsulating each frame with a VLAN ID header and trailer, 802.1Q embeds its tagging information within the Layer 2 frame. This method is referred to as single-tagging or internal tagging.
802.1Q also introduces the concept of a native VLAN on a trunk. Frames belonging to this VLAN are not encapsulated with any tagging information. In the event that an end station is connected to an 802.1Q trunk link, the end station can receive and understand only the native VLAN frames. This provides a simple way to offer full trunk encapsulation to the devices that can understand it, while giving normal access stations some inherent connectivity over the trunk.
You can manually configure trunk links on Catalyst switches for either ISL or 802.1Q mode. In addition, Cisco has implemented a proprietary, point-to-point protocol called Dynamic Trunking Protocol (DTP) that negotiates a common trunking mode between two switches. The negotiation covers the encapsulation (ISL or 802.1Q) as well as whether the link becomes a trunk at all.
In the switchport mode command, you can set the trunking mode to any of the following:
1. trunk-This setting places the port in permanent trunking mode. The corresponding switch port at the other end of the trunk should be similarly configured because negotiation is not allowed. You should also manually configure the encapsulation mode.
2. dynamicdesirable (the default)-The port actively attempts to convert the link into trunking mode. If the far-end switch port is configured to trunk, dynamic desirable, or dynamic auto mode, trunking is successfully negotiated.
3. dynamicauto-The port converts the link into trunking mode. If the far-end switch port is configured to trunk or dynamic desirable

, trunking is negotiated. Because of the passive negotiation behavior, the link never becomes a trunk if both ends of the link are left to the dynamic auto default.
802.1Q also introduces the concept of a native VLAN on a trunk. Frames belonging to this VLAN are not encapsulated with any tagging information. In the event that an end station is connected to an 802.1Q trunk link, the end station can receive and understand only the native VLAN frames. This provides a simple way to offer full trunk encapsulation to the devices that can understand it, while giving normal access stations some inherent connectivity over the trunk.
switchporttrunk allowed vlan, defines which VLANs can be trunked over the link. By default, a switch transports all active VLANs (1 to 4094) over a trunk link. There might be times when the trunk link should not carry all VLANs. For example, broadcasts are forwarded to every switch port on a VLAN-including the trunk link because it, too, is a member of the VLAN. If the VLAN does not extend past the far end of the trunk link, propagating broadcasts across the trunk makes no sense.

QUESTION 36
The core Certkiller switches are configured to use 802.1Q trunks. Which three statements are correct with regard to the IEEE 802.1Q standard? (Select three)

A. The IEEE 802.1Q frame format adds a 4 byte field to a Ethernet frame
B. The packet is encapsulated with a 26 byte header and a 4 byte FCS
C. The protocol uses point-to-multipoint connectivity
D. The protocol uses point-to-point connectivity
E. The IEEE 802.1Q frame uses multicast destination of 0x01-00-0c-00-00
F. The IEEE 802.1Q frame retains the original MAC destination address

Answer: A, D, F

Explanation:
The IEEE 802.1Q protocol can also carry VLAN associations over trunk links. However, this frame identification method is standardized, allowing VLAN trunks to exist and operate between equipment from multiple vendors.
In particular, the IEEE 802.1Q standard defines an architecture for VLAN use, services provided with VLANs, and protocols and algorithms used to provide VLAN services.
Like Cisco ISL, IEEE 802.1Q can be used for VLAN identification with Ethernet trunks. Instead of encapsulating each frame with a VLAN ID header and trailer, 802.1Q embeds its tagging information within the Layer 2 frame. This method is referred to as single-tagging or internal tagging.
802.1Q also introduces the concept of a native VLAN on a trunk. Frames belonging to this VLAN are not encapsulated with any tagging information. In the event that an end station is connected to an 802.1Q trunk link, the end station can receive and understand only the native VLAN frames. This provides a simple way to offer full trunk encapsulation to the devices that can understand it, while giving normal access stations some inherent connectivity over the trunk.

QUESTION 37
Switch CK1 has been configured with DTP using the desirable option. Which statement describes Dynamic Trunking Protocol (DTP) desirable mode?

A. The interface actively attempts to convert the link to a trunk link.
B. The interface is put into permanent trunking mode but prevented from generating DTP frames.
C. The interface is put into permanent trunking mode and negotiates to convert the link into a trunk link.
D. The interface is put into a passive mode, waiting to convert the link to a trunk link.
E. None of the above

Answer: A

Explanation:
In the switchport mode command, you can set the trunking mode to any of the following:
1. trunk-This setting places the port in permanent trunking mode. The corresponding switch port at the other end of the trunk should be similarly configured because negotiation is not allowed. You should also manually configure the encapsulation mode.
2. dynamicdesirable (the default)-The port actively attempts to convert the link into trunking mode. If the far-end switch port is configured to trunk, dynamic desirable, or dynamic auto mode, trunking is successfully negotiated.
3. dynamicauto-The port converts the link into trunking mode. If the far-end switch port is configured to trunk or dynamic desirable, trunking is negotiated.
Because of the passive negotiation behavior, the link never becomes a trunk if both ends of the link are left to the dynamic auto default.

QUESTION 38
DRAG DROP
You work as a network technician at Certkiller .com. You are required to put the DTP mode next to the correct description using the boxes shown below:

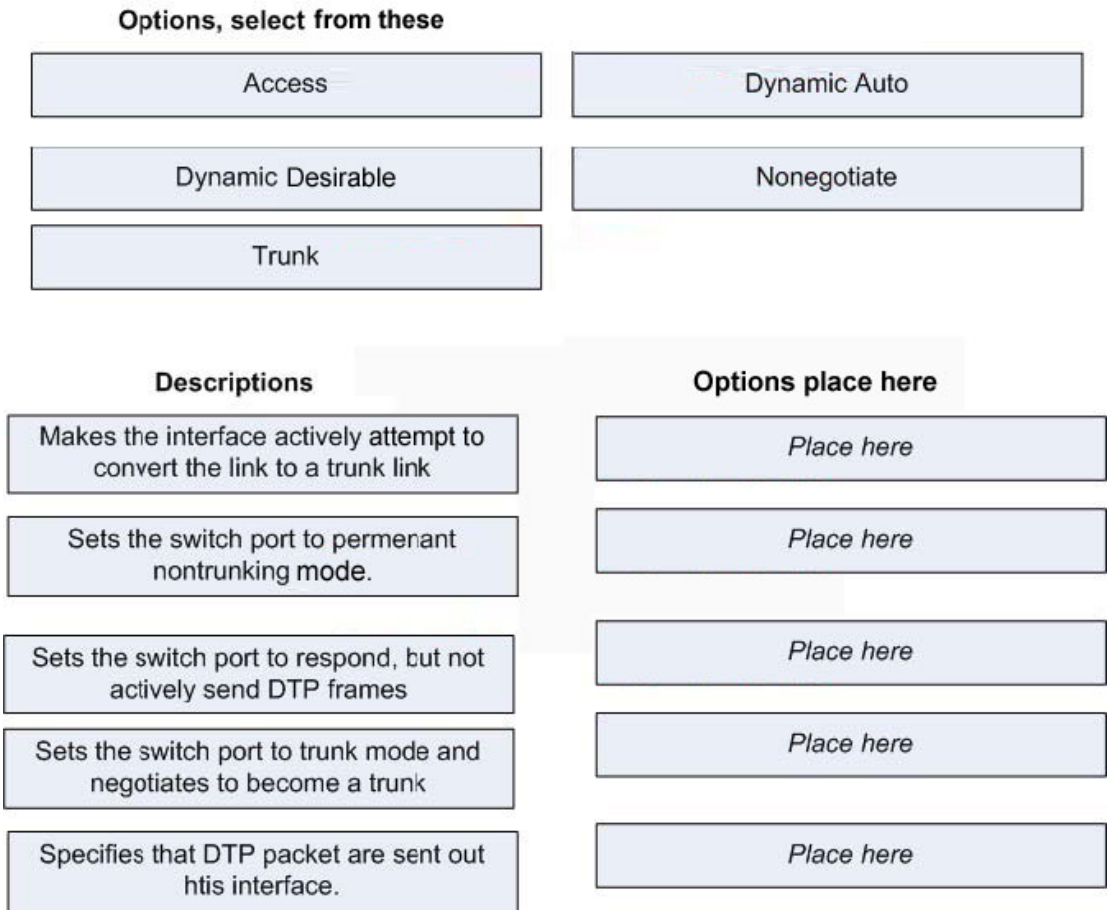

Answer:

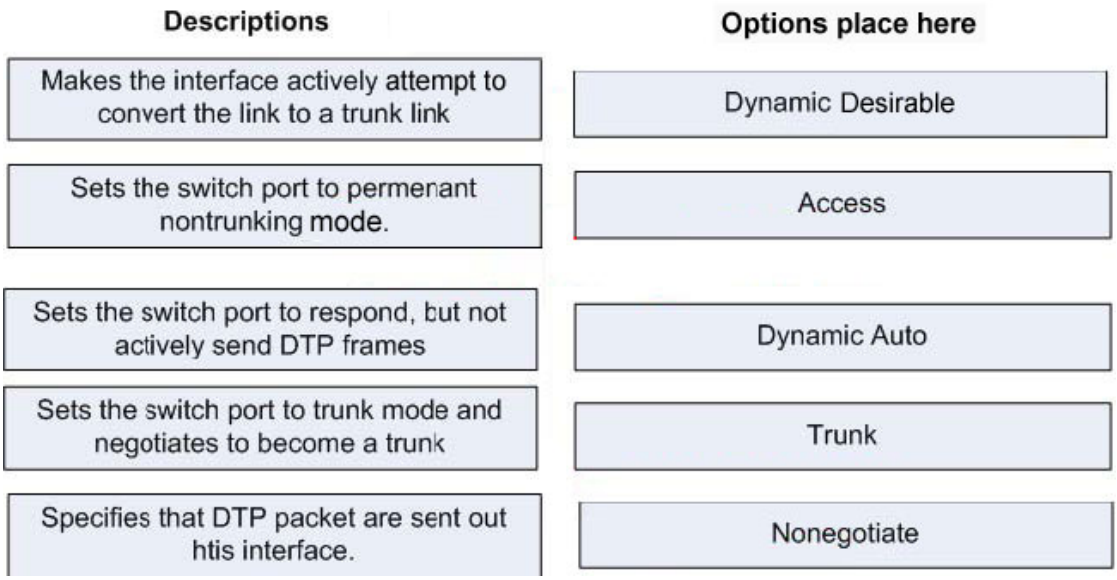

Explanation:

1. trunk-This setting places the port in permanent trunking mode. The corresponding switch port at the other end of the trunk should be similarly configured because negotiation is not allowed. You should also manually configure the encapsulation mode.

2. dynamicdesirable (the default)-The port actively attempts to convert the link into trunking mode. If the far-end switch port is configured to trunk, dynamic desirable, or dynamic auto mode, trunking is successfully negotiated.
3. dynamicauto-The port converts the link into trunking mode. If the far-end switch port is configured to trunk or dynamic desirable, trunking is negotiated.
Because of the passive negotiation behavior, the link never becomes a trunk if both ends of the link are left to the dynamic auto default.
4. negotiate(the default)-The encapsulation is negotiated to select either ISL or IEEE 802.1Q, whichever is supported by both ends of the trunk. If both ends support both types, ISL is favored.
5. Access : Puts the interface into access mode that mean interface is in non-trunking mode.

QUESTION 39
You need to configure a new Certkiller switch for trunking. Which switch command enables a trunking protocol that appends a four byte CRC to the packet?

A. Certkiller Switch(config-if)#switchport trunk encapsulation dot1q
B. Certkiller Switch(config-if)#switchport trunk encapsulation itef
C. Certkiller Switch(config-if)#switchport trunk encapsulation fddi
D. Certkiller Switch(config-if)#switchport trunk encapsulation isl
E. None of the above

Answer: D

Explanation:
The Inter-Switch Link (ISL) protocol is a Cisco proprietary method for preserving the source VLAN identification of frames passing over a trunk link. ISL performs frame identification in Layer 2 by encapsulating each frame between a header and trailer. Any Cisco switch or router device configured for ISL can process and understand the ISL VLAN information. ISL is primarily used for Ethernet media, although Cisco has included provisions to carry Token Ring, FDDI, and ATM frames over Ethernet ISL. (A Frame-Type field in the ISL header indicates the source frame type.)
When a frame is destined out a trunk link to another switch or router, ISL adds a 26-byte header and a 4-byte trailer to the frame. The source VLAN is identified with a 10-bit VLAN ID field in the header. The trailer contains a cyclic redundancy check (CRC) value to ensure the data integrity of the new encapsulated frame. Figure 6-3 shows how Ethernet frames are encapsulated and forwarded out a trunk link. Because tagging information is added at the beginning and end of each frame, ISL is sometimes referred to as double tagging.

QUESTION 40
While using a packet analyzer, you notice four additional bytes being added to the packets in the Certkiller network. Which protocol inserts a four byte tag into the Ethernet frame and recalculates CRC value?

A. DTP
B. VTP
C. 802.1Q
D. ISL
E. None of the above

Answer: C

Explanation:
The IEEE 802.1Q protocol can also carry VLAN associations over trunk links. However, this frame identification method is standardized, allowing VLAN trunks to exist and operate between equipment from multiple vendors.
In particular, the IEEE 802.1Q standard defines an architecture for VLAN use, services provided with VLANs, and protocols and algorithms used to provide VLAN services.
Like Cisco ISL, IEEE 802.1Q can be used for VLAN identification with Ethernet trunks. Instead of encapsulating each frame with a VLAN ID header and trailer, 802.1Q embeds its tagging information within the Layer 2 frame. This method is referred to as single-tagging or internal tagging.
802.1Q also introduces the concept of a native VLAN on a trunk. Frames belonging to this VLAN are not encapsulated with any tagging information. In the event that an end station is connected to an 802.1Q trunk link, the end station can receive and understand only the native VLAN frames. This provides a simple way to offer full trunk encapsulation to the devices that can understand it, while giving normal access stations some inherent connectivity over the trunk.

QUESTION 41
You need to configure a new Certkiller switch to support DTP. Which DTP switchport mode parameter sets the switch port to actively send and respond to DTP negotiation frames?

A. Access
B. No negotiate
C. Trunk
D. Dynamic desirable
E. Dynamic auto
F. None of the above

Answer: D

Explanation:
dynamicdesirable (the default)-The port actively attempts to convert the link into trunking mode. If the far-end switch port is configured to trunk, dynamic desirable, or dynamic auto mode, trunking is successfully negotiated.

QUESTION 42
A new Certkiller switch was just configured using the "switchport trunk native vlan

7" command. What does this interface command accomplish?

A. Causes the interface to apply ISL framing for traffic on VLAN 7
B. Configures the trunking interface to forward traffic from VLAN 7
C. Configures the interface to be a trunking port and causes traffic on VLAN 7 to be 802.1q tagged
D. Configures the trunking interface to send traffic from VLAN 7 untagged
E. None of the above

Answer: D

Explanation:
In 802.1Q trunking, all VLAN packets are tagged on the trunk link to indicate the VLAN to which they belong. Frames belonging to the Native VLAN are sent untagged on the trunk link. The Native VLAN contains ports not assigned to other VLANs that by default belong to VLAN 1. VLAN 1 is the Native VLAN by default, but VLANs other than VLAN 1 may be designated as the Native VLAN. However, the Native VLAN must be the same on trunked switches in 802.1Q trunking. If a VLAN other than VLAN 1 is to be the Native VLAN, it needs to be identified on the trunk ports. In the interface configuration mode of the trunk port(s), the IOS-based command to designate the Native VLAN is switchport trunk native.
Switch(config-if)#switchport trunk native vlan vlan-id

QUESTION 43
You need to connect two Certkiller core switches via an ISL trunk. Which statement is true regarding the configuration of ISL trunks?

A. A Catalyst switch cannot have ISL and IEEE 802.1q trunks enabled.
B. All Catalyst switches support ISL trunking.
C. A Catalyst switch will report giants if one side is configured for ISL while the other side is not.
D. ISL trunking requires that native VLANs match.
E. None of the above

Answer: C

Explanation:
The Inter-Switch Link (ISL) protocol is a Cisco proprietary method for preserving the source VLAN identification of frames passing over a trunk link. ISL performs frame identification in Layer 2 by encapsulating each frame between a header and trailer. Any Cisco switch or router device configured for ISL can process and understand the ISL VLAN information. ISL is primarily used for Ethernet media, although Cisco has included provisions to carry Token Ring, FDDI, and ATM frames over Ethernet ISL. (A Frame-Type field in the ISL header indicates the source frame type.)
When a frame is destined out a trunk link to another switch or router, ISL adds a 26-byte header and a 4-byte trailer to the frame. The source VLAN is identified with a 10-bit

VLAN ID field in the header. The trailer contains a cyclic redundancy check (CRC) value to ensure the data integrity of the new encapsulated frame. Figure 6-3 shows how Ethernet frames are encapsulated and forwarded out a trunk link. Because tagging information is added at the beginning and end of each frame, ISL is sometimes referred to as double tagging.

QUESTION 44
Two Certkiller switches are connected as shown in the diagram below:

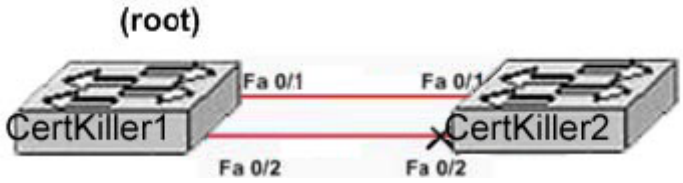

Study the exhibit above carefully. VLAN 1 and VLAN 2 are configured on the trunked links between Certkiller 1 and Certkiller 2. Port Fa 0/2 on Certkiller 2 is currently in a blocking state for both VLANs. What should be done to load balance VLAN traffic between Certkiller 1 and Certkiller 2?

A. Make the bridge ID of Certkiller 2 lower than the ID of Certkiller 1.
B. Lower the port priority for VLAN 1 on port 0/1 for Certkiller 1.
C. Enable HSRP on the access ports.
D. Lower the port priority for VLAN 1 on port 0/2 for Certkiller 1.
E. None of the above.

Answer: D

Explanation:
Load Sharing Using STP Port Priorities
When two ports on the same switch form a loop, the STP port priority setting determines which port is enabled and which port is in a blocking state. The priorities on a parallel trunk port can be set so that the port carries all the traffic for a given VLAN. The trunk port with the higher priority (lower values) for a VLAN is forwarding traffic for that VLAN. The trunk port with the lower priority (higher values) for the same VLAN remains in a Blocking state for that VLAN. One trunk port sends or receives all traffic for the VLAN.

QUESTION 45
Which of the following technologies would an Internet Service Provider use to support overlapping customer VLAN ID's over transparent LAN services?

A. 802.1q tunneling
B. ATM
C. SDH
D. IP Over Optical Networking
E. ISL

Answer: A

Explanation:
Understanding How 802.1Q Tunneling Works:
The 802.1Q tunnelling feature supports secure virtual private networks (VPNs). 802.1Q tunnelling enables service providers to keep traffic from different customers segregated in the service provider infrastructure while significantly reducing the number of VLANs required to support the VPNs. 802.1Q tunnelling allows multiple customer VLANs to be carried by a single VLAN on the Catalyst 6000 family switch without losing their unique VLAN IDs.
When you configure 802.1Q tunnelling on the Catalyst 6000 family switch, traffic to be tunnelled comes into the switch from an 802.1Q trunk port on a neighboring device and enters the switch through a port configured to support 802.1Q tunnelling (a tunnel port). When the tunnel port receives traffic from an 802.1Q trunk port, it does not strip the 802.1Q tags from the frame header but, instead, leaves the 802.1Q tags intact and puts all the received 802.1Q traffic into the VLAN assigned to the tunnel port. The VLAN assigned to the tunnel port then carries the tunnelled customer traffic to the other neighboring devices participating in the tunnel port VLAN. When the tunnelled traffic is received by an 802.1Q trunk port on a neighboring device, the 802.1Q tag is stripped and the traffic is removed from the tunnel.
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps700/products_configuration_guide_chapter09186a008007f

---

QUESTION 46
If you were to configure an ISL Ethernet trunk between two Cisco switches, named CK1 and CK2 , what would you have to include at the end of the link for the trunk to operate correctly? (Select two)

A. An identical VTP mode.
B. An identical speed/duplex.
C. An identical trunk negotiation parameter.
D. An identical trunk encapsulation parameter.

Answer: B, D

Explanation:
In order for a trunk to be operational, the speed and duplex settings must match at each end of the trunk, and both switches must use the same trunking encapsulation (802.1Q or ISL).
Incorrect Answers:
A: It is common for switches to have trunk links operating, while the VTP modes differ. For example, a switch configured with VTP mode server can have a trunk connected to a switch with VTP mode client.
C: This is incorrect, as there are a number of configurations that are supported where the trunk negotiation parameters differ between switches. For example, switch CK1 could

have the trunk configured for "on" while switch CK2 could have the switch trunk configured for "desirable" and the trunk would be operational.

QUESTION 47
DRAG DROP
Drag-and-drop the technology term on the left to the correct options column on the right (not all of the options will be used.)

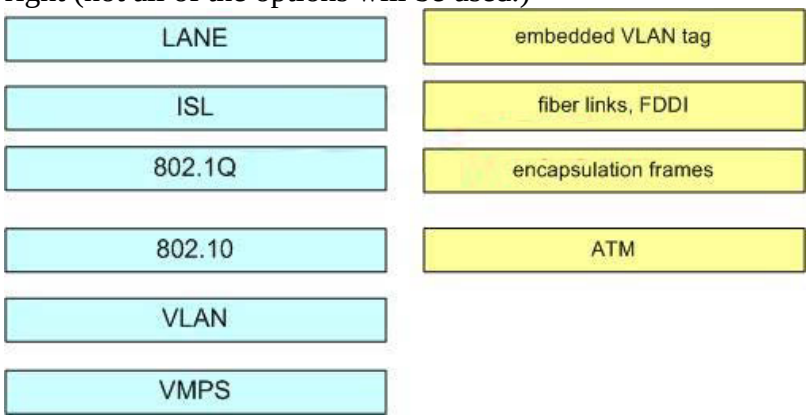

Answer:

Explanation:
LANE - ATM
ISL - Encapsulation frames
802.1Q - embedded VLAN tag
802.10 - Fiber links, FDDI
VLAN
VMPS
1. LANE - LAN Emulation - An IEEE standard method for transporting VLANs over Asynchronous Transfer Mode (ATM) networks.
2. ISL - A Cisco Proprietary encapsulation protocol for interconnection multiple switches.
3. 802.1Q - An IEEE standard method for identifying VLANs by inserting a VLAN identifier into the frame header. This process is called frame tagging.
4. 802.10 - A Cisco Proprietary method of transporting VLAN information inside the standard 802.10 frame (Fiber Distributed Data Interface [FDDI]).
Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 99

QUESTION 48
You are the network administrator at Certkiller and switch CK1 is configured as shown below:
Interface gigethernet 0/1
Switchport mode trunk
Switchport trunk encapsulation dot1q
Switchport trunk native vlan 5
If untagged frames are arriving on interface gigethernet 0/1 of CK1 , which of the following statement are correct?

A. Untagged frames are automatically assumed to be in VLAN 5.
B. Untagged frames are defaulted to VLAN 1 traffic.
C. Untagged frames are dropped because all packets are tagged when dot1q trunked.
D. Untagged frames are determined on the other switch
E. Untagged frames are not supported on 802.1Q trunks.

Answer: A

Explanation:
Each physical port has a parameter called PVID. Every 802.1Q port is assigned a PVID value that is of its native VLAN ID (default is VLAN 1). All untagged frames are assigned to the LAN specified in the PVID parameter. When a tagged frame is received by a port, the tag is respected. If the frame is untagged, the value contained in the PVID is considered as a tag. All untagged frames will be assigned to the native VLAN. The native VLAN is 1 by default, but in this case the native VLAN is configured as VLAN 5 so choice A is correct.

QUESTION 49
If you were to set up a VLAN trunk over a Fast Ethernet link on switch CK1 , which trunk mode would you set the local port to on CK1 if you wanted it to respond to requests from its link partner (CK2) and become a trunk?

A. Auto
B. Negotiate
C. Designate
D. Nonegotiate

Answer: A

Explanation:
Only ports in desirable and auto mode will negotiate a channel (either desirable-auto or desirable-desirable). Ports in on mode will only form a functional channel with other ports in on mode (they will not negotiate a channel with ports in desirable or auto mode).
Reference: Cisco, Troubleshooting Tips
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/trbl_ja.htm

---

QUESTION 50
Which of the following trunking modes are unable to request their ports to convert their links into trunk links? (Select all that apply)

A. Negotiate
B. Designate
C. Nonegotiate
D. Auto
E. Manual

F. Off

Answer: C, D

Explanation:
Auto is a trunking mode but does not actively negotiate a trunk. It requires opposite side to be trunk or desirable, and will only respond to requests from the other trunk link. No-negotiate will configure the link to be unable to dynamically become a trunk; since no requests will be sent it will not respond to requests from other trunk links from a different switch.
Incorrect Answers:
A, B, E, F: These choices are wrong because they are not valid trunking modes

---

QUESTION 51
ISL is being configured on a Certkiller switch. Which of the following choices are true regarding the ISL protocol? (Select two)

A. It can be used between Cisco and non-Cisco switch devices.
B. It calculates a new CRC field on top of the existing CRC field.
C. It adds 4 bytes of protocol-specific information to the original Ethernet frame.
D. It adds 30 bytes of protocol-specific information to the original Ethernet frame.

Answer: B, D

Explanation:
ISL adds a total of 30bytes to the Ethernet frame. A 26 byte header (10bytes identifies the VLAN ID) and a 4 byte trailer (containing a separate CRC).
Incorrect Answers:
A: This is incorrect because ISL is Cisco proprietary and can only be used on Cisco devices. For configuring a trunk to a non-Cisco switch, 802.1Q encapsulation should be used.
C: This is incorrect because it is contradictory to D. 30 byes are added with ISL, not 4 bytes. This choice describes what is used in 802.1Q frames, not ISL

QUESTION 52
You are the network administrator tasked with designing a switching solution for the Certkiller network. Which of the following statements describing trunk links are INCORRECT? (Select all that apply)

A. The trunk link belongs to a specific VLAN.
B. Multiple trunk links are used to connect multiple devices.
C. A trunk link only supports native VLAN.
D. Trunk links use 802.10 to identify a VLAN.
E. The native VLAN of the trunk link is the VLAN that the trunk uses if that link fails for any reason.

Answer: A, B, C, D

Explanation:
A trunk is a point-to-point link that transmits and receives traffic between switches or between switches and routers. Trunks carry the traffic of multiple VLANs and can extend VLANs across an entire network. 100BaseT and Gigabit Ethernet trunks use Cisco ISL (the default protocol) or industry-standard IEEE 802.1Q to carry traffic for multiple VLANs over a single link. Frames received from users in the administratively-defined VLANs are identified or tagged for transmission to other devices. Based on rules you define, a unique identifier (the tag) is inserted in each frame header before it is forwarded. The tag is examined and understood by each device before any broadcasts or transmission to other switches, routers, or end stations. When the frame reaches the last switch or router, the tag is removed before the frame is transmitted to the target end station.
Incorrect Answers:
E: This statement is true, as untagged frames are always used with the native VLAN. The native VLAN is VLAN 1 by default in Cisco switches.

QUESTION 53
A Certkiller switch port is configured as a trunk using 802.1Q encapsulation. Which three statements regarding the IEEE 802.1Q standard are true? (Select three)

A. The packet is encapsulated with a 26 byte header and a 4 byte FCS.
B. The IEEE 802.1Q frame format adds a 4 byte field to an Ethernet frame.
C. The IEEE 802.1Q frame retains the original MAC destination address.
D. The IEEE 802.1Q frame uses multicast destination of 0x01-00-0c-00-00
E. The 802.1Q protocol uses point-to-point connectivity.
F. The 802.1Q protocol uses point-to-multipoint connectivity.

Answer: B, C, E

Explanation:
802.1Q frames add 4 bytes to the Ethernet frame. The original MAC address is left unaltered so the destination MAC is not changed. Trunks are always defined in a point to point configuration, with two switch ports used as the endpoints.
Incorrect Answers:
A: This describes the frame that is added to an ISL encapsulated frame, not an 802.1Q frame.
D: The destination MAC address is not altered when trunks are configured.
F: All trunks are always configured in a point to point fashion, there is no method available to support point to multipoint trunk configurations.

QUESTION 54
Which DTP switchport mode parameter would you use to set a switch port to actively send and respond to DTP negotiation frames on switch CK1 ?

A. access
B. trunk
C. no negotiate
D. dynamic desirable
E. dynamic auto

Answer: D

Explanation:
There are five DTP switchport modes, and you should be familiar with all of them.
Access: This puts the interface (access port) into permanent nontrunking mode.
The interface will generate DTP frames, negotiating with the neighboring interface 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: The interface actively attempts 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. This is the default mode for all Ethernet interfaces. If the neighboring interface is set to the access or non-negotiate mode, the link will become a non-trunking link.
Dynamicauto: This command makes the interface willing to convert the link to a trunk link if the neighboring interface is set to trunk or desirable mode. Otherwise, the link will become a non-trunking link.
Switchport mode trunk: This command 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.
Switchport nonegotiate: Prevents the interface from generating DTP frames. You can use this command only when the interface switchport mode is access or trunk. You must manually configure the neighboring interface as a trunk interface to establish a trunk link; otherwise the link will be a non-trunking link.

QUESTION 55
A switch port on CK1 is being configured to support 802.1Q trunking. Which of the following are true about 802.1Q trunking? (Select one)

A. Both switches must be in the same VTP domain.
B. The encapsulation type of both ends of the trunk does not have to match.
C. The native VLAN on both ends of the trunk must be VLAN 1.
D. 802.1Q trunking can only be configured on a Layer 2 port.
E. In 802.1Q trunking, all VLAN packets are tagged on the trunk link, except the native VLAN.

Answer: E

Explanation:
E is correct because, "frames from the native VLAN of an 802.1Q trunk are not tagged with the VLAN number."

Incorrect Answers:
B: This is incorrect because the encapsulations types do have to match or it won't work properly. You can't use 802.1Q on one side and ISL on the other. C is incorrect because the native VLAN doesn't necessarily have to be VLAN 1.
C: By default, the native VLAN is VLAN 1 but this can be effectively changed to a different VLAN and the trunk will still be functional.
Reference: http://www.cisco.com/warp/public/473/27.html

QUESTION 56
A Certkiller switch is configured for 802.1Q trunking. What are valid characteristics of IEEE 802.1Q? (Select all that apply)

A. Use frame tagging.
B. None of the answers
C. It is a method for identifying VLANs
D. It inserts VLAN identifier into the frame header

Answer: A, C, D

Explanation:
802.1Q uses an internal tagging mechanism. Internal means that a tag is inserted within the frame (with ISL, the frame is encapsulated instead):

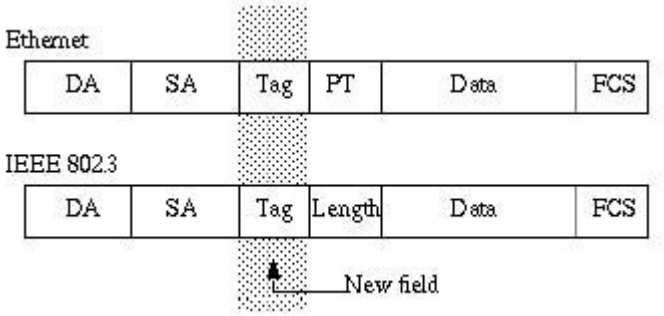

Note that on an 802.1Q trunk, one VLAN is NOT tagged. This VLAN, named the native VLAN, must be configured the same on each side of the trunk. This way, we can deduce to which VLAN a frame belongs when we receive a frame with no tag.
The tagging mechanism implies a modification of the frame; the trunking device inserts a 4-byte tag and recomputes the frame check sequence (FCS):

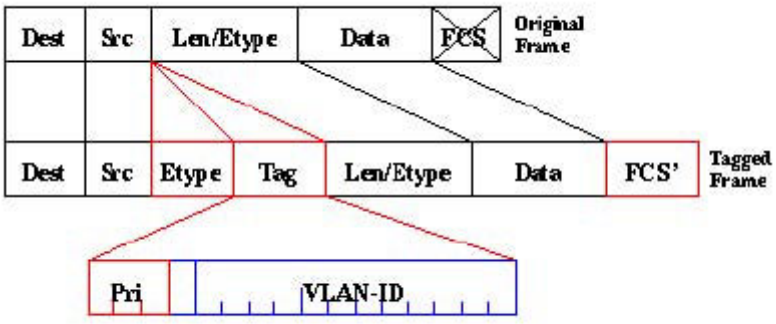

The EtherType field identifying the 802.1Q frame is 0x8100. In addition to the 12-bit VLAN-ID, 3 bits are reserved for 802.1p priority tagging.
Also, note that inserting a tag into a frame that already has the maximum Ethernet size creates a 1522 byte frame that can be considered as a "baby giant" by the receiving equipment. The 802.3 committee is extending the maximum standard frame size to address this issue.
.

QUESTION 57
What are the reasons as to why an administrator would deploy Dynamic Trunking Protocol (DTP) on the Certkiller switched LAN? (Select all that apply)

A. For supporting auto-negotiation of IEEE 802.1Q trunks
B. For supporting auto-negotiation of ISL
C. For managing trunk negotiation in 2500 router supervisor engine software R 4.2 or later
D. For managing trunk negotiation in Catalyst supervisor engine software R 4.2 or later.
E. None of the above.

Answer: A, B, D

Explanation:
DTP was developed for the specific purpose of supporting automatic trunk negotiation for 802.1Q and ISL trunks. It is used only with Cisco Catalyst switches.
Incorrect Answers:
DTP is supported only on Cisco Catalyst switches. It is not supported on Cisco 2500 series routers.

QUESTION 58
A fast Ethernet port on switch CK1 is configured as a trunk. What is true of this trunk link?

A. A trunk link only supports the native VLAN for a given port.
B. A trunk link uses 802.10 to identify VLANs over an Ethernet backbone.
C. A trunk link connects multiple devices on a single subnet to a switch port.

D. The native VLAN of the trunk link is the VLAN to which the port will belong if that link becomes non-trunk.
E. All of the above.

Answer: C

Explanation:
Trunks are used to connect multiple VLANs together. Individual switches configured with VLANs over the entire LAN subnet are connected to each other via a trunk port.
Incorrect Answers:
A: By default all VLANs within the range of 1 to 1000 is allowed to traverse the trunk port.
B: 802.10 is the standard used on FDDI networks and is not related to Ethernet VLAN trunks.
D: This is wrong because Native VLAN Number of the native VLAN for the trunk link (for 802.1Q trunks, the VLAN for which untagged traffic can be transmitted and received over the trunk; for ISL trunks, packets are tagged on all VLANs, including the native VLAN).

QUESTION 59
You are a technician at Certkiller and your newly appointed trainee asks you what the Dynamic Trunking Protocol (DTP) mode 'desirable' means. What would your reply be?

A. The interface is put into permanent trunking mode but prevented from generating DTP frames.
B. The interface actively attempts to convert the link to a trunk link.
C. The interface is put into a passive mode, waiting to convert the link to a trunk link.
D. The interface is put into permanent trunking mode and negotiates to convert the link into a trunk link.

Answer: B

Explanation:
The DTP mode of desirable configured the trunk port to actively attempt to convert the link to a trunk link.
Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 105

QUESTION 60
What would happen to a frame if a VLAN port configured as a trunk on the Catalyst switch CK1 were to receive an untagged frame?

A. The frame will cause an error message to be sent.
B. The frame will be dropped.
C. The frame will be processed as a native VLAN frame.
D. The frame will be tagged, and then processed as a native VLAN frame.

Answer: C

Explanation:
On an IEEE 802.1Q trunk port, all transmitted and received frames are tagged except for those on the VLAN configured as the native VLAN for the port. Frames on the native VLAN are always transmitted untagged and are normally received untagged. The default native VLAN is VLAN 1.
Reference:
http://www.cisco.com/en/US/products/hw/optical/ps2006/products_module_configuration_guide_chapter09186a

---

QUESTION 61
Switch CK1 has a trunk link configured with IEEE 802.1Q encapsulation. What is the maximum Ethernet frame size on this trunk port?

A. 1496 Bytes
B. 1500 Bytes
C. 1518 Bytes
D. 1522 Bytes
E. 1548 Bytes

Answer: D

Explanation:
The 802.1q tag is 4 bytes; hence the resulting ethernet frame can be as large as 1522 bytes (1518 for the maximum Ethernet frame size plus the 4 byte 802.1Q tag). The minimum size of the Ethernet frame with 802.1q tagging is 68 bytes.
Reference:
http://www.cisco.com/en/US/tech/ CK3 89/ CK3 90/technologies_tech_note09186a0080094665.shtml

---

QUESTION 62
The original frame is encapsulated and an additional header is added before the frame is carried over a trunk link. At the receive end, the header is removed and the frame is forwarded to the assigned VLAN. This describes which technology?

A. DISL
B. DTP
C. IEEE802.1Q
D. ISL
E. MPLS

Answer: D

QUESTION 63
Assuming you have an IOS based switch; which command would you execute if you

wanted to specify IEEE 802.1Q encapsulation on a trunked port?

A. Switch(config-if)#switchport trunk encapsulation dot1q
B. Switch(config-if)#switchport encapsulation dot1q
C. Switch(config-if)#switchport trunk encapsulation isl
D. Switch(config)#switchport 0/1 trunk encapsulation isl
E. None of the above

Answer: A

Explanation:
Ethernet Trunk Encapsulation Types:
1. switchport trunk encapsulation isl - Specifies ISL encapsulation on the trunk link.
2. switchport trunk encapsulation dot1q - Specifies 802.1Q encapsulation on the trunk link.
3. switchport trunk encapsulation negotiate - Specifies that the interface negotiate with the neighboring interface to become an ISL (preferred) or 802.1Q trunk, depending on the configuration and capabilities of the neighboring interface.
The trunking mode, the trunk encapsulation type, and the hardware capabilities of the two connected interfaces determine whether a link becomes an ISL or 802.1Q trunk.
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps700/products_configuration_guide_chapter09186a008007f

---

QUESTION 64
Two Certkiller switches are connected as shown below:

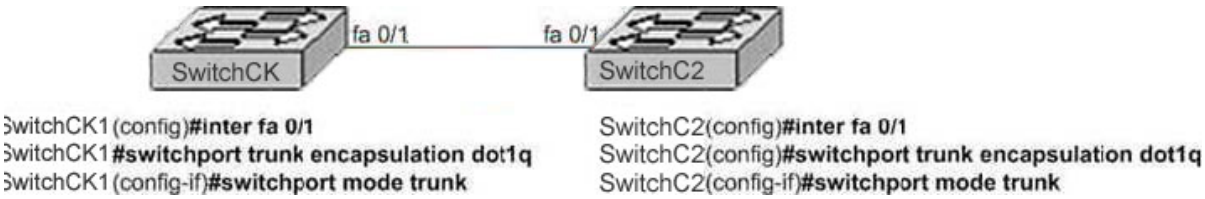

Which statements are true regarding the configuration of the above pair of switches? (Select two)

A. The trunk is currently using the ISL trunking protocol.
B. The trunk is currently using the 802.1q trunking protocol.
C. By default, all VLANs will be transmitted across this trunk.
D. By default, Switch CK1 and Switch CK2 's Fast Ethernet 0/1 port will not generate DTP messages.
E. By default, the trunk can only support one VLAN, and only that single VLAN is transmitted across the trunk.

Answer: B, C

Explanation:
The second line in each configuration (#switchport trunk encapsulation dot1q) proves

that B is correct, as dot1q is Cisco IOS for 802.1q trunking.
Since the interface fa/01 is configured (#interface fa 0/1) and the mode is set to trunk (#switchport mode trunk) in both switches, there is no need for dynamic trunking protocol since the trunk is already set. By default, all VLANs will be able to cross the trunk, unless explicitly configured not to do so.

QUESTION 65
Switches Certkiller 1 and Certkiller 2 are connected as shown in the diagram below:

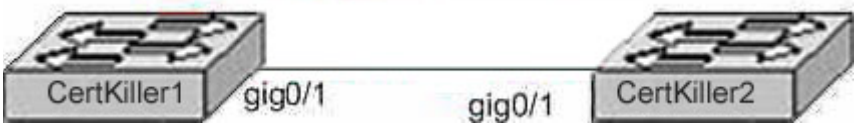

Use the following output taken from each port
Certkiller 1:
show config:
interface GigabitEthernet0/1
switchport trunk encapsulation dot1q
switchport mode dynamic auto
no ip address
show interface gig0/1 switchport:
Name: Gi0/1
Switchport: Enabled
Administrative Mode: dynamic auto
Operational Mode: down
Administrative Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Certkiller 2:
show interface gig0/1 switchport:
Name: Gi0/1
Switchport Enabled
Administrative Mode: dynamic auto
Operational Mode: down
Administrative Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
In accordance with the above exhibit: what's preventing the two switches from trunking on the link between them?

A. There is no IP address denied.
B. no shutdown needs to be entered on the interfaces.
C. Both sides are in auto negotiation mode.
D. ISL should be used instead of 802.1q.
E. Access mode VLAN must be different from native mode VLAN.

Answer: C

According to Cisco table Auto & Auto results in NO trunk formation. At least one end of the trunk should be set to on or desirable in order for the trunk to operate correctly.

QUESTION 66
You have just configured an ISL trunk line over Ethernet media between two Cisco Switches, each switch having identical modules, software revisions, and VLAN configurations. Which of the following variables are NOT necessary for the ISL trunk to operate properly? (Select all that apply)

A. Identical trunk negotiation parameters at each end of the link
B. Identical duplex at each end of the link
C. Identical speed at each end of the link
D. Identical native VLAN parameters at each end of the link

Answer: A, D

Explanation:
In order for a trunk connection to function properly, it is not necessary for the trunking negotiation parameters to be identical. For example, one end could be configured as "on" and the other could be configured for "auto-negotiate" and the trunk would be operational. Similarly, it is not necessary for the native VLAN parameters to be the same at each end.
Incorrect Answers:
B, C: One of the requirements for trunking to work is to have both sides of the trunk agree on the speed and duplex settings. Both sides must be configured with identical speed and duplex settings.

QUESTION 67
An ISL trunk connects switches CK1 and CK2 . What is the numerical range of user-configurable ISL VLANs on these switches?

A. 1-1001
B. 0-4095
C. there is no range
D. 0 - 1000
E. None of the above

Answer: A

Explanation:
The valid range of user-configurable ISL VLANs is 1-1001. The valid range of VLANs specified in the IEEE 802.1Q standard is 0-4095. In a network environment with non-Cisco devices connected to Cisco switches through 802.1Q trunks, you must map 802.1Q VLAN numbers greater than 1000 to ISL VLAN numbers. 802.1Q VLANs in the range 1-1000 are automatically mapped to the corresponding ISL VLAN. 802.1Q VLAN

numbers greater than 1000 must be mapped to an ISL VLAN in order to be recognized and forwarded by Cisco switches.

QUESTION 68
An ISL trunk connects switches CK1 and CK2 . What is true about the Inter-Switch Link (ISL) protocol? (Select two)

A. ISL can be used between Cisco and non-Cisco switch devices.
B. ISL calculates a new CRC field on top of the existing CRC field.
C. ISL adds 4 bytes of protocol-specific information to the original Ethernet frame.
D. ISL adds 30 bytes of protocol-specific information to the original Ethernet frame.

Answer: B, D

Explanation:
B: A second frame check sequence (FCS) field lies at the end of the frame.
D: ISL is an external tagging process: new 26-byte ISL header is added to the original Ethernet frame. A second 4-byte frame check sequence (FCS) field is added at the end of the frame so 30 bytes of total overhead is added.
Incorrect Answers:
A: Cisco's propriety version of frame tagging is ISL. ISL can only be used between Cisco routers.
C: 30 bytes are added to the Ethernet frame, not 4 bytes. 4 bytes are added using 802.1Q encapsulation.

QUESTION 69
Which of the commands below enables a trunking protocol that appends a four byte CRC to the packet when applied to the Certkiller switch?

A. Switch(config-if)#switchport trunk encapsulation dot1q
B. Switch(config-if)#switchport trunk encapsulation ietf
C. Switch(config-if)#switchport trunk encapsulation fddi
D. Switch(config-if)#switchport trunk encapsulation isl
E. None of the above

Answer: D

Explanation:
ISL is made up of three major components: a header, the original Ethernet frame, and a frame check sequence (FCS) at the end. With ISL, an Ethernet frame is encapsulated with a header that transports VLAN IDs between switches and routers. The 26-byte header containing a 10-bit VLAN ID is added to each frame. In addition, a 4-byte tail is added to the frame to perform a cyclic redundancy check (CRC). This CRC is in addition to any frame checking that the Ethernet frame performs.
Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 99

QUESTION 70
Which statement is true regarding the configuration of ISL trunks?

A. All catalyst switches support ISL trunking.
B. A Catalyst switch will report giants if one side is configured for ISL while the other side is not.
C. ISL trunking requires that native VLANs match.
D. A Catalyst switch cannot have ISL and IEEE 802.1q trunks enabled.
E. None of the above

Answer: B

Explanation:
The 802.1q tag is 4 bytes; hence the resulting ethernet frame can be as large as 1522 bytes. The minimum size of the Ethernet frame with 802.1q tagging is 68 bytes.
ISL frames are the standard MTU used in Ethernet frames, which is 1518 bytes. If one end of the trunk is configured for ISL frames of up to 1518 bytes will be expected on it, while the other end will be sending frames up to 1522 bytes in length. On the ISL configured end, these incoming frames will be considered as giants. This is just one of the reasons why ISL and 802.1Q are incompatible.
Incorrect Answers:
A: Not every Cisco switch model supports ISL.
C: In ISL, it is not necessary for the native VLANs to match.
D: Although it is true that each end of a trunk should be configured using the same encapsulation types, it is possible for a switch to have an ISL trunk configured on one port and an 802.1Q trunk on another port.

QUESTION 71
Two Certkiller switches are connected via a trunk using VTP. Which VTP information does a Catalyst switch advertise on its trunk ports when using VTP? (Select two)

A. STP root status
B. VTP mode
C. Negotiation status
D. Management domain
E. Configuration revision number

Answer: D, E

Explanation:
The role of the VLAN Trunking Protocol (VTP) is to maintain VLAN configuration consistency across the entire network. VTP is a messaging protocol that uses Layer 2 trunk frames to manage the addition, deletion, and renaming of VLANs on a network-wide basis from a centralized switch that is in the VTP server mode. VTP is responsible for synchronizing VLAN information within a VTP domain. This reduces the

need to configure the same VLAN information on each switch.
Using VTP, each Catalyst Family Switch advertises the following on its trunk ports:
1. Management domain
2. Configuration revision number
3. Known VLANs and their specific parameters

QUESTION 72
You need to investigate a VTP problem between two Certkiller switches. The lack of which two prevents VTP information from propagating between switches? (Select two)

A. A root VTP server
B. A trunk port
C. VTP priority
D. VLAN 1
E. None of the above

Answer: B, D

Explanation:
In Switch tow types of links are available, access and trunk. The interface is in access mode can carry the information of only one VLAN and trunk can carry the information of more than one VLAN. VTP carry the information of more than one vlan so Switch port should be in trunk mode. VLAN1 is the default VLAN on Cisco Switch, by default all interface belongs to VLAN 1.

QUESTION 73
CK1 and CK2 are switches that communicate via VTP. What is the default VTP advertisement intervals in Catalyst switches that are in server or client mode?

A. 30 seconds
B. 5 minutes
C. 1 minute
D. 10 seconds
E. 5 seconds
F. None of the above

Answer: B

Explanation:
Periodic (default is 5 minutes) VTP advertisements are sent out each trunk port with the multicast destination MAC address 01-00-0C-CC-CC-CC. VTP advertisements contain the following configuration information:
1. VLAN IDs (ISL and 802.1Q)
2. Emulated LAN names (ATM LANE)
3. 802.10 SAID values (FDDI)

4. VTP domain name
5. VTP configuration revision number
6. VLAN configuration, including the maximum transmission unit (MTU) size for each VLAN
7. Frame format

QUESTION 74
On the network shown below, VTP has been enabled on the trunk links between all switches within the TEST domain. An administrator has recently enabled VTP pruning. Port 1 on Switch Certkiller 1 and port 2 on Switch Certkiller 4 are assigned to VLAN 2. A broadcast is sent from the host connected to Switch Certkiller 1. Where will the broadcast propagate?

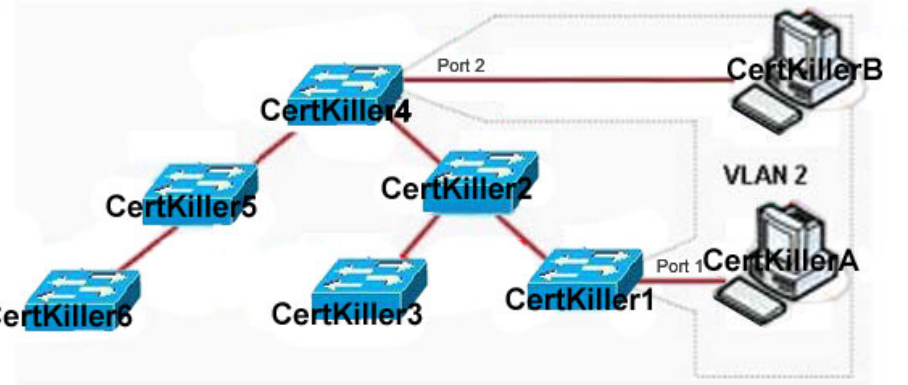

A. Switches Certkiller 1, Certkiller , and Certkiller 4 will receive the broadcast, but only Switch Certkiller 4 will forward it out port 2.
B. Only Switch Certkiller 4 will receive the broadcast and will forward it out port 2.
C. Every switch in the network receives the broadcast and will forward it out all ports.
D. Every switch in the network receives the broadcast, but only Switch Certkiller 4 will forward it out port 2.

Answer: A

Explanation:
The default behavior of a switch is to propagate broadcast and unknown packets across the network. This behavior results in a large amount of unnecessary traffic crossing the network.
VTP pruning increases bandwidth efficiency by reducing unnecessary flooding of traffic, such as broadcast, multicast, unknown, and flooded unicast packets. VTP pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the appropriate network devices. By default, VTP pruning is disabled.
Enabling VTP pruning on a VTP server enables pruning for the entire management domain. VTP pruning takes effect several seconds after it is enabled. By default, VLANs 2 through 1000 or 2 through 1001 are pruning eligible, depending upon the platform.
VTP pruning does not prune traffic from VLANs that are pruning ineligible. VLAN 1 is always pruning ineligible and VLAN 1 cannot be removed from a trunk. However, the

"VLAN 1 disable on trunk" feature available on Catalyst 4000, 5000, and 6000 family switches enables the pruning of user traffic, but not protocol traffic such as CDP and VTP, for VLAN 1 from a trunk. Use the vtp pruning command to make VLANs pruning eligible on a Cisco IOS-based switch.
Switch(vlan)#vtp pruning
Once pruning is enabled, use the switchport trunk pruning command to make a specific VLAN pruning ineligible.
Switch(config)#interface fastethernet 0/3
Switch(config-if)#switchport trunk pruning vlan remove vlan 5

QUESTION 75
You want to configure switch CK1 to propagate VLAN information across the Certkiller network using VTP. What must be configured on a Cisco switch in order to advertise VLAN information?

A. VTP mode
B. VTP password
C. VTP revision number
D. VTP pruning
E. VTP domain name
F. None of the above

Answer: E

Explanation:
If the switch being installed is the first switch in the network, the management domain will need to be created. However, if the network has other switches running VTP, then the new switch will join an existing management domain. Verify the name of the management domain. If the management domain has been secured, verify and configure the password for the domain.
To create a management domain or to add a switch to a management domain, use the vtp domain command in the global configuration mode or VLAN configuration mode.
Switch(config)#vtp domain name
Switch(vlan)#vtp domain

QUESTION 76
The Certkiller switches have all been upgraded to use VTP version 2. What are two benefits provided in VTP Version 2 that are not available in VTP Version 1? (Select two)

A. VTP version 2 supports Token Ring VLANs
B. VTP version 2 allows VLAN consistency checks
C. VTP version 2 allows active redundant links when used with spanning tree
D. VTP version 2 reduces the amount of configuration necessary
E. VTP version 2 saves VLAN configuration memory

Answer: A, B

Explanation:
Two different versions of VTP can run in the management domain, VTP Version 1 and VTP Version 2. The two versions are not interoperable in the same VTP domain. The major difference between the two versions is version 2 introduces support for Token Ring VLANs.
If all switches in a VTP domain can run VTP Version 2, version 2 only needs to be enabled on one VTP server switch. The version number is propagated to the other VTP Version 2-capable switches in the VTP domain. Version 2 should not be enabled unless every switch in the VTP domain supports version 2.
The VTP version can be configured from global configuration mode or VLAN database mode on a Cisco IOS-based switch. From there, the VTP version can be changed with the vtp command.
VTP Configuration in global configuration mode:
Switch#config terminal
Switch(config)#vtp version 2
VTP Configuration in VLAN configuration mode:
Switch#vlan database
Switch(vlan)#vtp v2-mode

QUESTION 77
The Certkiller network administrator needs to enable VTP pruning within the Certkiller network. What action should a network administrator take to enable VTP pruning on an entire management domain?

A. Enable VTP pruning on any switch in the management domain
B. Enable VTP pruning on any client switch in the domain
C. Enable VTP pruning on a VTP server in the management domain
D. Enable VTP pruning on every switch in the domain
E. None of the above

Answer: C

Explanation:
The default behavior of a switch is to propagate broadcast and unknown packets across the network. This behavior results in a large amount of unnecessary traffic crossing the network.
VTP pruning increases bandwidth efficiency by reducing unnecessary flooding of traffic, such as broadcast, multicast, unknown, and flooded unicast packets. VTP pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the appropriate network devices. By default, VTP pruning is disabled.
Enabling VTP pruning on a VTP server enables pruning for the entire management domain. VTP pruning takes effect several seconds after it is enabled. By default, VLANs 2 through 1000 or 2 through 1001 are pruning eligible, depending upon the platform.

VTP pruning does not prune traffic from VLANs that are pruning ineligible. VLAN 1 is always pruning ineligible and VLAN 1 cannot be removed from a trunk. However, the "VLAN 1 disable on trunk" feature available on Catalyst 4000, 5000, and 6000 family switches enables the pruning of user traffic, but not protocol traffic such as CDP and VTP, for VLAN 1 from a trunk. Use the vtp pruning command to make VLANs pruning eligible on a Cisco IOS-based switch.
Switch(vlan)#vtp pruning
Once pruning is enabled, use the switchport trunk pruning command to make a specific VLAN pruning ineligible.
Switch(config)#interface fastethernet 0/3
Switch(config-if)#switchport trunk pruning vlan remove vlan 5

QUESTION 78
VTP is configured on switch CK1 . Which of the following features were added in VTP version 2 that were not previously supported in VTP version 1? (Select two)

A. Supports Token Ring VLANs.
B. Allows VLAN consistency checks.
C. Saves VLAN configuration memory.
D. Reduces the amount of configuration necessary.
E. Allows active redundant links when used with spanning tree.

Answer: A, B

Explanation:
VTP Version 2 includes the following improvements: Token Ring VLAN support, TLV support, transparent mode, and Consistency checks.
Incorrect Answers:
C, D: These were not improvements added to VTP Version 2.
E: STP detects and prevents loops by logically disabling the redundant path ports so there are no active redundant links.

QUESTION 79
The Certkiller switches are configured to use VTP. What's true about the VLAN trunking protocol (VTP)? (Select two)

A. VTP messages will not be forwarded over nontrunk links.
B. VTP domain names need to be identical. However, case doesn't matter.
C. A VTP enabled device which receives multiple advertisements will ignore advertisements with higher configuration revision numbers.
D. A device in "transparent" VTP v.1 mode will not forward VTP messages.
E. VTP pruning allows switches to prune VLANs that do not have any active ports associated with them.

Answer: A, D

Explanation:
VTP messages are only transmitted across trunk links.
If the receiving switch is in transparent mode, the configuration is not changed. Switches in transparent mode do not participate in VTP. If you make VTP or VLAN configuration changes on a switch in transparent mode, the changes are not propagated to the other switches in the network.
Incorrect Answers:
B: The VTP domain name is case sensitive and it must be identical with the domain name configured on the VTP server.
C: This is incorrect because if a VTP client receives an advertisement with a higher revision number, it won't ignore it. In fact, the advertisement with a higher revision level takes precedence when the switch is configured in client mode.
E: VTP pruning enhances network bandwidth use by reducing unnecessary flooded traffic, such as broadcast, multicast, unknown, and flooded unicast packets. VTP pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the appropriate network devices. It does not prune the individual VLANs.

QUESTION 80
Switch CK1 and CK2 both belong to the Certkiller VTP domain. What's true about the switch operation in VTP domains? (Select all that apply)

A. A switch can only reside in one management domain
B. A switch is listening to VTP advertisements from their own domain only
C. A switch is listening to VTP advertisements from multi domains
D. A switch can reside in one or more domains
E. VTP is no longer supported on Catalyst switches

Answer: A, B

Explanation:
A VTP domain is made up of one or more interconnected devices that share the same VTP domain name. A switch can be configured to be in only one VTP domain, and each VLAN has a name that is unique within a management domain.
Typically, you use a VTP domain to ease administrative control of your network or to account for physical boundaries within your network. However, you can set up as many or as few VTP domains as are appropriate for your administrative needs. Consider that VTP is transmitted on all trunk connections, including ISL, IEEE 802.1Q, 802.10, and LANE.
Switches can only belong to one management domain with common VLAN requirements, and they only care about the neighbors in their own domains.
Reference:
CCNP Switching Exam Certification Guide: David Hucaby & Tim Boyles, Cisco Press 2001, ISBN 1-58720 000-7 page 114

QUESTION 81
VTP devices in a network track the VTP revision number. What is a VTP configuration revision number?

A. A number for identifying changes to the network switch.
B. A number for identifying changes to the network router.
C. A number for identifying changes to the network topology.
D. None of the above.

Answer: C

Explanation:
The configuration revision number is a 32-bit number that indicates the level of revision for a VTP packet. Each VTP device tracks the VTP configuration revision number assigned to it, and most of the VTP packets contain the VTP configuration revision number of the sender.
This information is used to determine whether the received information is more recent than the current version. Each time you make a VLAN change in a VTP device, the configuration revision is incremented by one. In order to reset the configuration revision of a switch, change the VTP domain name and then change it back to the original name.
Incorrect Answers:
A: Not all switch configuration changes will impact the VTP revision number. Only changes made to the VLAN configuration will cause an increment in the revision number.
B: VTP revision numbers are only used on network switches configured for VTP and are not used by Cisco routers.
Reference: Understanding and Configuring VLAN trunk protocol (VTP) Document ID: 10558http://www.cisco.com/warp/public/473/21.html

---

QUESTION 82
Switch CK1 is configured to use the VLAN Trunking Protocol (VTP). What does CK1 advertise in its VTP domain?

A. The VLAN ID of all known VLANs, the management domain name, and the total number of trunk links on the switch.
B. The VLAN ID of all known VLANs, a 1-bit canonical format (CF1 Indicator), and the switch configuration revision number.
C. The management domain name, the switch configuration revision number, the known VLANs, and their specific parameters.
D. A 2-byte TPID with a fixed value of 0x8100 for the management domain number, the switch configuration revision number, the known VLANs, and their specific parameters.
E. None of the above.

Answer: C

Explanation:

"Each switch participating in VTP advertises VLAN information, revision numbers, and VLAN parameters on its trunk ports to notify other switches in the management domain. VTP advertisements are sent as multicast frames. The switch intercepts frames sent to the VTP multicast address and processes them with its supervisory processor VTP frames are forwarded out trunk links as a special case.
The following global configuration information is distributed in VTP advertisements:
1. VLAN IDs (ISL and 802.1Q)
2. Emulated LAN names (for ATM LANE)
3. 802.10 SAID values (FDDI)
4. VTP domain name
5. VTP configuration revision number
6. VLAN configuration, including maximum transmission unit (MTU) size for each VLAN
7. Frame format
Reference: CCNP Switching Exam Certification Guide: page 115, David Hucaby & Tim Boyles, Cisco Press 2001, ISBN 1-58720 000-7
Incorrect Answers:
A: The total number of trunk links is not advertised.
B: A CFI is not advertised.
D: The TPID is not advertised. The value of 0x8100 is used to identify an 802.1Q trunking tag.

QUESTION 83
VTP switches use advertisements to exchange information with each other. Which of the following advertisement types are associated with VTP? (Select all that apply)

A. Domain advertisements
B. Advertisement requests from clients
C. Subset advertisements
D. Summary advertisements

Answer: B, C, D

Explanation:
VTP advertisements include:
1. Summary Advertisements - These go out every 5 minutes or ever time the VLAN topology changes, and lists of information about the management domain (VTP version, domain name, configuration revision number, timestamp, MD5 encryption hash code, & number of subset advertisements incoming). When there is a configuration change, summary advertisements are complimented by or more subset advertisements.
2. Subset advertisements - These are sent out by VTP domain servers after a configuration change. They list the specifics of the change (VLAN creation / deletion / suspension / activation / name change / MTU change) and the VLAN parameters (VLAN status, VLAN type, MTU, VLAN name, VLAN number, SAID value).
3. Avertisement Requests from Clients- VTP clients request specific VLAN information that they're lacking (ie. Client switch is reset and loses its database, or VTP domain

membership changes) so they can be responded by summary and subset advertisements.
Reference: CCNP Switching Exam Certification Guide: pages 116-117 David Hucaby & Tim Boyles, Cisco Press 2001, ISBN 1-58720 000-7

QUESTION 84
Switch CK1 is part of the Certkiller VTP domain. What's true of VTP Pruning within this domain? (Select all that apply)

A. it does not prune traffic from VLANs that are pruning-ineligible
B. VLAN 1 is always pruning-eligible
C. it will prune traffic from VLANs that are pruning-ineligible
D. VLAN 2 is always pruning-ineligible
E. None of the above.

Answer: A

Explanation:
By definition, pruning-ineligible VLANs can not be pruned. You can make specific VLANs pruning ineligible with the clear vtp pruneeligible vlan_range command. By default, VLANs 2-1000 are pruning-eligible. Since the default VLAN for any switch port in a Catalyst switch is VLAN 1, it is not eligible for pruning.
Incorrect Answers:
B: VLAN 1 is always pruning-ineligible
C: The opposite is true.
D: By default, VLANs 2-1000 are eligible to be pruned.

QUESTION 85
What action should you execute if you wanted to enable VTP pruning on your entire management domain?

A. Enable VTP pruning on any client switch in the management domain.
B. Enable VTP pruning on any switch in the management domain.
C. Enable VTP pruning on every switch in the management domain.
D. Enable VTP pruning on a VTP server in the management domain.
E. Disable VTP pruning on a VTP server in the management domain.

Answer: D

Explanation:
Enabling VTP pruning on a VTP server allows pruning for the entire management domain. Enabling this on the VTP server will mean that the VTP pruning configuration will be propagated to all VTP client switches within the domain. VTP pruning takes effect several seconds after you enable it. By default, VLANs 2 through 1000 are pruning-eligible.
Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 117

QUESTION 86
Switch CK1 is configured with VTP. Which two VTP modes will make CK1 capable of creating and deleting VLANs on itself? (Select two)

A. Client
B. Server
C. Transparent
D. Pass-through
E. No-negotiate

Answer: B, C

Explanation:
VTP Modes
You can configure a switch to operate in any one of these VTP modes:
1. Server-In VTP server mode, you can create, modify, and delete VLANs and specify other configuration parameters (such as VTP version) for the entire VTP domain. VTP servers advertise their VLAN configuration to other switches in the same VTP domain and synchronize their VLAN configuration with other switches based on advertisements received over trunk links. VTP server is the default mode.
2. Client-VTP clients behave the same way as VTP servers, but you cannot create, change, or delete VLANs on a VTP client.
3. Transparent-VTP transparent switches do not participate in VTP. A VTP transparent switch does not advertise its VLAN configuration and does not synchronize its VLAN configuration based on received advertisements. However, in VTP version 2, transparent switches do forward VTP advertisements that they receive out their trunk interfaces.
If you configure the switch as VTP transparent, you can create and modify VLANs but the changes affect only the individual switch.
Incorrect Answers:
A: Clients can not modify, add, or delete any VLAN information.
D, E: These options are not valid VTP modes.

QUESTION 87
When the Catalyst switch CK1 is enabled to use VTP, which information does it advertise on its trunk ports? (Select two)

A. VTP mode
B. STP root status
C. Negotiation status
D. Management domain
E. Configuration revision number

Answer: D, E

Explanation:
The VTP protocol maintains VLAN configuration consistency throughout the network by

distributing VLAN information to the network. VLAN information is sent to network devices in advertisements that contain the VTP management domain name, the current configuration revision number, the VLANs that the server knows about, and certain VLAN parameters. Any time you change a VLAN, VTP automatically sends an advertisement to update all other network devices.
The following global configuration information is distributed in VTP advertisements:
1. VLAN IDs (ISL and 802.1Q)
2. Emulated LAN names (for ATM LANE)
3. 802.10 SAID values (FDDI)
4. VTP domain name
5. VTP configuration revision number
6. VLAN configuration, including maximum transmission unit (MTU) size for each VLAN
7. Frame format
Reference: Cisco, Configuring VTP
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/rel_6_1/config/vtp.htm

---

QUESTION 88
Is the following statement True or False?
MLS requires that MLS components be in the same VTP domain.

A. False
B. There is not enough information to determine
C. True
D. It could be true or false, depending on the type of switch.

Answer: C

Explanation:
MLS requires that MLS components, including the end stations, must be in the same Virtual Trunking Protocol (VTP) domain. VTP is a Layer 2 protocol used for managing VLANs on several Catalyst switches from a central switch. It allows an administrator to create or delete a VLAN on all switches in a domain without having to do so on every switch in that domain. The MultiLayer Switching Protocol (MLSP), which the MLS-SE and the MLS-RP use to communicate with one another, does not cross a VTP domain boundary.

QUESTION 89
Which of the following VTP modes receives and forwards VTP updates, but does NOT participate in VTP synchronization?

A. Client
B. Server
C. Transparent
D. Pass-through
E. None of the above

Answer: C

Explanation:
Transparent-VTP transparent network devices do not participate in VTP. A VTP transparent network device does not advertise its VLAN configuration and does not synchronize its VLAN configuration based on received advertisements. However, in VTP version 2, transparent network devices do forward VTP advertisements that they receive out their trunking LAN ports.

QUESTION 90
How can VTP pruning enhance network bandwidth?

A. By limiting the spreading of VLAN information.
B. By reducing unnecessary flooding of traffic to inactive VLANs.
C. By disabling periodic VTP updates.
D. By restricting unicast traffic to across VTP domains.
E. By updating unicast traffic periodically.
F. None of the above

Answer: B

Explanation:
VTP pruning enhances network bandwidth use by reducing unnecessary flooded traffic, such as broadcast, multicast, unknown, and flooded unicast packets. VTP pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the appropriate network devices.
Reference:
http://www.cisco.com/en/US/products/hw/routers/ps368/products_configuration_guide_chapter09186a00800916

QUESTION 91
Which of the following statements is NOT true regarding VTP?

A. Switches in VTP transparent mode will simply forward advertisements without processing them.
B. VTP reduces administrative overhead.
C. VTP pruning reduces overall network traffic.
D. VTP pruning is on by default.
E. All of the above are true statements.
F. None of the above

Answer: D

Explanation:
By default, VTP pruning is disabled.

For VTP pruning to be effective, all devices in the management domain must either support VTP pruning or, on devices that do not support VTP pruning, you must manually configure the VLANs allowed on trunks.
Incorrect Answers:
A: This statement is true. Transparent VTP switches do not participate in the VTP process, but they do forward VTP information to other switches.
B, C: VTP pruning enhances network bandwidth use by reducing unnecessary flooded traffic, such as broadcast, multicast, unknown, and flooded unicast packets. VTP pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the appropriate network devices.

QUESTION 92
By what condition can a VTP version 2 switch operate in the same domain as a switch running VTP version 1?

A. VTP version 2 is disabled on the VTP version 2-capable switch
B. VTP version 2 is enabled on the VTP version 2-capable switch
C. VTP version 1 is disabled on the VTP version 2-capable switch
D. None of the above. VTP version 1 and version 2 are not compatible.
E. None of the above

Answer: A

Explanation:
A VTP version 2-capable switch can operate in the same VTP domain as a switch running VTP version 1 provided VTP version 2 is disabled on the VTP version 2-capable switch (VTP version 2 is disabled by default). With VTP version 2 disabled, the switch will revert to version 1 to become backward compatible.

QUESTION 93
Is the following statement True or False?
If you modified a VTP transparent switch, the changes you implement will affect all the switches in the network?

A. True
B. There is not enough information to determine
C. False

Answer: C

Explanation:
According to Cisco:
If you configure the switch as VTP transparent, you can create and modify VLANs but the changes affect only the individual switch. Transparent switches do not participate in VTP. Only changes made to a VTP switch in server mode will be propagated to all the other client switches within the network.

QUESTION 94
Which of the following message types are associated with VTP header fields on a Cisco switched network? (Select all that apply)

A. Summary advertisements
B. Advertisement requests
C. VTP Join messages
D. Subset advertisement
E. None of the above.

Answer: A, B, C, D

Explanation:
The format of the VTP header can vary depending on the type of VTP message. However, they all contain the following fields in the header:
VTP protocol version: 1 or 2
VTP Message Types: Summary advertisements, Subset advertisement, Advertisement requests, VTP join messages, Management domain length, and Management domain name.

QUESTION 95
In order for the Certkiller network to use VTP, which of the following conditions have to be met? (Select all that apply)

A. Trunking must be enabled between all Catalyst switches.
B. The Catalyst switches must be non-adjacent for trunking to be possible between them
C. The Catalyst switches must be adjacent.
D. Each Catalyst switch in a domain should be assigned the same VTP domain name.
E. None of the above

Answer: A, C, D

Explanation:
According to the online documentation provided by Cisco:
In order to use VTP, you must assign a VTP domain name to each switch. VTP information will remain only within the same VLAN domain. The following are conditions for a VTP domain:
-Each Catalyst switch in a domain should be assigned the same VTP domain name.
-The Catalyst switches must be adjacent.
-Trunking must be enabled between all Catalyst switches.
If any one of the previous conditions is not met, the VTP domain is broken and information will not travel between the two separate parts.

QUESTION 96
The Certkiller switches are all VTP enabled. What is true about VTP? (Select all

that apply)

A. VTP version 2 is supported in supervisor engine software release 3.1(1) and later.
B. you must decide whether to use VTP version 1 or version 2.
C. VTP version 1 is supported in supervisor engine software release 2.1 or later
D. VTP version 1 is supported in ATM software release 3.1 or later.
E. None of the above

Answer: A, B, C, D

Explanation:
According to Cisco Documentation:
If you use VTP in your network, you must decide whether to use VTP version 1 or version 2. VTP version 1 is supported in supervisor engine software release 2.1 or later and ATM software release 3.1 or later. VTP version 2 is supported in supervisor engine software release 3.1(1) and later.
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps679/products_configuration_guide_chapter09186a008007d

---

QUESTION 97
In the Certkiller switched network VTP pruning has been enabled. What is the purpose of VTP pruning?

A. Enhancing network integrity
B. Enhancing network bandwidth use
C. Deploying AAA
D. Enhancing network security
E. Enhancing network load balancing
F. None of the above

Answer: B

Explanation:
According to Cisco:
VTP pruning enhances network bandwidth use by reducing unnecessary flooded traffic, such as broadcast, multicast, unknown, and flooded unicast packets. VTP pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the appropriate network devices.

QUESTION 98
Switches CK1 and CK2 are both configured for transparent mode in the VTP domain. Which statement accurately describes these transparent VTP switches? (Select all that apply):

A. They do not synchronize VLAN configuration based on received advertisements

B. They do not participate in VTP
C. They do not advertise VLAN configuration
D. None of the above

Answer: A, B, C

Explanation:
VTP transparent switches do not participate in VTP. A VTP transparent switch does not advertise its VLAN configuration and does not synchronize its VLAN configuration based on received advertisements. However, in VTP version 2, transparent switches do forward VTP advertisements that they receive out their trunk ports.

QUESTION 99
You are configuring switch CK1 for VTP. Which of the following are valid VTP operating modes that can be configured on CK1 ? (Select all that apply)

A. Server
B. Frontend
C. Client
D. Transparent
E. Backbone

Answer: A, C, D

Explanation:
There are only three VTP operating modes:
1) Server: These switches have full control in the creation and modification of VLANs. Servers advertise out all the VLAN information they receive, and they configure themselves in accord with whatever information they hear. Switches are in server mode by default.
2) Client: These switches listen to VTP advertisements, they modify their configuration as a result of what they hear, and they forward out VTP information to neighbouring switches; but they don't have the ability to: create, change, or delete VLANs.
3) Transparent: These switches don't participate in the VTP process. They don't advertise their VLAN configurations and they don't synchronize their database when they receive advertisements. In VTP version 1 a switch doesn't relay information it gets to the other switches but in VTP version 2 they do.

QUESTION 100
Which of the following are true regarding the default values of a switch that is configured for VTP pruning? (Select two).

A. VLAN 1-1000 are pruning-eligible
B. VLAN 2-1000 are pruning-eligible
C. VLAN 1 is pruning-eligible
D. VLAN 1 is pruning-ineligible

E. VLAN 1-1023 is pruning-eligible
F. VLAN 1-1023 is pruning-ineligible

Answer: B, D

Explanation:
By default, VLANs 2-1000 are pruning-eligible. Since the default VLAN for any switch port in a Catalyst switch is VLAN 1, it is not eligible for pruning.

QUESTION 101
Which of the following tasks are NOT functions performed by VTP switches?
(Select all that apply)

A. To reduce parallel load sharing
B. To propagate global VLAN information
C. To provide routing randomness
D. To set the trunk priority levels of adjacent switches.
E. To ensure that there is a trunk operating in the network.
F. None of the above

Answer: A, C, D, E

Explanation:
VTP is a Layer 2 messaging protocol that maintains VLAN configuration consistency by managing the addition, deletion, and renaming of VLANs within a VTP domain. A VTP domain (also called a VLAN management domain) is made up of one or more network devices that share the same VTP domain name and that are interconnected with trunks. VTP minimizes misconfigurations and configuration inconsistencies that can result in a number of problems, such as duplicate VLAN names, incorrect VLAN-type specifications, and security violations.
The fundamental function of VTP is to manage, maintain, and propagate VLAN information throughout the enterprise network. All of the choices, besides B, describe functions that are not performed by VTP.

QUESTION 102
You need to configure switch CK1 for pruning. Which VTP command would you use if you wanted to allow pruning?

A. show vtp
B. set vtp
C. set vtp domain
D. set vtp pruneeligible
E. None of the above.

Answer: D

Explanation:
Use the set vtp command to set the options for VTP.
set vtp [domain domain_name] [mode {client | server | transparent}] [passwd passwd] [pruning{enable | disable}] [v2 {enable | disable}]
The pruning keyword is used to enable or disable VTP pruning for the VTP domain. VTP pruning causes information about each pruning-eligible VLAN to be removed from VTP updates if there are no stations belonging to that VLAN out a particular switch port. Use the set vtp pruneeligible and clear vtp pruneeligible commands to specify which VLANs should or should not be pruned when pruning is enabled for the domain.
Reference:
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_5_5/cmd_refr/set_v.htm

---

QUESTION 103
If you have just configured a Catalyst switch to operate in VTP mode, and that switch is configured to not advertise VLAN configuration information. Which VTP mode has been configured on this switch?

A. Client
B. Server
C. Host
D. Transparent
E. Native

Answer: D

Explanation:
You can configure a switch to operate in any one of these VTP modes:
Server-In VTP server mode, you can create, modify, and delete VLANs and specify other configuration parameters (such as VTP version and VTP pruning) for the entire VTP domain. VTP servers advertise their VLAN configuration to other switches in the same VTP domain and synchronize their VLAN configuration with other switches based on advertisements received over trunk links. VTP server is the default mode.
Client-VTP clients behave the same way as VTP servers, but you cannot create, change, or delete VLANs on a VTP client.
Transparent-VTP transparent switches do not participate in VTP. A VTP transparent switch does not advertise its VLAN configuration and does not synchronize its VLAN configuration based on received advertisements. However, in VTP version 2, transparent switches do forward VTP advertisements that they receive out their trunk ports.

QUESTION 104
SIMULATION
The Certkiller network is displayed in the diagram below:

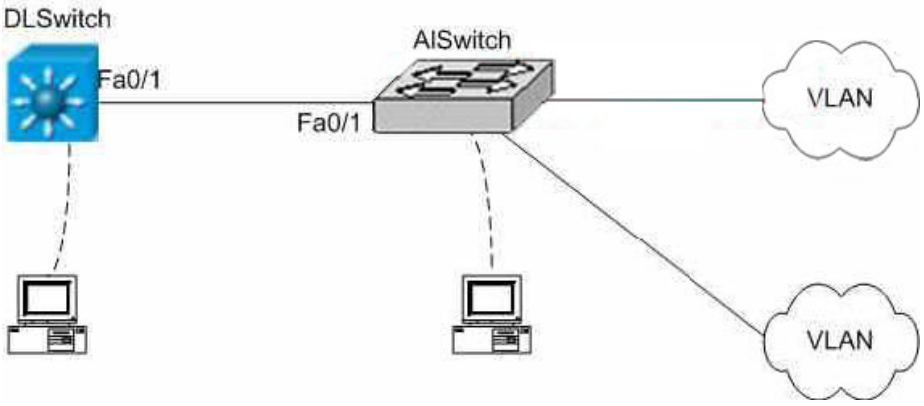

You have just been hired by Certkiller .com to help their main office expand. The main offices have enhanced their wiring closets with some Layer 3 switches. The new distribution layer switch has been installed and a new access layer switch cabled next to it. Your task is to configure the distribution layer and access layer switch with VTP to share VLAN information, then to configure inter-VLAN routing on the distribution layer switch to route traffic between the different VLANs that are configured on the access layer switches.
VTP Domain Distribution
VLAN Ids 20 31
IP Addresses 172.16.71.1/24 172.16.132.1/24
These are your specific tasks:
1. Configure the VTP information with the distribution layer switch as the VTP server
2. Configure the VTP information with the access layer switch as a VTP client
3. Configure VLANs on the distribution layer switch
4. Configure inter-VLAN routing on the distribution layer switch
5. Specific VLAN port assignments will be made as users are added to the access layer switches in the future.
6. All VLANs and VTP configurations are to completed in the global configuration
To configure the switch click on the host icon that is connected to the switch be way of a serial console cable.

Answer:
LAB configuration:
switch#conf t
switch(config)#vtp mode server
switch(config)#vtp domain CISCO
switch(config)#vlan 20
switch(config)#vlan 31
switch(config)#int vlan 20
switch(if-config)#ip add 172.64.20.1 255.255.255.0
switch(if-config)#no shut
switch(if-config)#int vlan 31
switch(if-config)#ip add 192.162.31.1 255.255.255.0
switch(if-config)#no shut
switch(if-config)#exit

switch#ip routing
switch#sh run
switch#copy run start
switch#conf t
switch(config)#vtp mode client
vtp domain CISCO
switch(config)#exit
switch#show run
switch#copy run start
Alternative #1
VTP Domain Distribution
VLAN Ids 20 31
IP Addresses 172.16.16.1/24 172.16.193.1/24
Alternative #12
VTP Domain Distribution
VLAN Ids 30 21
IP Addresses 172.16.203.1/24 172.16.93.1/24

QUESTION 105
The following commands were entered on a Certkiller switch:
Switch(config)# vtp mode transparent
Switch(config)# vtp version 2
What is the result of these commands?

A. VLAN configuration information is saved in RAM only.
B. VLANs cannot be created, modified or deleted via command line interface.
C. VLAN configuration information received via VTP advertisements are forwarded to other switches within the management domain.
D. VLAN configuration information is synchronized with information within VTP advertisements received from other switches in the management domain.

Answer: C

Explanation:
VTPv2 will allow the switch to be in transparent mode which will forward VTP info. The command series above put the switch in VTP transparent mode. This Certkiller switch does not actively participate in VTP, it doesn't advertise its VLAN configuration to other switches, and when other switches advertise their VLAN configuration it doesn't consider that information. It will, however, pass incoming VLAN information that was received to other switches within the VTP domain.

QUESTION 106
The Certkiller network is displayed in the diagram below:

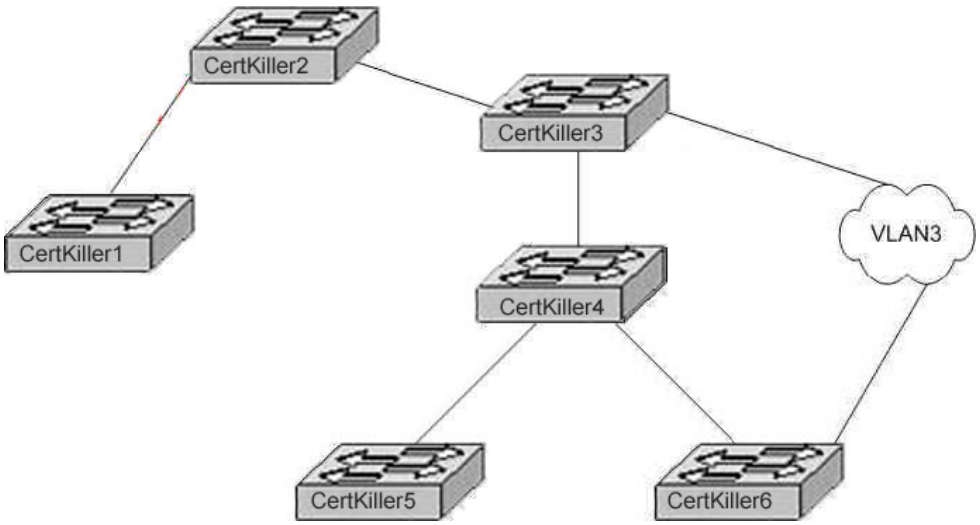

The network in the above exhibit is configured with VLANs 1,2,3,4, & 5 and 802.1 Q. trunking is enabled between all switches. However, access ports for Certkiller 3 and Certkiller 6 are the only access ports for VLAN 3. What could an administrator do to make sure that other switches don't receive unnecessary broadcast packets destined for VLAN 3, while still allowing all the other VLAN packets to cross?

A. Configure VTP pruning.
B. Configure Certkiller 3 and Certkiller 6 as transparent switches.
C. Configure Certkiller 1, Certkiller 2, Certkiller 4 and Certkiller 5 as transparent switches.
D. Nothing is required.
Only Certkiller 3 and Certkiller 6 will receive VLAN3 packets by default.

Answer: A

Explanation:
VTP pruning enhances network bandwidth use by reducing unnecessary flooded traffic, such as broadcast, multicast, unknown, and flooded unicast packets. VTP pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the appropriate network devices. By configuring VTP pruning, traffic will not flow to switches destined for VLANs that they are not attached to.

QUESTION 107
You are configuring VTP on a non IOS switch named CK1 , and you enter the following command:
set vtp pruneeligible
What is this command useful for?

A. For determining management domain name
B. For verifying configuration.
C. For enabling VTP pruning.
D. For selecting VTP version.
E. For verifying configuration set

Answer: C

Explanation:
Use the set vtp command to set the options for VTP.
set vtp [domain domain_name] [mode {client | server | transparent}] [passwd passwd] [pruning{enable | disable}] [v2 {enable | disable}]
The pruning keyword is used to enable or disable VTP pruning for the VTP domain. VTP pruning causes information about each pruning-eligible VLAN to be removed from VTP updates if there are no stations belonging to that VLAN out a particular switch port. Use the set vtp pruneeligible and clear vtp pruneeligible commands to specify which VLANs should or should not be pruned when pruning is enabled for the domain.
Reference:
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_5_5/cmd_refr/set_v.htm

---

QUESTION 108
Switch CK1 is configured as a VTP server. What is true when you enable VTP pruning on a VTP server?

A. It is not possible without a root re-election
B. It enables pruning for the entire management domain.
C. It cannot be done on a VTP server
D. It enables pruning for the individual switch.

Answer: B

Explanation:
Enabling VTP pruning on a VTP server enables pruning for the entire management domain. VTP pruning takes effect several seconds after you enable it. By default, VLANs 2 through 1000 are pruning-eligible. VTP pruning does not prune traffic from VLANs that are pruning-ineligible. VLAN 1 is always pruning-ineligible; traffic from VLAN 1 cannot be pruned.

QUESTION 109
One of the configurable VTP commands is displayed below:
Clear vtp pruneeligible vlan_range
What is the purpose of this above command?

A. Verify the VTP pruning configuration.
B. Make specific VLANs pruning-eligible on the device.
C. Make specific VLANs pruning-ineligible on the device.
D. Enable VTP pruning in the management domain.
E. Verify that the appropriate VLANs are being pruned on trunk ports.

Answer: C

Explanation:
This command makes specific VLANs pruning-ineligible on the device. (By default, VLANs 2-1000 are pruning-eligible.)
Note: VLAN 1 is not pruning eligible.
Reference:
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/rel_5_2/config/vtp.htm

---

QUESTION 110
When setting up multiple VTP domains, what should be considered in order to maintain VLAN database consistency? (Select two)

A. Do not configure any switches as a VTP server
B. Ensure that all switches not authorized to make changes are in client mode
C. Always configure switches using VTP server mode when adding them to the existing network
D. Allow only one VTP server in each domain so that adding and deleting VLANs can be centralized to one location.

Answer: B, D

Explanation:
B: Switches not authorized to make changes should be run as VTP clients. VTP clients receive information from VTP servers and send and receive updates, but they cannot make any changes.
D: You need at least one server in your VTP domain to propagate VLAN information throughout the domain. You are able to use several VTP servers in a domain. However, only allowing one VTP server would help keep the VLAN database consistent.
Incorrect Answers:
A: Switches can very well be used as VTP servers. VTP server mode is the default for all Catalyst switches.
C: It is more prudent to configure switches using VTP client mode. They will not be able to update information in the VLAN domain database.

QUESTION 111
VTP is running on the Certkiller network. In which VTP modes can a full list of all VLANs be maintained? (Select two)

A. VTP Bypass
B. VTP Client
C. VTP Transparent
D. VTP Restore
E. VTP Server

Answer: B, E

Explanation:

VTP-capable devices can be configured to operate in the following three modes:
The VTP Server maintains a full list of all VLANs within the VTP domain. Information is stored in nonvolatile random-access memory (NVRAM). The server can add, delete, and rename VLANs.
The VTP Client also maintains a full list of all VLANs. However, it will not store in NVRAM. The client can not add, delete, or rename VLANs. Any changes made must be received from a VTP server advertisement.
The VTP Transparent mode does not participate in VTP. However, it will pass on a VTP advertisement. VLAN, as defined, is only local to the switch and is stored in NVRAM.

QUESTION 112
You wish to configure VTP on switch CK1 . What do you have to do before you can create a VLAN on a VTP server?

A. The VTP server ID must be cleared
B. The VTP membership list must be refreshed
C. The priority must be cleared
D. The management domain name must be specified

Answer: D

Explanation:
By default, the switch is in VTP server mode and is in the no-management domain state until the switch receives an advertisement for a domain over a trunk link or you configure a management domain. You cannot create or modify VLANs on a VTP server until the management domain name is specified or learned.

QUESTION 113
What must you do if you wish to configure VTP in secure mode within the Certkiller LAN?

A. Assign a management domain password to the VTP Server in the domain.
B. Assign a management domain password to each switch in the domain.
C. Assign a management domain password to the root switch in the domain.
D. None of the above.

Answer: B

Explanation:
If you configure VTP in secure mode, the management domain will not function properly if you do not assign a management domain password to each switch in the domain. All switches must be configured with the password in order for VTP to function properly in a network.

QUESTION 114
If you configure a switch as a VTP server offline, then connect it to a network, what

could happen to the network?

A. Cause a loss of VLAN information
B. Destabilize the spanning tree
C. Revert to simplex mode
D. Revert to duplex mode
E. Ignore the configuration revision numbers created on the other VTP servers
F. Revert to client mode

Answer: A

Explanation:
When connecting a new switch to your network you can accidentally change your current VLAN database if the new switch has a higher VLAN Trunking Protocol (VTP) revision number. If the newly inserted switch has no VLANs configured and the revision number is higher and is configured as a VTP server, it will override the configuration of the other switches within the network, deleting all of the configured VLANs. To avoid this, you must clear the VTP revision number on the new switch. The easiest way is to change the VTP domain name to "something_else" and back to "your_VTP_domain" on the new switch. This sets the VTP revision number to 0 and you can connect the switch to the network without any problem.

QUESTION 115
Which of the following are true if you configure a password for VTP? (Select all that apply)

A. It is carried in all summary-advertisement VTP packets
B. It needs to be the same on all switches in the VTP domain
C. It needs to be configured on all switches in the VTP domain
D. It is translated using an algorithm in a 24 bytes word
E. None of the above

Answer: A, B, C

Explanation:
According to the online documentation provided by Cisco:
If you configure a password for VTP, it needs to be configured on all switches in the VTP domain and it needs to be the same password. The VTP password you configure is translated using an algorithm in a 16 bytes word (MD5 value) carried in all summary-advertisement VTP packets.
Incorrect Answers:
D: The algorithm uses a 16 byte word, not 24 bytes.

QUESTION 116
Is the following statement True or False?
With VTP, if an administrator makes configuration changes centrally on one or

more switches, those changes will be automatically communicated to all the other switches on the network.

A. There is not enough information to determine
B. True
C. False

Answer: B

Explanation:
This statement is true. Before you create virtual LANs (VLANs), you must decide whether to use VTP in your network. With VTP, you can make configuration changes centrally on one or more switches and those changes are automatically communicated to all the other switches in the network. Changes made to VTP servers are propagated to all other switches within the VTP domain.

QUESTION 117
You have to enter a new switch into the existing Certkiller VTP domain without altering the configurations of the systems currently on this domain. Which of the following answer choices describes one of the conditions required to ensure that the new switch will not change the existing VTP domain configuration?

A. The switch must be in client mode.
B. The switch must be in a mode other than the client mode.
C. The VTP domain must not have a password assigned to it.
D. The trunk links must not be configured for ISL
E. None of the above

Answer: A

Explanation:
You can configure a switch to operate in any one of these VTP modes:
Server-In VTP server mode, you can create, modify, and delete VLANs and specify other configuration parameters (such as VTP version and VTP pruning) for the entire VTP domain. VTP servers advertise their VLAN configuration to other switches in the same VTP domain and synchronize their VLAN configuration with other switches based on advertisements received over trunk links. VTP server is the default mode.
Client-VTP clients behave the same way as VTP servers, but you cannot create, change, or delete VLANs on a VTP client.
Transparent-VTP transparent switches do not participate in VTP. A VTP transparent switch does not advertise its VLAN configuration and does not synchronize its VLAN configuration based on received advertisements. However, in VTP version 2, transparent switches do forward VTP advertisements that they receive out their trunk ports.
Only by using the client or transparent modes can you ensure that the other switches within the domain are left unaffected.

QUESTION 118
What method could you use to eliminate unicast and broadcast traffic that is flooded to a VLAN unnecessarily?

A. VTP pruning
B. MLS-SE
C. VTP trunking
D. VTP compression
E. All of the above

Answer: A

Explanation:
VTP ensures that all switches in the VTP domain are aware of all VLANs. There are occasions, however, when VTP can create unnecessary traffic. All unknown unicasts and broadcasts in a VLAN are flooded all over the VLAN. All switches in the network receive all broadcasts, even in situations where few users are connected in that VLAN. VTP pruning is a feature used to eliminate (prune) this unnecessary traffic.

QUESTION 119
You are a network engineer and you've been assigned the task of configuring an Ethernet media trunk between two Cisco switches. Assuming that these two switches have the same modules, software revisions and VLAN configurations, which of the following are true for the trunk to operate? (Choose all that apply)

A. The link must be a point to point for ISL encapsulation.
B. The link must be 100Mbs or slower
C. The link must use the IEEE 802.1Q trunk protocol.
D. The link may use the IEEE 802.1Q trunk protocol.
E. The link can operate at 10, 1000, or 100 Mbs interfaces

Answer: A, D, E

Explanation:
A: This statement is true. ISL trunks must be configured on point to point links; point-to-multipoint configurations are not supported.
D: This statement is also true. 802.1Q can be used, but it does not have to be.
E: Trunks can operate at 10, 100, or 1000 Mbps interfaces
Incorrect Answers:
B: Trunks can operate at 10, 100, or 1000 Mbps interfaces.
C: The industry standard method of trunking is 802.1Q. As an alternative to this, the Cisco proprietary ISL method is also an option for setting up trunks.

QUESTION 120
The following output was seen on a Certkiller switch:

```
CertKiller2# show interfaces gigabitethernet0/1 switchport
Name: Gi0/1
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
Negotiation of Trunking: Off
Access Mode VLAN: 7 (Core)
Trunking Native Mode VLAN: 1 (default)
Voice VLAN: none
Administrative private-vlan host-association: none
Administrative private-vlan mapping: none
Administrative private-vlan trunk native VLAN: none
Administrative private-vlan trunk encapsulation: dot1q
Administrative private-vlan trunk normal VLANs: none
Administrative private-vlan trunk private VLANs: none
Operational private-vlan: none
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
Capture Mode Disabled
Capture VLANs Allowed: ALL
Protected: false
Unknown unicast blocked: disabled
Unknown multicast blocked: disabled

Appliance trust: none
```

Refer to the show interface Gi0/1 switchport command output shown in the exhibit. Based on the information shown above, which two statements are true about this interface? (Select two)

A. This interface is a member of a voice VLAN.
B. This interface is a dot1q trunk passing all configured VLANs.
C. This interface is configured for access mode.
D. This interface is a member of VLAN7.
E. This interface is a member of VLAN1.

Answer: C, D

Explanation: In Exhibit, Operation mode is in static access and Access mode VLAN is 7 so it means this port is operating on access mode as a member of VLAN 7.

QUESTION 121
Two Certkiller switches are connected as shown below:

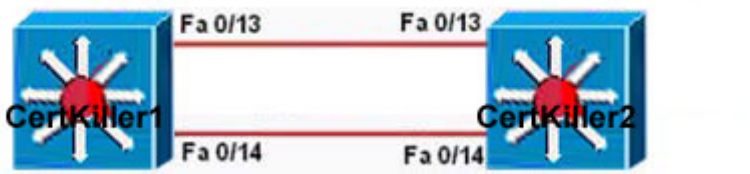

The command "switchport mode access" is issued on interface FastEthernet0/13 on switch Certkiller 1. What will be the result?

A. The command will be rejected by the switch.

B. Interfaces FastEthernet0/13 and FastEthernet0/14 will no longer be bundled.
C. Interfaces FastEthernet0/13 and FastEthernet0/14 will only allow traffic from the native VLAN.
D. Interfaces FastEthernet0/13 and FastEthernet0/14 will continue to pass traffic for VLANs 88,100,360.
E. Dynamic Trunking Protocol will be turned off on interfaces FastEthernet0/13 and FastEthernet0/14.

Answer: B

Explanation:
EtherChannel bundles can consist of up to eight physical ports of the same Ethernet media type and speed. Some configuration restrictions exist to ensure that only similarly configured links are bundled.
Generally, all bundled ports must first belong to the same VLAN. If used as a trunk, bundled ports must all be in trunking mode, have the same native VLAN, and pass the same set of VLANs. Each of the ports should also have the same speed and duplex settings before they are bundled. Bundled ports must also be configured with identical Spanning Tree settings.

QUESTION 122
Certkiller 1 configuration exhibit:

```
CertKiller1#show vtp status
  VTP Version                     : 2
  Configuration Revision          : 0
  Maximum VLANs supported Locally : 1005
  Number of existing VLANs        : 11
  VTP Operating Mode              : Transparent
  VTP Domain Name                 : certkillera
  VTP Pruning Mode                : Enabled
  VTP V2 Mode                     : Enabled
  VTP Traps Generation            : Disabled
  MD5 digest                      : 0x97 0xD6 0xD5 0x0E 0xFA 0x42 0x74
  0xFO
  Configuration last modilled by 0.0.0.0 at 3-1-93 00:57:21
CertKiller1#show vtp password
  VTP Password: certkillerpwd1
```

Certkiller 2 configuration exhibit:

```
CertKiller2#show vtp status
  VTP Version                     : 2
  Configuration Revision          : 1
  Maximum VLANs supported Locally : 1005
  Number of existing VLANs        : 13
  VTP Operating Mode              : Server
  VTP Domain Name                 : certkillera
  VTP Pruning Mode                : Enabled
  VTP V2 Mode                     : Enabled
  VTP Traps Generation            : Disabled
  MD5 digest                      : 0x61 0x57 0xFC 0x33 0xF1 0x77 0x55
  0x57
  Configuration last modified by 0.0.0.0 at 3-2-93 03:00:01
  Local updater ID is 0.0.0.0 (no valid interface found)
CertKiller2#show vtp password
  VTP Password: certkillerpwd2
```

Study the exhibits shown above carefully. Switch Certkiller 1 is not applying VLAN updates from switch Certkiller 2. What are three reasons why this is not occurring? (Select three)

A. The VTP domains are different.
B. The passwords do not match.
C. Switch Certkiller 2 is in server mode.
D. Switch Certkiller 1 is in transparent mode.
E. VTP trap generation is disabled on both switches.
F. The MD5 digests do not match.

Answer: A, B, D

Explanation:
Determine the VTP mode of operation of the switch and include the mode when setting the VTP domain name information on the switch. If you leave the switch in server mode, be sure to verify that the configuration revision number is set to 0 before adding the switch to the VTP domain. It is generally recommended that you have several servers in the domain, with all other switches set to client mode for purposes of controlling VTP information.
It is also highly recommended that you use secure mode in your VTP domain. Assigning a password to the domain will accomplish this. This will prevent unauthorized switches from participating in the VTP domain. From the privileged mode or VLAN configuration mode, use the vtp password password command.

QUESTION 123
The "show vlan" command was issued on a Certkiller device as shown below:

CertKiller1# show vlan

VLAN	Name	Status	Ports
1	default	active	Fa0/1, Fa0/2, Fa0/3, Fa0/4 Fa0/5, Fa0/6,Fa0/7,Fa0/8 Fa0/9, Fa0/10,Fa0/11,Fa0/12 Fa0/15, Fa0/16, Fa0/17, Fa0/18 Fa0/19, Fa0/20, Fa0/21, Fa0/22 Fa0/23, Fa0/24, Gi0/1, Gi0/2
2	VLAN0002	active	
3	VLAN0003	active	
4	VLAN0004	active	
5	VLAN0005	active	
1002	fddi-default	act/unsup	
1003	trcrf-default	act/unsup	
1004	fddinet-default	act/unsup	
1005	trbrf-default	act/unsup	

Study the exhibit carefully. Based upon the output on switch Certkiller 1, what can we conclude about interfaces Fa0/13 and Fa0/14?

A. That interfaces Fa0/13 and Fa0/14 have a domain mismatch with another switch
B. That interfaces Fa0/13 and Fa0/14 have a duplex mismatch with another switch
C. That interfaces Fa0/13 and Fa0/14 are trunk interfaces
D. That interfaces Fa0/13 and Fa0/14 are down
E. That interfaces Fa0/13 and Fa0/14 are in VLAN 1
F. None of the above

Answer: C

Explanation:
trunk-This setting places the port in permanent trunking mode. The corresponding switch port at the other end of the trunk should be similarly configured because negotiation is not allowed. You should also manually configure the encapsulation mode.
show vlan: This commands shows the vlan, ports belonging to VLAN means that port on access mode. It doesn't shows the port on trunk mode.

QUESTION 124
The following output was shown on a Certkiller device:

Study the exhibit carefully. On the basis of the output generated by the show commands, which two statements are true? (Select two)

A. There are no native VLANs configured on the trunk.
B. VLAN 2 will not be encapsulated with an 802.1q header.
C. VLAN 1 will not be encapsulated with an 802.1q header.
D. Because it has not been assigned to any VLAN, interface gigabitethernet 0/1 does not appear in the show vlan output.
E. Because it is configured as a trunk interface, interface gigabitethernet 0/1 does not appear in the show vlan output.
F. All interfaces on the switch have been configured as access ports.

Answer: C, E

Explanation:
The IEEE 802.1Q protocol can also carry VLAN associations over trunk links. However, this frame identification method is standardized, allowing VLAN trunks to exist and operate between equipment from multiple vendors.
In particular, the IEEE 802.1Q standard defines an architecture for VLAN use, services provided with VLANs, and protocols and algorithms used to provide VLAN services.
Like Cisco ISL, IEEE 802.1Q can be used for VLAN identification with Ethernet trunks. Instead of encapsulating each frame with a VLAN ID header and trailer, 802.1Q embeds its tagging information within the Layer 2 frame. This method is referred to as single-tagging or internal tagging.

802.1Q also introduces the concept of a native VLAN on a trunk. Frames belonging to this VLAN are not encapsulated with any tagging information. In the event that an end station is connected to an 802.1Q trunk link, the end station can receive and understand only the native VLAN frames. This provides a simple way to offer full trunk encapsulation to the devices that can understand it, while giving normal access stations some inherent connectivity over the trunk.
show vlan: This commands shows the vlan, ports belonging to VLAN means that port on access mode. It doesn't shows the port on trunk mode.

QUESTION 125
Two Certkiller switches connect multiple VLANs as shown below:

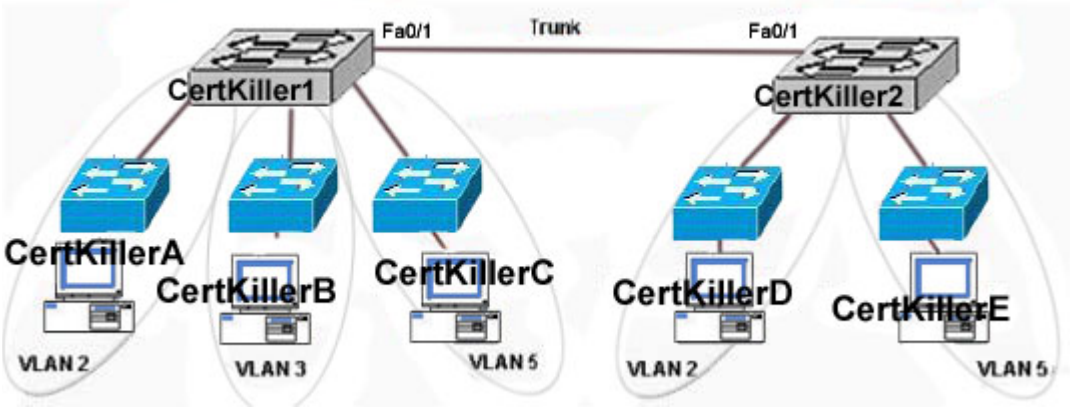

Certkiller 1 configuration exhibit:

```
CertKiller1# show interfaces fastethernet0/1 switchport
Name: Fa0/1
Switchport: Enabled
Administrative Mode: dynamic desirable
Operational Mode: down
Administrative Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: 1,2,3,10-1005
Pruning VLAN: Enabled: 2-4, 10-1001

Protected: false
Unknown unicast blocked: false
Unknown multicast blocked:false
```

Certkiller 2 configuration exhibit:

```
CertKiller2#show interfaces fastethernet0/1 switchport
Name: Fa0/1
Switchport: Enabled
Administrative Mode: dynamic desirable
Operational Mode: down
Administrative Trunking Encapsulation: dot1q
Negotiation of Trunking: On
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: 1,2,3,10-1005
Pruning VLAN: Enabled: 2-4, 10-1001

Protected: false
Unknown unicast blocked: false
Unknown multicast blocked:false
```

Refer to the exhibits and the show interfaces fastethernet0/1 switchport outputs.
Users in VLAN 5 on switch Certkiller 1 complain that they do not have connectivity to the users in VLAN 5 on switch Certkiller 2. What should be done to fix the

problem?

A. Define VLAN 5 in the allowed list for the trunk port on Certkiller 2
B. Configure the same number of VLANs on both switches.
C. Disable pruning for all VLANs in both switches.
D. Define VLAN 5 in the allowed list for the trunk port on Certkiller 1.
E. Create switch virtual interfaces (SVI) on both switches to route the traffic.
F. None of the above.

Answer: D

Explanation:
switchporttrunk allowed vlan, defines which VLANs can be trunked over the link. By default, a switch transports all active VLANs (1 to 4094) over a trunk link. There might be times when the trunk link should not carry all VLANs. For example, broadcasts are forwarded to every switch port on a VLAN-including the trunk link because it, too, is a member of the VLAN.
If the VLAN does not extend past the far end of the trunk link, propagating broadcasts across the trunk makes no sense.

QUESTION 126
In the Certkiller network, VLAN Trunking Protocol (VTP) is running with a domain name of CK1 . VLANs 1, 2, 3, 4, 5, 10, 20 are active on the network. Suddenly the whole network goes down. No traffic is being passed on VLANs 2, 3, 4, 5, 10, 20. However, traffic passes on VLAN 1 and indicates all switches are operational. Right before the network problem occurred; a switch named Certkiller 13 was taken out of the lab and added to the network. What three configuration issues on Certkiller 13 could be causing the network outage? (Select three)

A. Certkiller 13 has a higher VTP configuration revision than the current VTP revision.
B. Certkiller 13 is configured as a VTP server with a different domain name.
C. Certkiller 13 is configured as a VTP server with the domain name CK1 .
D. Certkiller 13 has a lower VTP configuration revision than the current VTP revision.
E. Certkiller 13 is not configured to participate in VTP.
F. Certkiller 13 is configured with only VLAN1.

Answer: A, C, F

Explanation:
VTP Modes:
1. Server
By default, a Catalyst switch is in the VTP server mode and in the "no management domain" state until the switch receives an advertisement for a domain over a trunk link or a VLAN management domain is configured. A switch that has been put in VTP server mode and had a domain name specified can create, modify, and delete VLANs. VTP servers can also specify other configuration parameters such as VTP version and VTP

pruning for the entire VTP domain. VTP information is stored in NVRAM.
2. Client
The VTP client maintains a full list of all VLANs within the VTP domain, but it does not store the information in NVRAM. VTP clients behave the same way as VTP servers, but it is not possible to create, change, or delete VLANs on a VTP client. Any changes made must be received from a VTP server advertisement. Client will make contact with the VTP server in between 5 minutes, it copies the advertisements from that VTP server having highest Revision number. So, before connecting any switch into LAN verify that new switch is in which mode, what is the revision number, is that highest than other switch operated in server mode?
3. Transparent
VTP transparent switches do not participate in VTP. A VTP transparent switch does not advertise its VLAN configuration, and does not synchronize its VLAN configuration based on received advertisements. However, in VTP Version 2, transparent switches do forward VTP advertisements that the switches receive out their trunk ports. VLANs can be configured on a switch in the VTP transparent mode, but the information is local to the switch (VLAN information is not propagated to other switches) and is stored in NVRAM

QUESTION 127
Refer to the following outputs shown on a Certkiller switch:
Exhibit #1

```
CertKiller1# show ip interface vlan 10
Vlan10 is up, line protocol is up
  Internet address is 10.1.1.1/24
  Broadcast address is 255.255.255.255
  Address delermined by setup command
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Multicast reserved groups joined: 224.0.0.10
  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 unreachables 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 CEF Fast 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
  Probe proxy name replies are disabled
  Policy routing is disabled
  Network address translation is disabled
  WCCP Redirect outbound is disabled
  WCCP Redirect inbound is disabled
  WCCP Redirect tesclde is disabled
  BGP Policy Mapping sideabled
```

Exhibit #2

Study the exhibits carefully. What command was issued on this Certkiller switch between Exhibit #1 and Exhibit #2?

A. no router eigrp 1
B. mls qos
C. ip routing
D. router eigrp 1
E. no mls qos
F. no ip routing
G. None of the above

Answer: F

Explanation:
IP routing is Enabled, so to disabled IP routing use:
#no ip routing

QUESTION 128
The Certkiller administrator has issue the "show vlan id 5" command. What will this command display? (Select two)

A. Ports in VLAN 5
B. Utilization
C. VLAN information on port 0/5
D. Filters
E. MTU and type

Answer: A, E

Explanation:
#show vlan id 5 : Shows all ports belonging to VLAN 5 and MTU of ports and type.

QUESTION 129
You're a network administer and you issue the command (show port 3/1) on an Ethernet port. To your surprise you notice a non-zero entry in the 'Giants' column. What could be the cause of this?

A. IEEE 802.1Q
B. IEEE 802.10
C. Misconfigured NIC
D. User configuration
E. All of the above

Answer: A
The 802.1Q standard can create an interesting scenario on the network. Recalling that the maximum size for an Ethernet frame as specified by IEEE 802.3 is 1518 bytes, this means that if a maximum-sized Ethernet frame gets tagged, the frame size will be 1522 bytes, a number that violates the IEEE 802.3 standard. To resolve this issue, the 802.3 committee created a subgroup called 802.3ac to extend the maximum Ethernet size to 1522 bytes.
Note: The show port command is used to display port status and counters. Giants denote the number of received giant frames (frames that exceed the maximum IEEE 802.3 frame size) on the port.
Reference: Trunking between Catalyst 4000, 5000, and 6000 Family Switches Using 802.1q Encapsulation
http://www.cisco.com/warp/public/473/27.html

QUESTION 130
You have a trunk link operating between two switches and you're experiencing problems with frames leaking between the two VLANs. Each switch has identical modules, software revisions and VLAN configuration information. Spanning tree protocol is disabled on all VLANs. What is probably causing this problem? (Select all that apply)?

A. The link is using IEEE 802.1Q protocol
B. The link is using IEEE 802.1E protocol

C. Spanning tree is disabled
D. Not enough information to determine.
E. The native VLAN information is identical at each end of the link.
F. The native VLAN information is different at each end of the link.

Answer: A, F

Explanation:
While internal to a switch, VLAN numbers and identification are carried in a special extended format that allows the forwarding path to maintain VLAN isolation from end to end without any loss of information. Instead, outside of a switch, the tagging rules are dictated by standards such as ISL or 802.1Q.
ISL is a Cisco proprietary technology and is in a sense a compact form of the extended packet header used inside the device: since every packet always gets a tag, there is no risk of identity loss and therefore of security weaknesses.
On the other hand, the IEEE committee that defined 802.1Q decided that because of backward compatibility it was desirable to support the so-called native VLAN, that is to say, a VLAN that is not associated explicitly to any tag on an 802.1Q link. This VLAN is implicitly used for all the untagged traffic received on an 802.1Q capable port.
This capability is desirable because it allows 802.1Q capable ports to talk to old 802.3 ports directly by sending and receiving untagged traffic. However, in all other cases, it may be very detrimental because packets associated with the native VLAN lose their tags, for example, their identity enforcement, as well as their Class of Service (802.1p bits) when transmitted over an 802.1Q link.
For these sole reasons-loss of means of identification and loss of classification-the use of the native VLAN should be avoided.
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps708/products_white_paper09186a008013159f.shtml

---

QUESTION 131
What command could you enter to display the trunking status of a module/port in the switch? (Type in the answer below):

Answer: show trunk

Explanation:
Use the show trunk command to display trunking information for the switch.
show trunk [mod_num[/port_num]] [detail]mod_num (Optional) Number of the module.
/port_num (Optional) Number of the port.
detail (Optional) Keyword to show detailed information about the specified trunk port.

QUESTION 132
You are troubleshooting a Catalyst 5000 trunk in the Certkiller network. What should you do if there's a disagreement about the VLANs configured to use the trunk?

A. Reload the active VLAN configuration
B. Clear the affected port and bring it up again.
C. Explicitly set the trunk for the VLAN to be on.
D. Remove all the VLANs set

Answer: B

Explanation:
In this situation you may want to set or clear the VLANS on both ends. A trunk is a point-to-point link between one or more Ethernet switch interfaces and another networking device such as a router or a switch. Trunks carry the traffic of multiple VLANs over a single link and allow you to extend VLANs across an entire network. Two trunking encapsulations are available on all Ethernet interfaces:
Inter-Switch Link (ISL)-ISL is a Cisco-proprietary trunking encapsulation
802.1Q-802.1Q is an industry-standard trunking encapsulation
When a trunk is first brought up using either of these methods, it may be beneficial to clear the port immediately after.

QUESTION 133
You want to check your Catalyst switch port to see if there's an active link state. Which of the following actions would NOT be useful? (Select all that apply)

A. Switch fan
B. Port's link LED of the Switching Module.
C. Switch RP
D. Switch slot
E. Switch backplane

Answer: A, C, D, E

Explanation:
To find out if there is an active link state on a Catalyst port, check the port's link LED. As an alternative, using various show commands will also aid in troubleshooting port problems. However, checking the fan, RP, slot, or backplane will not help when looking for the status of an individual port's link status.

QUESTION 134
Which of the following commands could you use to clear a switch of its current configuration?

A. the "clear config all" command
B. the "del config all" command
C. the "erase config all" command
D. the "clean config all" command
E. None of the above

Answer: A

Explanation:
In a Cisco switch, the 'clear config all' command will purge the existing configuration on a switch and start you off with a brand new default configuration. All the other commands: del, erase, and clean; aren't valid IOS commands.

QUESTION 135
Which kind of management can be performed from the console port of a Cisco 6500 switch?

A. Physical management of the switch.
B. Logical management of the switch.
C. In-band management of the switch.
D. Out-of-band management of the switch.

Answer: D

Explanation:
When you configure a switch or a router from the console, it is considered 'out of band' because you don't get in there from any of the paths that the network device is a part of. Modems are often attached to the console port, providing for remote out of band management of the device.

QUESTION 136
A VTP domain named Certkiller has six active VLANs. Without notice, all VLANs except VLAN1 fail. Just prior to the failure, Switch Certkiller 2 was added to the network.
Which three issues on Switch Certkiller 2 could be the cause? Select three.

A. Switch Certkiller 2 is configured for only VLAN1.
B. Switch Certkiller 2 is a VTP server in a different domain.
C. Switch Certkiller 2 is a VTP server in the Certkiller domain.
D. Switch Certkiller 2 is not a VTP domain.
E. Switch Certkiller 2 has a lower VTP configuration revision number than the current VTP revision.
F. Switch Certkiller 2 has a higher VTP configuration revision number than the current VTP revision.

Answer: A, C, F

Explanation: A VTP server in a given domain with the highest revision number will overwrite the VTP configuration of all other switch in the same VTP domain. Cisco best practices advises one to configure the correct VTP domain, VTP password, VTP mode, (server, client, transparent), and VTP revision number before adding any new switch to a network. The default VTP mode is server. A network can have more than one VTP

domain. Each VTP domain has it own server(s) that do not influence clients in other VTP domains.

QUESTION 137
You work as a network Technician at Certkiller .com. A new workstation has consistently been unable to obtain an IP address from the DHCP server when the workstation boots. Older workstations function normally, and the new workstation obtains an address when manually forced to renew its address.
What should be configured on the switch to allow the workstation to obtain an IP address at boot?

A. UplinkFast on the switch port connected to the server
B. BackboneFast on the switch port connected to the server
C. PortFast on the switch port connected to the workstation
D. trunking on the switch

Answer: C

Explanation:
Spanning tree PortFast is a Catalyst feature that causes a switch or trunk port to enter the spanning tree Forwarding state immediately, bypassing the Listening and Learning states. IOS-based switches only use PortFast on access ports connected to end stations.
When a device is connected to a port, the port normally enters the spanning tree Listening state. When the Forward Delay timer expires, the port enters the Learning state. When the Forward Delay timer expires a second time, the port is transitioned to the Forwarding or Blocking state. When PortFast is enabled on a switch or trunk port, the port is immediately transitioned to the Forwarding state. As soon as the switch detects the link, the port is transitioned to the Forwarding state (less than 2 seconds after the cable is plugged in).

QUESTION 138
A topology change has occurred in the Certkiller network causing an STP change.
Which two statements concerning STP state changes are true? (Select two)

A. If a forwarding port receives an inferior BPDU, it will transition to listening.
B. Upon bootup, a port transitions from blocking to listening because it assumes itself as root.
C. Upon bootup, a port transitions from listening to forwarding because it assumes itself as root.
D. Upon bootup, a port transitions from blocking to forwarding because it assumes itself as root.
E. If a blocked port receives no BPDUs by the max_age time limit, it will transition to listening.
F. If a forwarding port receives no BPDUs by the max_age time limit, it will transition to listening.

Answer: B, E

Explanation:
As a first step in the STP process, the switches need to elect a single Root Bridge by looking for the bridge with the lowest BID. This process of selecting the bridge with the lowest BID is sometimes called the "root war." A BID is an 8-byte identifier that is composed of two subfields, the Bridge Priority and a MAC address. The process of how to configure a switch to become the Root Bridge will be examined in the following topics. Normally, the default settings should not be allowed to determine the location of the Root Bridge. How did the bridges learn that Cat-A has the lowest BID? This is accomplished through the exchange of BPDUs. As discussed earlier, BPDUs are special frames that bridges use to exchange spanning tree information with each other. By default, BPDUs are sent out every two seconds. BPDUs propagate between bridges, which includes switches and all routers configured for bridging. BPDUs do not carry end-user traffic.

QUESTION 139
Switch CK1 is configured with the RSTP feature. What will occur when a non edge switch port that is configured for Rapid Spanning Tree does not receive a BPDU from its neighbor for three consecutive hello time intervals?

A. RSTP information is automatically aged out.
B. The port sends a TCN to the root bridge.
C. The port moves to listening state.
D. The port becomes a normal spanning tree port.
E. None of the above

Answer: A

Explanation:
The IEEE 802.1D Spanning Tree Protocol was designed to keep a switched or bridged network loop free, with adjustments made to the network topology dynamically. A topology change typically takes 30 seconds, where a port moves from the Blocking state to the Forwarding state after two intervals of the Forward Delay timer. As technology has improved, 30 seconds has become an unbearable length of time to wait for a production network to failover or "heal" itself during a problem.
The IEEE 802.1w standard was developed to take 802.1D's principle concepts and make the resulting convergence much faster. This is also known as the Rapid Spanning Tree Protocol (RSTP). RSTP defines how switches must interact with each other to keep the network topology loop free, in a very efficient manner. Like 802.1D, RSTP's basic functionality can be applied as a single or multiple instances. This can be done as the IEEE 802.1s Multiple Spanning Tree (MST).
BPDUs are sent out every switch port at Hello Time intervals, regardless of whether BPDUs are received from the Root. In this way, any switch anywhere in the network can play an active role in maintaining the topology. Switches can also expect to receive regular BPDUs from their neighbors. When three BPDUs are missed in a row, that

neighbor is presumed to be down, and all information related to the port leading to the neighbor is immediately aged out. This means that a switch can detect a neighbor failure in three Hello intervals (default 6 seconds), versus the Max Age Timer interval (default 20 seconds) for 802.1D.

QUESTION 140
You have been tasked with optimizing the STP operation in the Certkiller switched LAN. Which two statements about the various implementations of STP are true? (Select two)

A. Common Spanning Tree maintains a separate spanning-tree instance for each VLAN configured in the network.
B. Per-VLAN Spanning Tree (PVST) supports 802.1Q trunking.
C. Rapid Spanning Tree Protocol (RSTP) includes features equivalent to Cisco PortFast, UplinkFast, and BackboneFast for faster network reconvergence.
D. Multiple Spanning Tree (MST) assumes one spanning-tree instance for the entire Layer 2 network, regardless of the multiple number of VLANs.
E. The Spanning Tree Protocol (STP) is an evolution of the IEEE 802.1w standard.
F. Per-VLAN Spanning Tree Plus(PVST+) is an enhancement to 802.1Q specification and is supported only on Cisco devices.

Answer: C, F

Explanation:
In 802.1D, BPDUs basically originate from the Root Bridge and are relayed by all switches down through the tree. It is because of this propagation of BPDUs that 802.1D convergence must wait for steady-state conditions before proceeding.
RSTP uses the 802.1D BPDU format for backward-compatibility. However, some previously unused bits in the Message Type field are used. The sending switch port identifies itself by its RSTP role and state. The BPDU version is also set to 2, to distinguish RSTP BPDUs from 802.1D BPDUs. Also, RSTP uses an interactive process so that two neighboring switches can negotiate state changes. Some BPDU bits are used to flag messages during this negotiation.
BPDUs are sent out every switch port at Hello Time intervals, regardless of whether BPDUs are received from the Root. In this way, any switch anywhere in the network can play an active role in maintaining the topology. Switches can also expect to receive regular BPDUs from their neighbors. When three BPDUs are missed in a row, that neighbor is presumed to be down, and all information related to the port leading to the neighbor is immediately aged out. This means that a switch can detect a neighbor failure in three Hello intervals (default 6 seconds), versus the Max Age Timer interval (default 20 seconds) for 802.1D. Because RSTP distinguishes its BPDUs from 802.1D BPDUs, it can coexist with switches still using 802.1D. Each port attempts to operate according to the STP BPDU that is received. For example, when an 802.1D BPDU (version 0) is received on a port, that port begins to operate according to the 802.1D rules. However, each port has a measure that locks the protocol in use for the duration of the migration

delay timer. This keeps the protocol type from flapping or toggling during a protocol migration. After the timer expires, the port is free to change protocols if needed.

QUESTION 141
You have noticed a large number of topology changes in the Certkiller Rapid Spanning Tree Protocol network. Which statement is true about RSTP topology changes?

A. Only nonedge ports moving to the blocking state generate a TC BPDU.
B. Any loss of connectivity generates a TC BPDU.
C. Any change in the state of the port generates a TC BPDU.
D. Only nonedge ports moving to the forwarding state generate a TC BPDU.
E. If either an edge port or a nonedge port moves to a block state, then a TC BPDU is generated.
F. None of the above.

Answer: D

Explanation:
The IEEE 802.1D Spanning Tree Protocol was designed to keep a switched or bridged network loop free, with adjustments made to the network topology dynamically. A topology change typically takes 30 seconds, where a port moves from the Blocking state to the Forwarding state after two intervals of the Forward Delay timer. As technology has improved, 30 seconds has become an unbearable length of time to wait for a production network to failover or "heal" itself during a problem.
Topology Changes and RSTP
Recall that when an 802.1D switch detects a port state change (either up or down), it signals the Root Bridge by sending topology change notification (TCN) BPDUs. The Root Bridge must then signal a topology change by sending out a TCN message that is relayed to all switches in the STP domain. RSTP detects a topology change only when a nonedge port transitions to the Forwarding state. This might seem odd because a link failure is not used as a trigger. RSTP uses all of its rapid convergence mechanisms to prevent bridging loops from forming. Therefore, topology changes are detected only so that bridging tables can be updated and corrected as hosts appear first on a failed port and then on a different functioning port.
When a topology change is detected, a switch must propagate news of the change to other switches in the network so they can correct their bridging tables, too. This process is similar to the convergence and synchronization mechanism-topology change (TC) messages propagate through the network in an ever-expanding wave.

QUESTION 142
You want to make a number of Certkiller switches part of the same Spanning Tree. What must be the same to make multiple switches part of the same Multiple Spanning Tree (MST)?

A. VLAN instance mapping and revision number

B. VLAN instance mapping and member list
C. VLAN instance mapping, revision number, member list, and timers
D. VLAN instance mapping, revision number, and member list
E. None of the above

Answer: A

Explanation:
By default, a switch operates in the Per VLAN Spanning Tree Plus (PVST+) mode using traditional 802.1D STP. Therefore, RSTP cannot be used until a different Spanning Tree mode (MST or RPVST+) is enabled. Remember that RSTP is just the underlying mechanism that a Spanning Tree mode can use to detect topology changes and converge a network into a loop-free topology.
MST is built on the concept of mapping one or more VLANs to a single STP instance. Multiple instances of STP can be used (hence the name MST), with each instance supporting a different group of VLANs.
Each could be tuned to result in a different topology, so that Instance 1 would forward on the left uplink, while Instance 2 would forward on the right uplink. Therefore, VLAN A would be mapped to Instance 1, and VLAN B to Instance 2.
In most networks, a single MST region is sufficient, although you can configure more than one region. Within the region, all switches must run the instance of MST that is defined by the following
attributes:
1. MST configuration name (32 characters)
2. MST configuration revision number (0 to 65535)
3. MST instance-to-VLAN mapping table (4096 entries)

QUESTION 143
An access layer switch in the Certkiller network has just received a BPDU from a neighboring device. What two things will occur when an edge port receives a BPDU? (Select two)

A. The port immediately transitions to the err-disable state.
B. The port becomes a normal STP switch port.
C. The switch generates a Topology Change Notification (TCN) BPDU.
D. The port immediately transitions to the Forwarding state.

Answer: B, C

Explanation:
Recall that when an 802.1D switch detects a port state change (either up or down), it signals the Root Bridge by sending topology change notification (TCN) BPDUs. The Root Bridge must then signal a topology change by sending out a TCN message that is relayed to all switches in the STP domain.
RSTP detects a topology change only when a nonedge port transitions to the Forwarding state. This might seem odd because a link failure is not used as a trigger. RSTP uses all of

its rapid convergence mechanisms to prevent bridging loops from forming. Therefore, topology changes are detected only so that bridging tables can be updated and corrected as hosts appear first on a failed port and then on a different functioning port.

QUESTION 144
Please refer to the Certkiller Switched LAN shown below:

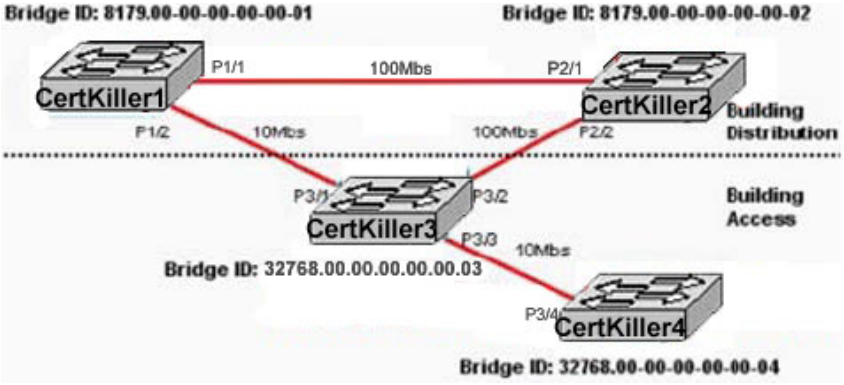

Given the diagram shown above and assuming that STP is enabled on all switch devices, which two statements are true? (Select two.)

A. P2/2 will be elected the nondesignated port.
B. P3/2 will be elected the nondesignated port.
C. P3/1 will be elected the nondesignated port.
D. Certkiller 2 will be elected the root bridge.
E. Certkiller 1 will be elected the root bridge.
F. Certkiller 3 will be elected the root bridge.

Answer: C, E

Explanation:
The Root Bridge in a network using RSTP is elected just as with 802.1D-by the lowest Bridge ID. After all switches agree on the identity of the Root, the following port roles are determined:
1. RootPort-The one switch port on each switch that has the best root path cost to the Root. This is identical to 802.1D. (By definition, the Root Bridge has no Root Ports.)
2. Designated Port-The switch port on a network segment that has the best root path cost to the Root.
3. AlternatePort-A port that has an alternate path to the Root, different than the path the Root Port takes. This path is less desirable than that of the Root Port. (An example of this is an accesslayer switch with two uplink ports; one becomes the Root Port, the other is an Alternate Port.)
4. BackupPort-A port that provides a redundant (but less desirable) connection to a segment where another switch port already connects. If that common segment is lost, the switch might or might not have a path back to the Root.

QUESTION 145
The Certkiller switches have been implemented with the extended ID system feature.

Which two statements are true when the extended system ID feature is enabled?
(Select two)

A. The BID is made up of the bridge priority (four bits), the system ID (12 bits), and a bridge MAC address (48 bits).
B. The BID is made up of the system ID (six bytes) and bridge priority value (two bytes).
C. The BID is made up of the bridge priority value(two bytes) and bridge MAC address (six bytes).
D. The system ID value is a unique MAC address allocated from a pool of MAC addresses assigned to the switch or module.
E. The system ID value is a hex number used to measure the preference of a bridge in the spanning-tree algorithm.
F. The system ID value is the VLAN ID (VID).

Answer: C, F

Explanation:
Each switch has a unique Bridge ID that identifies it to other switches. The Bridge ID is an 8-byte value consisting of the following fields:
1. Bridge Priority (2 bytes)-The priority or weight of a switch in relation to all other switches. The priority field can have a value of 0 to 65,535 and defaults to 32,768 (or 0x8000) on every Catalyst switch.
2. MAC Address (6 bytes)-The MAC address used by a switch can come from the Supervisor module, the backplane, or a pool of 1024 addresses that are assigned to every Supervisor or backplane depending on the switch model. In any event, this address is hardcoded, unique and the user cannot change it.

QUESTION 146
A change in the Certkiller spanning tree network is being propagated to all STP devices. How are STP timers and state transitions affected when a topology change occurs in an STP environment?

A. All ports will temporarily transition to the learning state for a period equal to the max age timer plus the forward delay interval.
B. The default hello time for configuration BDPUs will be reduced for the period of the max age timer.
C. The default aging time for MAC address entries will be reduced for a period of the max age timer plus the forward delay interval.
D. All ports will transition temporarily to the learning state for a period equal to the forward delay interval.
E. None of the above.

Answer: C

Explanation:
1. Hello Time-The time interval between Configuration BPDUs sent by the Root Bridge.

The Hello Time value configured in the Root Bridge switch determines the Hello Time for all nonroot switches because they just relay the Configuration BPDUs as they are received from the root. However, all switches have a locally configured Hello Time that is used to time TCN BPDUs when they are retransmitted. The IEEE 802.1D standard specifies a default Hello Time value of 2 seconds.
2. Forward Delay-The time interval that a switch port spends in both the Listening and Learning states. The default value is 15 seconds.
3. Max (maximum) Age-The time interval that a switch stores a BPDU before discarding it. While executing the STP, each switch port keeps a copy of the "best" BPDU that it has heard. If the BPDU's source loses contact with the switch port, the switch notices that a topology change occurred after the Max Age time elapses and the BPDU is aged out. The default Max Age value is 20 seconds.

QUESTION 147
The Certkiller network administrator is fine tuning the STP parameters on the Catalyst switches. In which states does the Spanning Tree protocol (STP) get affected by the forward delay parameter? (Select two)

A. Forwarding
B. Listening
C. Blocking
D. Disabled
E. Learning

Answer: B, E

Explanation:
The following states utilize information from the forward delay timer:
Listen - The switch listens for a period of time called the fwd delay (forward delay)
Learn - The switch learns for a period of time called the fwd delay (forward delay)
Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 139

QUESTION 148
If two paths to a root switch share the exact same path cost, what information will spanning tree use to determine the root port?

A. The lowest time to receive BPDUs.
B. The lowest Port ID.
C. The lowest sender bridge ID.
D. The highest MAC address on the receiving port.
E. None of the above.

Answer: C

Explanation:
A Root Bridge is chosen based on the results of the BPDU process between the switches.

Initially, every switch considers itself the Root Bridge! When a switch first powers up on the network, it sends out a BPDU with its own BID as the Root BID. When the other switches receive the BPDU, they compare the BID to the one they already have stored as the Root BID. If the new Root BID has a lower value, they replace the saved one. But if the saved Root BID is lower, a BPDU is sent to the new switch with this BID as the Root BID. When the new switch receives the BPDU, it realizes that it is not the Root Bridge and replaces the Root BID in its table with the one it just received. The result is that the switch that has the lowest BID is elected by the other switches as the Root Bridge.
Based on the location of the Root Bridge, the other switches determine which of their ports has the lowest path cost to the Root Bridge. These ports are called Root Ports and each switch (other than the current Root Bridge) must have one.
The switches determine who will have Designated Ports. A Designated Port is the connection used to send and receive packets on a specific segment. By having only one Designated Port per segment, all looping issues are resolved!
Designated Ports are selected based on the lowest path cost to the Root Bridge for a segment. Since the Root Bridge will have a path cost of "0", any ports on it that are connected to segments will become Designated Ports. For the other switches, the path cost is compared for a given segment. If one port is determined to have a lower path cost, then it becomes the Designated Port for that segment. If two or more ports have the same path cost, then the switch with the lowest BID is chosen.

QUESTION 149
What is true of a topology change in an STP environment?

A. The default aging time for MAC address entries will be reduced for a period of the max_age timer plus the forward_delay interval.
B. All ports will transition temporarily to the learning state for a period equal to the forward_delay interval.
C. All ports will temporarily transition to the learning state for a period equal to the max_age timer plus the forward_delay interval.
D. The default hello_timer for configuration BDPUs will be reduced for the period of the max_age timer.

Answer: C

Explanation:
In normal STP operation, a bridge keeps receiving configuration BPDUs from the root bridge on its root port. However, it never sends out a BPDU toward the root bridge. In order to achieve that, a special BPDU called the topology change notification (TCN) BPDU has been introduced. Therefore, when a bridge needs to signal a topology change, it starts to send TCNs on its root port. The designated bridge receives the TCN, acknowledges it, and generates another one for its own root port. The process continues until the TCN hits the root bridge.

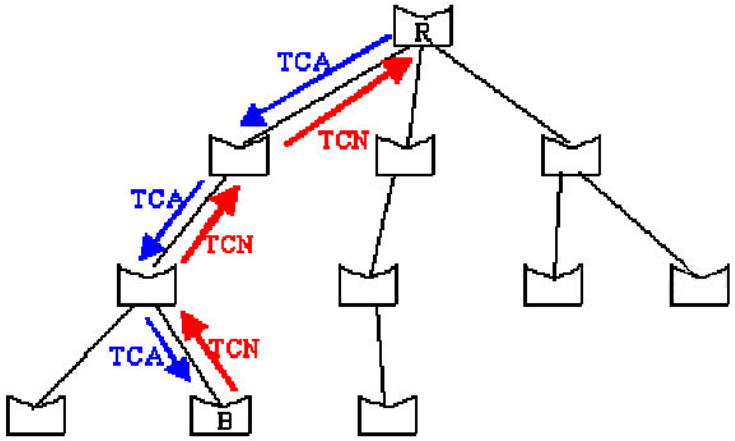

Bridge B notifies a topology change by sending a TCN on its root port. The TCN is acknowledged and forwarded up to the root bridge R.

The TCN is a very simple BPDU that contains absolutely no information that a bridge sends out every hello_time seconds (this is locally configured hello_time, not the hello_time specified in configuration BPDUs). The designated bridge acknowledges the TCN by immediately sending back a normal configuration BPDU with the topology change acknowledgement (TCA) bit set. The bridge that notifies the topology change does not stop sending its TCN until the designated bridge has acknowledged it. Therefore, the designated bridge answers the TCN even though it does not receive configuration BPDU from its root.

Broadcast the Event to the Network

Once the root is aware that there has been a topology change event in the network, it starts to send out its configuration BPDUs with the topology change (TC) bit set. These BPDUs are relayed by every bridge in the network with this bit set. As a result all bridges become aware of the topology change situation and it can reduce its aging time to forward_delay. Bridges receive topology change BPDUs on both forwarding and blocking ports.

The TC bit is set by the root for a period of max_age + forward_delay seconds, which is 20+15=35 seconds by default.

Reference:

Understanding Spanning-Tree Protocol Topology Changes

http://www.cisco.com/warp/public/473/17.html

QUESTION 150

Multiple Certkiller switches are connected together, forming a loop in the network to provide redundancy. Which of the following technologies provides loop avoidance?

A. VTP
B. MLS-RP
C. MLS-SE
D. VTP Pruning
E. STP
F. STP Trunking

G. None of the above

Answer: E

Explanation:
Spanning-Tree Protocol (STP) is a Layer 2 protocol designed to run on bridges and switches. The specification for STP is defined in IEEE 802.1d. The main purpose of STP is to ensure that you do not run into a loop situation when you have redundant paths in your network. STP detects/disables network loops and provides backup links between switches or bridges. It allows the device to interact with other STP compliant devices in your network to ensure that only one path exists between any two stations on the network.
Reference: http://www.zyxel.com/support/supportnote/ves1012/app/stp.htm

---

QUESTION 151
Before a port can participate in the STP process the ports have to change. In which sequence do the STP port states change through?

A. Initial, Learning, Updating, and Active
B. Blocking, Listening, Updating, and Active
C. Initial, Learning, Updating, and Forwarding
D. Blocking, Listening, Learning, and Forwarding

Answer: D

Explanation: The correct order is: blocking state (not participating), listening, learning (prepares to participate), and Forwarding.
Note: STP states:
1. Blocking-The Layer2 LAN port does not participate in frame forwarding
2. Listening-First transitional state after the blocking state when STP determines that the Layer2 LAN port should participate in frame forwarding
3. Learning-The Layer2 LAN port prepares to participate in frame forwarding
4. Forwarding-The Layer2 LAN port forwards frames
5. Disabled-The Layer2 LAN port does not participate in STP and is not forwarding frames.

QUESTION 152
In order for STP to run successfully on the Certkiller network, what standard the bridges and switches have to comply with?

A. 802.1c
B. 802.1e
C. 802.1x
D. 802.1f
E. 802.1d
F. 802.11

Answer: E

Explanation:
According to the online documentation provided by Cisco:
STP runs on bridges and switches that are 802.1d-compliant. There are different flavors of STP, with IEEE 802.1d being the most popular and widely implemented. STP is implemented on bridges and switches in order to prevent loops in the network. Use it in situations where you want redundant links, but not loops. Redundant links are important as backups in case of failover in a network. If your primary fails, the backup links are activated so that users can continue using the network. Without STP on the bridges and switches, such a situation could result in a loop.

QUESTION 153
Switches CK1 and CK2 are exchanging Bridge Protocol Data Unit (BPDU) information. Which of the following can result from a BPDU exchange? (Select all that apply)

A. One switch is elected as the root switch.
B. Ports included in the spanning tree are selected.
C. The shortest distance to the root switch is calculated
D. A designated bridge for each LAN segment is selected.
E. A root port is selected.

Answer: A, B, C, D, E

Explanation:
A BPDU exchange between devices results in the following:
One switch is elected as the root switch.
The shortest distance to the root switch is calculated for each switch based on the path cost.
A designated bridge for each LAN segment is selected. This is the switch closest to the root bridge through which frames are forwarded to the root.
A root port is selected. This is the port providing the best path from the bridge to the root bridge.
Ports included in the spanning tree are selected.

QUESTION 154
What determines the default spanning tree port path cost of STP devices within the Certkiller network?

A. The server speed settings
B. The available bandwidth.
C. The media speed of an interface.
D. The stored IOS settings
E. The interface number

Answer: C

Explanation:
The spanning tree port path cost default value is derived from the media speed of an interface. In the event of a loop, spanning tree considers port cost when selecting an interface to put into the forwarding state. You can assign lower cost values to interfaces that you want spanning tree to select first and higher cost values to interfaces that you want spanning tree to select last. If all interfaces have the same cost value, spanning tree puts the interface with the lowest interface number in the forwarding state and blocks other interfaces. The possible cost range is 1 through 200000000 (the default is media specific).

QUESTION 155
Switch CK1 is a non-root switch in the Certkiller network. By what method does a non-Root switch choose its Root Port?

A. It chooses the port with the lowest cumulative Root Path Cost to the Root Bridge.
B. The port receives an inferior BPDU from a neighboring switch on a shard LAN segment.
C. It chooses the port with the highest cumulative Root Path Cost to the Root Bridge.
D. The port receives a BPDU announcing a higher Root Path Cost from a neighboring switch on a shared LAN segment.
E. None of the above.

Answer: A

Explanation:
The spanning tree Protocol uses the information found in the BPDUs to determine which ports should be forwarding and which should be blocking. If costs are equal, the STP reads through BPDU until it finds a parameter that is not equal. The lower port ID becomes the forwarding port, and the higher port ID is placed in a blocked state. As the BPDU prepares to leave a port, it applies a port cost. The sum of all the port costs is the path cost. Spanning Tree looks first at the path cost to decide which ports should forward and which should block. The port that reports the lowest path cost is chosen to forward.
Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 155

QUESTION 156
What is the default transition time for a switch in the Certkiller switched LAN to move from blocking to forwarding state in the Spanning-Tree protocol?

A. 5 seconds
B. 50 seconds
C. 60 seconds
D. 90 seconds
E. 120 seconds

Answer: B

Explanation:
The default STP timers are shown below:
From blocking to listening 20 seconds
From listening to learning 15 seconds
From learning to forwarding 15 seconds
From blocking to forwarding state 50 seconds
Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 141

QUESTION 157
DRAG DROP
Drag the Spanning Tree Protocol state in the options column on the left to the matching
definition column on the right.

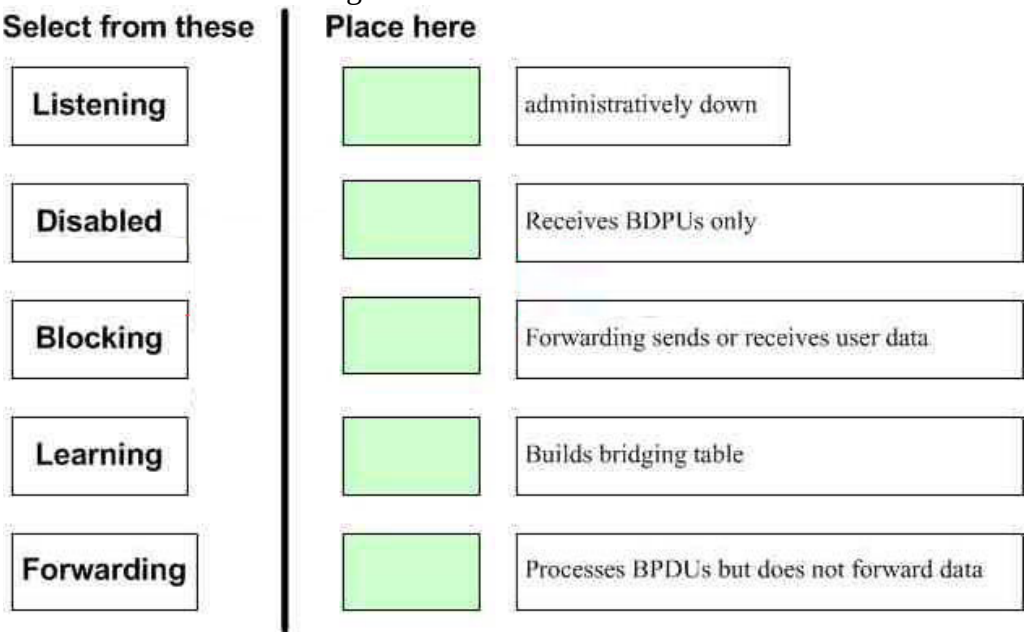

Answer:

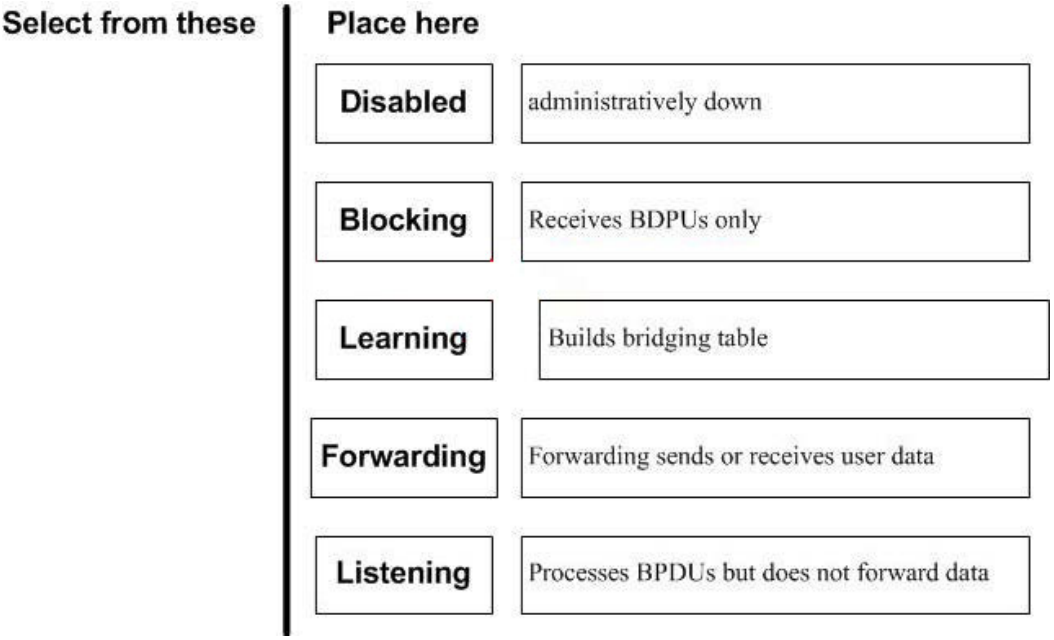

Explanation:
After the bridges have determined which ports are Root Ports, Designated Ports, and non-Designated Ports, STP is ready to create a loop-free topology. To do this, STP configures Root Ports and Designated Ports to forward traffic. STP sets non-Designated Ports to block traffic. Although Forwarding and Blocking are the only two states commonly seen in a stable network, there are actually five STP states.
This list can be viewed hierarchically in that bridge ports start at the Blocking state and work their way up to the Forwarding state. The Disabled state is the administratively shutdown STP state. It is not part of the normal STP port processing. After the switch is initialized, ports start in the Blocking state. The Blocking state is the STP state in which a bridge listens for BPDUs.
A port in the Blocking state does the following:
1. Discards frames received from the attached segment or internally forwarded through switching
2. Receives BPDUs and directs them to the system module
3. Has no address database
4. Does not transmit BPDUs received from the system module
5. Receives and responds to network management messages but does not transmit them
If a bridge thinks it is the Root Bridge immediately after booting or in the absence of BPDUs for a certain period of time, the port transitions into the Listening state. The Listening state is the STP state in which no user data is being passed, but the port is sending and receiving BPDUs in an effort to determine the active topology.
A port in the Listening state does the following:
1. Discards frames received from the attached segment or frames switched from another port
2. Has no address database
3. Receives BPDUs and directs them to the system module
4. Processes BPDUs received from the system module (Processing BPDUs is a separate action from receiving or transmitting BPDUs)

5. Receives and responds to network management messages
It is during the Listening state that the three initial convergence steps take place - elect a Root Bridge, elect Root Ports, and elect Designated Ports. Ports that lose the Designated Port election become non-Designated Ports and drop back to the Blocking state. Ports that remain Designated Ports or Root Ports after 15 seconds - the default Forward Delay STP timer value - progress into the Learning state. The lifetime of the Learning state is also governed by the Forward Delay timer of 15 seconds, the default setting.
The Learning state is the STP state in which the bridge is not passing user data frames but is building the bridging table and gathering information, such as the source VLANs of data frames. As the bridge receives a frame, it places the source MAC address and port into the bridging table. The Learning state reduces the amount of flooding required when data forwarding begins.
A port in the Learning state does the following:
1. Discards frames received from the attached segment
2. Discards frames switched from another port for forwarding
3. Incorporates station location into its address database
4. Receives BPDUs and directs them to the system module
5. Receives, processes, and transmits BPDUs received from the system module
6. Receives and responds to network management messages
If a port is still a Designated Port or Root Port after the Forward Delay timer expires for the Learning state, the port transitions into the Forwarding state. The Forwarding state is the STP state in which data traffic is both sent and received on a port. It is the "last" STP state. At this stage, it finally starts forwarding user data frames.
A port in the Forwarding state does the following:
1. Forwards frames received from the attached segment
2. Forwards frames switched from another port for forwarding
3. Incorporates station location information into its address database
4. Receives BPDUs and directs them to the system module
5. Processes BPDUs received from the system module
6. Receives and responds to network management messages

QUESTION 158
Switch CK1 is participating in the Spanning Tree Protocol (STP). What is true about STP Path Cost on a particular port of CK1 ?

A. It is known only to the local switch where the port resides.
B. It can be modified to help determine Root Bridge selection.
C. Modifying it can cause TCN BPDU to be sent to the Root Bridge.
D. When increased, it can provide higher bandwidth to a connecting port.
E. None of the above

Answer: A

Explanation:
With STP, first a root bridge is elected. Then, the shortest distance to the root bridge is calculated for each switch based on the path cost. This calculation is done locally on each

switch and the path cost for that switch is only used on the local switch.
Incorrect Answers:
B: Adjust the STP port priority, not the port path cost, can be done to influence the election of the root bridge.
C: A bridge considers it a topology change only when one of the following occurs:
1. When a port that was forwarding is going down (blocking for instance).
2. When a port transitions to forwarding and the bridge has a designated port. (This means that the bridge is not standalone.)
TCN BPDUs are only sent to other switches within the network if one of the above happens.
D: Simply adjusting the cost value of the port will not make the port faster or provide for additional bandwidth throughput.

QUESTION 159
Study the exhibit below:

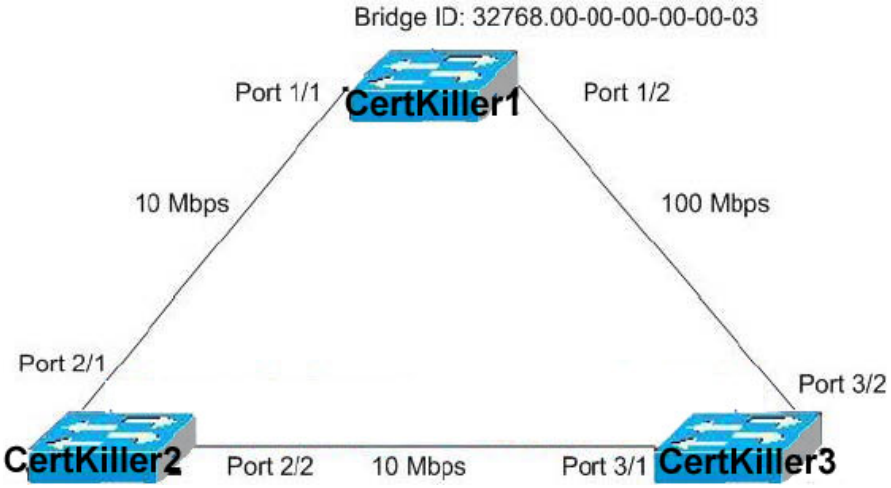

Switch Certkiller 2 Switch Certkiller 3
Bridge ID: 32768.00-00-00-00-00-01 Bridge ID: 32768.00-00-00-00-00-02
Given the network configuration above and assuming that STP is enabled, which port will be elected the non-designated port?

A. Port 1/1
B. Port 1/2
C. Port 2/1
D. Port 2/2
E. Port 3/1
F. Port 3/2

Answer: B:

Explanation:
For each VLAN, the switch with the highest bridge priority (the lowest numerical priority value) is elected as the root bridge. If all switches are configured with the default priority value (32,768), the switch with the lowest MAC address in the VLAN becomes

the root bridge.
The spanning tree root bridge is the logical center of the spanning tree topology in a switched network. All paths that are not required to reach the root bridge from anywhere in the switched network are placed in spanning tree blocking mode.
A spanning tree uses the information provided by BPDUs to elect the root bridge and root port for the switched network, as well as the root port and designated port for each switched segment.
In this example, since the priorities are set to the default, the switch with the lowest MAC address is used as the tie breaker. In this case, Certkiller 2 will become the root switch, which means that port 3/1 and 1/1 will become the root ports and must be in the forwarding state. That leaves the other port on switch Certkiller 2, port 1/2 as the non-designated port since this switch has the highest MAC address.

QUESTION 160
Switch CK1 is powered on for the first time in the Certkiller network. Upon initial bootup, which destination address does a CK1 use to send BPDUs?

A. A well-known STP multicast address.
B. The IP address of its default gateway.
C. The MAC addresses stored in the CAM table.
D. The MAC address of neighbors discovered via CDP
E. None of the above

Answer: A

Explanation:
Bridge protocol data units (BPDUs) are used by the spanning tree algorithm to determine information about the topology of the network BPDUs are used to send configuration messages using multicast frames. When STP devices are first powered on, a well known multicast destination MAC address is used to send the BPDU information.
Incorrect Answers:
B: This would only send the BPDU information to the router. The other switches in the network that are participating in STP need the BPDU information, not the router.
C: When a switch is first powered up, the CAM table will be empty.
D: Since STP is standards based, it does not use any Cisco proprietary protocols such as CDP to perform any of its functions. This will ensure inter-operability with switches from other vendors.

QUESTION 161
Switch CK3 is calculating the root path cost to the Root Bridge, CK1 . What is true regarding the Root Path cost?

A. It is the Path Cost of a particular Root Port.
B. It is the cost sent from the Root Bridge to all non-root bridges.
C. This value is the cumulative cost of all the links leading to the Root Bridge.
D. This value is the cumulative cost of all links sent from the Designated Port of the Root

Bridge.

Answer: C

Explanation:
The first stage in the STP process is the calculation stage. During this stage, each bridge on the network transmits BPDUs that allow the system to work out:
1. The identity of the bridge that is to be the Root Bridge - the central reference point from which the network is configured.
2. The Root Path Costs for each bridge - that is, the cost of the paths from each bridge to the Root Bridge. This value is found by adding up the cost of all of the links to the root bridge.
3. The identity of the port on each bridge that is to be the Root Port - the one that is connected to the Root Bridge using the most efficient path, that is, the one that has the lowest Root Path Cost. Note that the Root Bridge does not have a Root Port.
4. The identity of the bridge that is to be the Designated Bridge of each LAN segment - the one that has the lowest Root Path Cost from that segment. Note that if several bridges have the same Root Path Cost, the one with the lowest Bridge Identifier becomes the Designated Bridge.

QUESTION 162
At which layer of the OSI model does the Spanning Tree Protocol (STP) operate at?

A. Layer 5
B. Layer 4
C. Layer 3
D. Layer 2
E. Layer 1

Answer: D

Explanation:
Spanning-Tree Protocol (STP) is a Layer 2 (L2) protocol designed to run on bridges and switches. The specification for STP is called 802.1d. The main purpose of STP is to ensure that you do not run into a loop situation when you have redundant paths in your network. Loops are deadly to a network.

QUESTION 163
You want to influence the root switch election process within the Certkiller network. When setting up STP in this network, which switch should you configure as the root switch?

A. The most centralized switch
B. The most secure switch
C. The most updated switch
D. The most powerful switch

E. The switch that has the longest uptime.

Answer: A

Explanation:
Cisco recommends using the most centralized switch in the network as the root switch.
According to Cisco:
Before configuring STP, you need to select a switch to be the root of the spanning tree. It does not necessarily have to be the most powerful switch; it should be the most centralized switch on the network. All dataflow across the network will be from the perspective of this switch. It is also important that this switch be the least disturbed switch in the network. The backbone switches are often selected for this function, because they typically do not have end stations connected to them. They are also less likely to be disturbed during moves and changes within the network.

QUESTION 164
Which of the following factors are NOT used to determine the stable active spanning tree topology of the switched Certkiller network?

A. The port identifier associated with each Layer 2 interface
B. The port identifier associated with each Layer 3 interface
C. The spanning tree path cost to the root bridge
D. The unique bridge ID
E. All of the above are used.

Answer: B

Explanation:
The Spanning Tree Protocol does not use layer 3 information to determine the overall topology. Layer 3 interfaces do not participate in STP, since spanning tree is a layer 2 technology.
Incorrect Answers:
A, C, D: The stable active spanning tree topology of a switched network is determined by the following:
-The unique bridge ID (bridge priority and MAC address) associated with each VLAN on each switch
-The spanning tree path cost to the root bridge
-The port identifier (port priority and MAC address) associated with each Layer 2 interface

---

QUESTION 165
A failure has occurred in the Certkiller switched network, causing a loop. What causes bridging loops to occur in a LAN?

A. A failure in the route-switch module
B. A failure in the VLAN tunnel

C. A failure in the VTP trunk
D. A failure in the STA
E. None of the above

Answer: D

Explanation:
The primary function of the spanning-tree algorithm (STA) is to cut loops created by redundant links in bridged networks. The Spanning-Tree Protocol (STP) operates at Layer 2 of the OSI model and, by the means of bridge protocol data units (BPDUs) exchanged between bridges, elects the ports that will eventually forward or block traffic. This protocol can fail in some specific cases and troubleshooting the resulting situation can be very difficult, depending on the design of the network. We can even say that in this particular area, the most important part of the troubleshooting is done before the problem occurs. A failure in the STA generally leads to a bridging loop (not a spanning tree loop as you don't need STP to have a loop). Most customers calling the TAC for spanning tree problems are suspecting a bug, but experience proves that it is seldom the case. Even if the software is at stake, a bridging loop in a STP environment necessarily comes from a port that should block, but that is forwarding traffic.

QUESTION 166
The bridge priority of switch CK1 is being manually configured. In the STP root selection process, what happens to the switch with the lowest priority in the network?

A. It is withdrawn from the election process.
B. It loses the root bridge election process.
C. It wins the root bridge election process.
D. None of the above. The bridge priority is not used to determine the root bridge.

Answer: C

Explanation:
As the BPDU goes out through the network, each switch compares the BPDU it sent out to the one it received from its neighbors. From this comparison, the switches come to an agreement as to who the root switch is. The switch with the lowest priority in the network wins this election process.

QUESTION 167
If a layer 2 interface on switch CK1 uses the Spanning Tree Protocol (STP) which of the following states could it NOT possibly be in at any time?

A. Forwarding
B. Learning
C. Disabled
D. Blocking

E. Listening
F. None of the above

Answer: F

Explanation:
According to Cisco:
Each Layer 2 interface on a switch using spanning tree exists in one of the following five states:
Blocking-The Layer 2 interface does not participate in frame forwarding
Listening-First transitional state after the blocking state when spanning tree determines that the Layer 2 interface should participate in frame forwarding
Learning-The Layer 2 interface prepares to participate in frame forwarding
Forwarding-The Layer 2 interface forwards frames
Disabled-The Layer 2 interface does not participate in spanning tree and is not forwarding frames.

QUESTION 168
When a network engineer designs a switch topology, they assign higher priority values to interfaces that they want spanning tree to select first and lower priority values to interfaces that they want spanning tree to select last. However, if multiple interfaces have equal priority values, spanning tree puts the interface with the _____ interface number in the forwarding state.

A. Neutral
B. Highest
C. Lowest
D. Random
E. First

Answer: C

Explanation:
In the event of a loop, spanning tree considers port priority when selecting an interface to put into the forwarding state. You can assign higher priority values to interfaces that you want spanning tree to select first and lower priority values to interfaces that you want spanning tree to select last. If all interfaces have the same priority value, spanning tree puts the interface with the lowest interface number in the forwarding state and blocks other interfaces.

QUESTION 169
What is the default priority value assigned to a switch when STP is enabled?

A. 1
B. 255
C. 4096

D. 32,768
E. 65,536

Answer: D

Explanation:
Each switch has a unique Bridge ID that identifies it to other switches. The Bridge ID is an 8-byte value consisting of the following fields:
1. Bridge Priority (2 bytes)-The priority or weight of a switch in relation to all other switches. The priority field can have a value of 0 to 65,535 and defaults to 32,768 (or 0x8000) on every Catalyst switch.

QUESTION 170
DRAG DROP
As a Certkiller .com administrator you are required to drag the port states to their correct description.

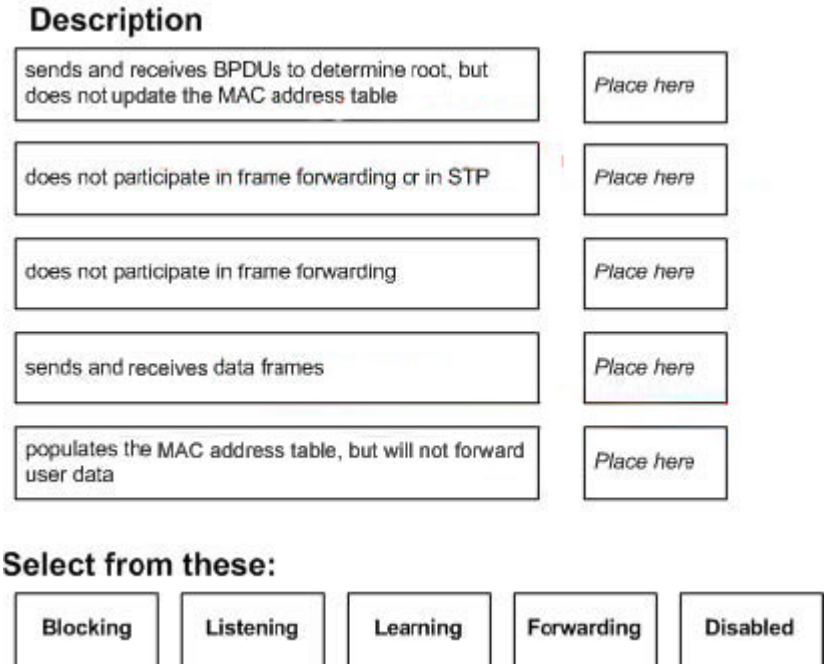

Answer:

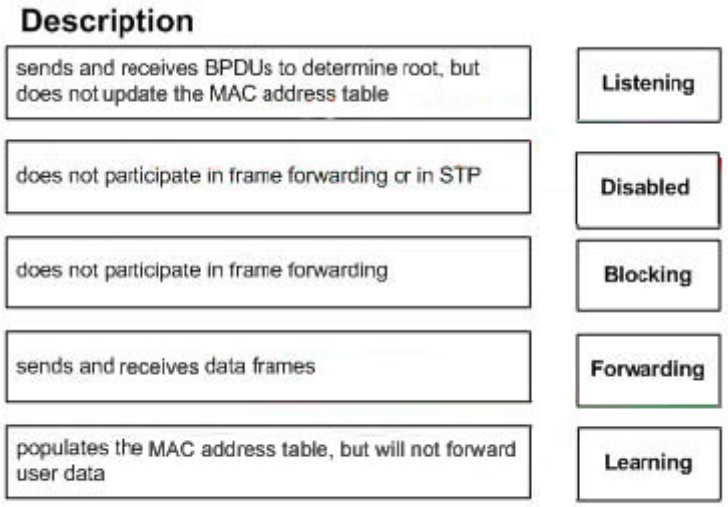

Select from these:

Explanation:

After the bridges have determined which ports are Root Ports, Designated Ports, and non-Designated Ports, STP is ready to create a loop-free topology. To do this, STP configures Root Ports and Designated Ports to forward traffic. STP sets non-Designated Ports to block traffic. Although Forwarding and Blocking are the only two states commonly seen in a stable network, there are actually five STP states.

This list can be viewed hierarchically in that bridge ports start at the Blocking state and work their way up to the Forwarding state. The Disabled state is the administratively shutdown STP state. It is not part of the normal STP port processing. After the switch is initialized, ports start in the Blocking state. The Blocking state is the STP state in which a bridge listens for BPDUs.

A port in the Blocking state does the following:

1. Discards frames received from the attached segment or internally forwarded through switching
2. Receives BPDUs and directs them to the system module
3. Has no address database
4. Does not transmit BPDUs received from the system module
5. Receives and responds to network management messages but does not transmit them

If a bridge thinks it is the Root Bridge immediately after booting or in the absence of BPDUs for a certain period of time, the port transitions into the Listening state. The Listening state is the STP state in which no user data is being passed, but the port is sending and receiving BPDUs in an effort to determine the active topology.

A port in the Listening state does the following:

1. Discards frames received from the attached segment or frames switched from another port
2. Has no address database
3. Receives BPDUs and directs them to the system module
4. Processes BPDUs received from the system module (Processing BPDUs is a separate action from receiving or transmitting BPDUs)
5. Receives and responds to network management messages

It is during the Listening state that the three initial convergence steps take place - elect a

Root Bridge, elect Root Ports, and elect Designated Ports. Ports that lose the Designated Port election become non-Designated Ports and drop back to the Blocking state. Ports that remain Designated Ports or Root Ports after 15 seconds - the default Forward Delay STP timer value - progress into the Learning state. The lifetime of the Learning state is also governed by the Forward Delay timer of 15 seconds, the default setting.
The Learning state is the STP state in which the bridge is not passing user data frames but is building the bridging table and gathering information, such as the source VLANs of data frames. As the bridge receives a frame, it places the source MAC address and port into the bridging table. The Learning state reduces the amount of flooding required when data forwarding begins.
A port in the Learning state does the following:
1. Discards frames received from the attached segment
2. Discards frames switched from another port for forwarding
3. Incorporates station location into its address database
4. Receives BPDUs and directs them to the system module
5. Receives, processes, and transmits BPDUs received from the system module
6. Receives and responds to network management messages
If a port is still a Designated Port or Root Port after the Forward Delay timer expires for the Learning state, the port transitions into the Forwarding state. The Forwarding state is the STP state in which data traffic is both sent and received on a port. It is the "last" STP state. At this stage, it finally starts forwarding user data frames.
A port in the Forwarding state does the following:
1. Forwards frames received from the attached segment
2. Forwards frames switched from another port for forwarding
3. Incorporates station location information into its address database
4. Receives BPDUs and directs them to the system module
5. Processes BPDUs received from the system module
6. Receives and responds to network management messages

QUESTION 171
What is the purpose of MST, according to the IEEE 802.1s standard?

A. It is the spanning-tree implementation used by non-Cisco 802.1Q switches.
B. It runs a separate instance of STP for each VLAN.
C. It allows a VLAN bridge to use multiple spanning trees to prevent Layer 2 loops.
D. It creates a single loop-tree structure that spans the entire Layer 2 network.

Answer: C

Explanation:
IEEE 802.1s MST Overview
MST allows you to build multiple spanning trees over trunks. You can group and associate VLANs to spanning tree instances. Each instance can have a topology independent of other spanning tree instances. This new architecture provides multiple forwarding paths for data traffic and enables load balancing. Network fault tolerance is improved because a failure in one instance (forwarding path) does not affect other

instances (forwarding paths).
In large networks, you can more easily administer the network and use redundant paths by locating different VLAN and spanning tree instance assignments in different parts of the network. A spanningtree instance can exist only on bridges that have compatible VLAN instance assignments. You must configure a set of bridges with the same MST configuration information, which allows them to participate in a specific set of spanning tree instances. Interconnected bridges that have the same MST configuration are referred to as an MST region.
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a008007e

---

QUESTION 172
What will occur when a nonedge switch port that is configured for Rapid Spanning Tree does not receive a BPDU from its neighbor for three consecutive hello time intervals?

A. RSTP information is automatically aged out.
B. The port sends a TCN to the root bridge.
C. The port moves to listening state.
D. The port becomes a normal spanning tree port.

Answer: A

Explanation:
The IEEE 802.1D Spanning Tree Protocol was designed to keep a switched or bridged network loop
free, with adjustments made to the network topology dynamically. A topology change typically
takes 30 seconds, where a port moves from the Blocking state to the Forwarding state after two
intervals of the Forward Delay timer. As technology has improved, 30 seconds has become an
unbearable length of time to wait for a production network to failover or "heal" itself during a
problem.
The IEEE 802.1w standard was developed to take 802.1D's principle concepts and make the
resulting convergence much faster. This is also known as the Rapid Spanning Tree Protocol (RSTP).
RSTP defines how switches must interact with each other to keep the network topology loop free,
in a very efficient manner. Like 802.1D, RSTP's basic functionality can be applied as a single or
multiple instances. This can be done as the IEEE 802.1s Multiple Spanning Tree (MST), covered in

this chapter, and also as the Cisco-proprietary, Rapid Per-VLAN Spanning Tree Protocol (RPVST+).
RSTP operates consistently in each, but replicating RSTP as multiple instances requires different
approaches.
In 802.1D, BPDUs basically originate from the Root Bridge and are relayed by all switches down
through the tree. It is because of this propagation of BPDUs that 802.1D convergence must wait for
steady-state conditions before proceeding.
RSTP uses the 802.1D BPDU format for backward-compatibility. However, some previously
unused bits in the Message Type field are used. The sending switch port identifies itself by its RSTP
role and state. The BPDU version is also set to 2, to distinguish RSTP BPDUs from 802.1D BPDUs.
Also, RSTP uses an interactive process so that two neighboring switches can negotiate state
changes. Some BPDU bits are used to flag messages during this negotiation.
BPDUs are sent out every switch port at Hello Time intervals, regardless of whether BPDUs are
received from the Root. In this way, any switch anywhere in the network can play an active role in
maintaining the topology. Switches can also expect to receive regular BPDUs from their neighbors.
When three BPDUs are missed in a row, that neighbor is presumed to be down, and all information
related to the port leading to the neighbor is immediately aged out. This means that a switch can
detect a neighbor failure in three Hello intervals (default 6 seconds), versus the Max Age Timer
interval (default 20 seconds) for 802.1D.
Because RSTP distinguishes its BPDUs from 802.1D BPDUs, it can coexist with switches still
using 802.1D. Each port attempts to operate according to the STP BPDU that is received. For example,
when an 802.1D BPDU (version 0) is received on a port, that port begins to operate according
to the 802.1D rules. However, each port has a measure that locks the protocol in use for the duration
of the migration delay timer. This keeps the protocol type from flapping or toggling during a protocol
migration. After the timer expires, the port is free to change protocols if needed.

QUESTION 173
The Certkiller network utilizes the Multiple Instance Spanning Tree protocol in its

switched LAN. Which three statements about the MIST protocol (IEEE 802.1S) are true? (Select three)

A. To verify the MST configuration, the show pending command can be used in MST configuration mode.
B. When RSTP and MSTP are configured; UplinkFast and BackboneFast must also be enabled.
C. All switches in the same MST region must have the same VLAN-to-instance mapping, but different configuration revision numbers.
D. All switches in an MST region, except distribution layer switches, should have their priority lowered from the default value 32768.
E. An MST region is a group of MST switches that appear as a single virtual bridge to adjacent CST and MST regions.
F. Enabling MST with the "spanning-tree mode mst" global configuration command also enables RSTP.

Answer: A, E, F

Explanation:
MST is built on the concept of mapping one or more VLANs to a single STP instance. Multiple instances of STP can be used (hence the name MST), with each instance supporting a different group of VLANs.
Each could be tuned to result in a different topology, so that Instance 1 would forward on the left uplink, while Instance 2 would forward on the right uplink. Therefore, VLAN A would be mapped to Instance 1,and VLAN B to Instance 2.
To implement MST in a network, you need to determine the following:
1. The number of STP instances needed to support the desired topologies.
2. Whether to map a set of VLANs to each instance.

QUESTION 174
Certkiller uses MSTP within their switched LAN. What is the main purpose of Multiple Instance Spanning Tree Protocol (MSTP)?

A. To enhance Spanning Tree troubleshooting on multilayer switches
B. To reduce the total number of spanning tree instances necessary for a particular topology
C. To provide faster convergence when topology changes occur in a switched network
D. To provide protection for STP when a link is unidirectional and BPDUs are being sent but not received
E. None of the above

Answer: B

Explanation:
MST is built on the concept of mapping one or more VLANs to a single STP instance. Multiple instances of STP can be used (hence the name MST), with each instance

supporting a different group of VLANs.
Each could be tuned to result in a different topology, so that Instance 1 would forward on the left uplink, while Instance 2 would forward on the right uplink. Therefore, VLAN A would be mapped to Instance 1,and VLAN B to Instance 2.
To implement MST in a network, you need to determine the following:
1. The number of STP instances needed to support the desired topologies.
2. Whether to map a set of VLANs to each instance.

QUESTION 175
Which of the following specifications is a companion to the IEEE 802.1w Rapid Spanning Tree Protocol (RSTP) algorithm, and warrants the use multiple spanning-trees?

A. IEEE 802.1s (MST)
B. IEEE 802.1Q (CST)
C. Cisco PVST+
D. IEEE 802.1d (STP)
E. None of the above

Answer: A

Explanation:
MST uses the modified RSTP version called the Multiple Spanning Tree Protocol (MSTP). MST extends the IEEE 802.1w rapid spanning tree (RST) algorithm to multiple spanning trees. This extension provides both rapid convergence and load balancing in a VLAN environment. MST converges faster than PVST+. MST is backward compatible with 802.1D STP, 802.1w (rapid spanning tree protocol [RSTP]), and the Cisco PVST+ architecture.
MST allows you to build multiple spanning trees over trunks. You can group and associate VLANs to spanning tree instances. Each instance can have a topology independent of other spanning tree instances. This new architecture provides multiple forwarding paths for data traffic and enables load balancing. Network fault tolerance is improved because a failure in one instance (forwarding path) does not affect other instances (forwarding paths).
In large networks, you can more easily administer the network and use redundant paths by locating different VLAN and spanning tree instance assignments in different parts of the network. A spanningtree instance can exist only on bridges that have compatible VLAN instance assignments. You must configure a set of bridges with the same MST configuration information, which allows them to participate in a specific set of spanning tree instances. Interconnected bridges that have the same MST configuration are referred to as an MST region.
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a008007e

---

QUESTION 176
Which of the following specification will allow you to: associate VLAN groups to STP instances so you can provide multiple forwarding paths for data traffic and enable load balancing?

A. IEEE 802.1d (STP)
B. IEEE 802.1s (MST)
C. IEEE 802.1Q (CST)
D. IEEE 802.1w (RSTP)

Answer: B

Explanation:
IEEE 802.1s MST Overview
MST extends the IEEE 802.1w rapid spanning tree (RST) algorithm to multiple spanning trees. This extension provides both rapid convergence and load balancing in a VLAN environment. MST converges faster than PVST+. MST is backward compatible with 802.1D STP, 802.1w (rapid spanning tree protocol [RSTP]), and the Cisco PVST+ architecture.
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a008007e

---

QUESTION 177
Which three items are configured in MST configuration submode? (Select three)

A. Region name
B. Configuration revision number
C. VLAN instance map
D. IST STP BPDU hello timer
E. CST instance map
F. PVST+ instance map

Answer: A, B, C

Explanation:
spanning-treemst configuration:
Use the spanning-tree mst configuration command to enter the MST configuration submode. Use the no form of this command to return to the default MST configuration.
Defaults:
The default value for the MST configuration is the default value for all its parameters:
1. No VLANs are mapped to any MST instance (all VLANs are mapped to the CIST instance).
2. The region name is an empty string.
3. The revision number is 0.
Usage Guidelines:

The MST configuration consists of three main parameters:
1. Instance VLAN mapping (see the instance command)
2. Region name (see the name command)
3. Configuration revision number (see the revision command)

QUESTION 178
By default, all VLANs will belong to which MST instance when using Multiple STP?

A. MST00
B. MST01
C. the last MST instance configured
D. none

Answer: A

Explanation:
Recall that the whole idea behind MST is the capability to map multiple VLANs to a smaller number of STP instances. Inside a region, the actual MST instances (MSTIs) exist alongside the IST. Cisco supports a maximum of 16 MSTIs in each region. IST always exists as MSTI number 0, leaving MSTI 1 through 15 available for use. By default all VLANs are belonged to MST00 instance.

QUESTION 179
Which MST configuration statement is correct?

A. MST configurations can be propagated to other switches using VTP.
B. After MST is configured on a Switch, PVST+ operations will also be enabled by default.
C. MST configurations must be manually configured on each switch within the MST region.
D. MST configurations only need to be manually configured on the Root Bridge.
E. MST configurations are entered using the VLAN Database mode on Cisco Catalyst switches.

Answer: C

Explanation:
MST configuration must be manually be configured on each switch within the MST region.

QUESTION 180
While logged into a Certkiller switch you issue the following command:
Certkiller Switch(config-mst)# instance 10 vlan 11-12
What does this command accomplish?

A. It enables a PVST+ instance of 10 for vlan 11 and vlan 12

B. It enables vlan 11 and vlan 12 to be part of the MST region 10
C. It maps vlan 11 and vlan 12 to the MST instance of 10.
D. It creates an Internal Spanning Tree (IST) instance of 10 for vlan 11 and vlan 12
E. It create a Common Spanning Tree (CST) instance of 10 for vlan 11 and vlan 12
F. It starts two instances of MST, one instance for vlan 11 and another instance for vlan 12.

Answer: C

Explanation:
MST extends the IEEE 802.1w rapid spanning tree (RST) algorithm to multiple spanning trees. This extension provides both rapid convergence and load balancing in a VLAN environment. MST converges faster than Per VLAN Spanning Tree Plus (PVST+) and is backward compatible with 802.1D STP, 802.1w (Rapid Spanning Tree Protocol [RSTP]), and the Cisco PVST+ architecture.
MST allows you to build multiple spanning trees over trunks. You can group and associate VLANs to spanning tree instances. Each instance can have a topology independent of other spanning tree instances. This architecture provides multiple forwarding paths for data traffic and enables load balancing. Network fault tolerance is improved because a failure in one instance (forwarding path) does not affect other instances.
Map the VLANs to an MST instance.
If you do not specify the vlan keyword, you can use the no keyword to unmap all the VLANs that were mapped to an MST instance.
If you specify the vlan keyword, you can use the no keyword to unmap a specified VLAN from an MST instance.
Switch(config-mst)# instance instance_number vlan vlan_range
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps663/products_configuration_guide_chapter09186a00800dd

---

QUESTION 181
Exhibit:

```
CertKiller1(config)# spanning-tree mst configuration
CertKiller1(config-mst)# instance 1 vlan 10-20
CertKiller1(config-mst)# name test
CertKiller1(config-mst)# revision 1
CertKiller1(config-mst)# exit
```

```
CertKiller2(config)# spanning-tree mst configuration
CertKiller2(config-mst)# instance 1 vlan 10-20
CertKiller2(config-mst)# name test
CertKiller2(config-mst)# revision 2
CertKiller2(config-mst)# exit
```

Refer to the show spanning-tree mst configuration output shown in the exhibit.
What should be changed in the configuration of the switch Certkiller 2 in order for it to participate in the same MST region?

A. Switch Certkiller 2 must be configured with a different VLAN range.
B. Switch Certkiller 2 must be configured with a different MST name.
C. Switch Certkiller 2 must be configured with the revision number of 1.
D. Switch Certkiller 2 must be configured with the revision number of 2.

Answer: C

Explanation:
MST is built on the concept of mapping one or more VLANs to a single STP instance. Multiple
instances of STP can be used (hence the name MST), with each instance supporting a different
group of VLANs.
In most networks, a single MST region is sufficient, although you can configure more than one
region. Within the region, all switches must run the instance of MST that is defined by the following
attributes:
MST configuration name (32 characters)
MST configuration revision number (0 to 65535)
MST instance-to-VLAN mapping table (4096 entries)
Example of configuration of MST
Switch(config)# spanning-tree mode mst
Switch(config)# spanning-tree mst configuration
Switch(config-mst)# name name
Switch(config-mst)# revision version
The configuration revision number gives you a means to track changes to the
MST region configuration. Each time you make changes to the configuration,
you should increase the number by one. Remember that the region configuration
(including the revision number) must match on all switches in the
region. Therefore, you also need to update the revision numbers on the other
switches to match.
Switch(config-mst)# instance instance-id vlan vlan-list
The instance-id (0 to 15) carries topology information for the VLANs listed
in vlan-list. The list can contain one or more VLANs separated by commas.
You can also add a range of VLANs to the list by separating numbers with a
hyphen. VLAN numbers can range from 1 to 4094. (Remember that by
default, all VLANs are mapped to instance 0, the IST.)
Switch(config-mst)# show pending
region configuration:
Switch(config-mst)# exit
So belong the routers in same MST region, MST attributes should be same, in Certkiller 2 router revision number is not same so to make belong the Certkiller 2 router on same MST region, revision number should be 1.

QUESTION 182
The network administrator maps VLAN 10 through 20 to MST instance 2. How will this information be propagated to all appropriate switches?

A. Information will be carried in the RSTP BPDUs.

B. It will be propagated in VTP updates.
C. Information stored in the Forwarding Information Base and the switch will reply on query.
D. Multiple Spanning Tree must be manually configured on the appropriate switches.

Answer: D

Explanation:
Recall that the whole idea behind MST is the capability to map multiple VLANs to a smaller number
of STP instances. Inside a region, the actual MST instances (MSTIs) exist alongside the IST. Cisco
supports a maximum of 16 MSTIs in each region. IST always exists as MSTI number 0, leaving
MSTI 1 through 15 available for use. MST must be manually configured on the all switch belongs to same MST region.

QUESTION 183
The following was configured on a Certkiller switch:
Certkiller Switch(config) # spanning-tree portfast bpdufilter default
What is the effect of this configuration command?

A. If BPDUs are received by a port configured for PortFast, they are ignored and none are sent.
B. The command will enable BPDU filtering on all ports regardless of whether they are configured for BPDU filtering at the interface level.
C. If BPDUs are received by a port configured for Portfast, the port will transition to forwarding state.
D. If BPDUs are received by a port configured for PortFast, then PortFast is disabled and the BPDUs are processed normally.
E. None of the above.

Answer: D

Explanation:
Spanning tree PortFast is a Catalyst feature that causes a switch or trunk port to enter the spanning tree Forwarding state immediately, bypassing the Listening and Learning states. IOS-based switches only use PortFast on access ports connected to end stations.
When a device is connected to a port, the port normally enters the spanning tree Listening state. When the Forward Delay timer expires, the port enters the Learning state. When the Forward Delay timer expires a second time, the port is transitioned to the Forwarding or Blocking state. When PortFast is enabled on a switch or trunk port, the port is immediately transitioned to the Forwarding state. As soon as the switch detects the link, the port is transitioned to the Forwarding state (less than 2 seconds after the cable is plugged in).

QUESTION 184
You need to verify the operation of the BPDU enhancements that were added to the Certkiller switches. Which two statements are true about BPDU port-guard and BPDU filtering? (Select two)

A. When globally enabled, BPDU port-guard and BPDU filtering apply only to PortFast enabled ports.
B. When a BPDU is received on a BPDU filtering enabled port, the interface goes into the STP blocking state.
C. BPDU port-guard can be enabled globally, whereas BPDU filtering must be enabled on a per-interface basis.
D. When a BPDU is received on a BPDU filtering enabled port, the interface goes into the err-disabled state.
E. When a BPDU is received on a BPDU port-guard enabled port, the interface goes into the err-disabled state.
F. When globally enabled, BPDU port-guard and BPDU filtering apply only to trunking-enabled ports.

Answer: A, E

Explanation:
STP configures a meshed topology into a loop-free, tree-like topology. When the link on a bridge port goes up, there is STP calculation done on that port. The result of the calculation will be the transition of the port into forwarding or Blocking state, depending on the position of the port in the network, and the STP parameters. This calculation and transition period usually takes about 30-50 seconds. At this time, no user data is passing via the port. Some user applications may timeout during this period.
To allow immediate transition of the port into Forwarding state, the STP PortFast feature is enabled. PortFast transitions the port into STP forwarding mode immediately upon linkup. The port still participates in STP in the event that the port is to be a part of a loop, if so it will eventually transition into STP blocking mode.
As long as the port is participating in STP, there is a possibility that some device attached to that port and also running STP with lower bridge priority than that of the current root bridge, will assume the root bridge function and affect active STP topology, thus rendering the network suboptimal. Permanent STP recalculation caused by the temporary introduction and subsequent removal of STP devices with low (zero) bridge priority represents a simple form of Denial of Service (DoS) attack on the network.
The STP PortFast BPDU guard enhancement is designed to allow network designers to enforce the STP domain borders and keep the active topology predictable. The devices behind the ports with STP PortFast enabled are not allowed to influence the STP topology. This is achieved by disabling the port with PortFast configured upon reception of BPDU. The port is transitioned into Errdisable state, and a message is printed on the console. The following is an example of the message printed out as a result of BPDU guard operation:

Configuring BPDU Guard

	Command	Purpose
Step 1	`configure terminal`	Enter global configuration mode.
Step 2	`spanning-tree portfast bpduguard default`	Globally enable BPDU guard. By default, BPDU guard is disabled.
Step 3	`interface` *interface-id*	Enter interface configuration mode, and specify the interface connected to an end station.
Step 4	`spanning-tree portfast`	Enable the Port Fast feature.
Step 5	`end`	Return to privileged EXEC mode.
Step 6	`show running-config`	Verify your entries.
Step 7	`copy running-config startup-config`	(Optional) Save your entries in the configuration file.

QUESTION 185
The following output was seen on a Certkiller switch:

```
CertKillerSwitch # show spanning-tree inconsistentports

Name                 Interface              Inconsistency
-------------------- ---------------------- ------------------
VLAN0001             FastEthernet0/1        Root Inconsistent
VLAN0001             FastEthernet0/2        Root Inconsistent

Number of incinsistent ports (Segments) In the System : 2
```

On the basis of the output of the "show spanning-tree inconsistentports" command shown above, which statement about interfaces FastEthernet 0/1 and FastEthernet 0/2 is true?

A. They have been configured with the spanning-tree guard loop command.
B. They have been configured with the spanning-tree guard root command.
C. They have been configured with the spanning-tree bpdufilter enable command.
D. They have been configured with the spanning-tree bpduguard disable command.
E. They have been configured with the spanning-tree bpduguard enable command.
F. They have been configured with the spanning-tree bpdufilter disable command.
G. None of the above.

Answer: B

Explanation:
The traditional STP does not provide any means for the network administrator to securely enforce the topology of the switched Layer 2 network. This may become especially important in networks with shared administrative control. For example, one switched network controlled by different administrative entities or companies.
Forwarding topology of the switched network is calculated, based among other parameters, on the root bridge position. Although any switch can be Root Bridge in the network, it is better to place the root bridge manually, (somewhere in the core layer) so the forwarding topology will be optimal. The standard STP does not allow the administrator to enforce the position of the root bridge. If a bridge is introduced into the

network with lower bridge priority, it will take the role of the root bridge.
The root guard ensures that the port on which it is enabled is the designated port (normally, root bridge ports are all designated, unless two or more ports of the root bridge are connected together). If the bridge receives superior STP Bridge Port Data Units (BPDUs) on a root guard enabled port, this port will be moved to a root-inconsistent STP state (effectively equal to listening state), and no traffic will be forwarded across this port. The position of the root bridge will be enforced.

QUESTION 186
Study the exhibit shown below carefully.

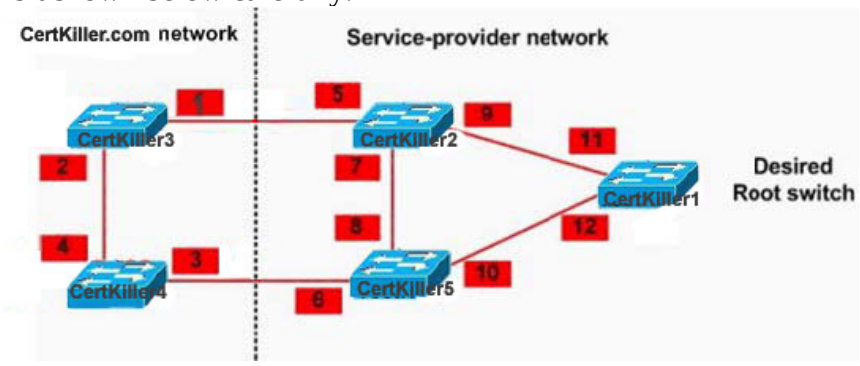

In this scenario the service provider wants to ensure that switch S1 is the root switch for its own network and the network of Certkiller .com. On which interfaces should root guard be configured to ensure that this happens?

A. Interfaces 11 and 12
B. Interfaces 1, 2, 3, and 4
C. Interfaces 5 and 6
D. Interfaces 5, 6, 7, and 8
E. Interfaces 1 and 2
F. Interfaces 1, 3, 5, and 6

Answer: C

Explanation:
The traditional STP does not provide any means for the network administrator to securely enforce the topology of the switched Layer 2 network. This may become especially important in networks with shared administrative control. For example, one switched network controlled by different administrative entities or companies.
Forwarding topology of the switched network is calculated, based among other parameters, on the root bridge position. Although any switch can be Root Bridge in the network, it is better to place the root bridge manually, (somewhere in the core layer) so the forwarding topology will be optimal. The standard STP does not allow the administrator to enforce the position of the root bridge. If a bridge is introduced into the network with lower bridge priority, it will take the role of the root bridge.
The root guard ensures that the port on which it is enabled is the designated port (normally, root bridge ports are all designated, unless two or more ports of the root bridge are connected together). If the bridge receives superior STP Bridge Port Data

Units (BPDUs) on a root guard enabled port, this port will be moved to a root-inconsistent STP state (effectively equal to listening state), and no traffic will be forwarded across this port. The position of the root bridge will be enforced.
Configuring Root Guard

	Command	Purpose
Step 1	`configure terminal`	Enter global configuration mode.
Step 2	`interface` *interface-id*	Enter interface configuration mode, and specify the interface to configure.
Step 3	`spanning-tree guard root`	Enable root guard on the interface. By default, root guard is disabled on all interfaces.
Step 4	`end`	Return to privileged EXEC mode.
Step 5	`show running-config`	Verify your entries.
Step 6	`copy running-config startup-config`	(Optional) Save your entries in the configuration file.

QUESTION 187
You want to optimize the STP timers configuration on the Certkiller switched LAN. Which of the below statements about STP timers are true? (Select three)

A. A switch is not concerned about its local configuration of the STP timers values. It will only consider the value of the STP timers contained in the BPDU it is receiving.
B. The root bridge passes the timer information in BPDUs to all routers in the Layer 3 configuration.
C. On a switched network with a small network diameter, the STP hello timer can be tuned to a lower value to decrease the load on the switch CPU.
D. If any STP timer value (hello, forward delay, max age) needs to be changed, it should at least be changed on the root bridge and backup root bridge.
E. To successfully exchange BPDUs between two switches, their STP timers value (hello, forward delay, max age) must be the same.
F. STP timers values (hello, forward delay, max age) are included in each BPDU.

Answer: A, D, F

Explanation:
STP operation is controlled by three timers. The Hello Time is the amount of time between the sending of Configuration BPDUs. The 802.1D standard specifies a default value of 2 seconds. This value controls Configuration BPDUs as the Root Bridge generates them. Other bridges propagate BPDUs from the Root Bridge as they are received.
If BPDUs stop arriving for the time interval ranging from 2 to 20 seconds because of a network disturbance, or if the Root Bridges stop sending periodic BPDUs during this time, the timer will expire. 2 to 20 seconds is the range between the expected receipt of a BPDU and the expiration of the Max Age time. If the outage lasts for more than 20 seconds, the default Max Age time, the bridge invalidates the saved BPDUs and begins looking for a new Root Port.

Forward Delay is the amount of time the bridge spends in the Listening and Learning states. This is a single value that controls both states. The default value of 15 seconds was originally derived assuming a maximum network size of seven bridge hops, a maximum of three lost BPDUs, and a Hello Time of 2 seconds. The Forward Delay timer also controls the bridge table age-out period after a change in the active topology.
Max Age is the STP timer that controls how long a bridge stores a BPDU before discarding it. Max Age is only an issue when the link failure is not on a directly connected link. When a failure occurs on a directly connected link, the switch knows there will not be any BPDUs coming in on that link, so Max Age is not considered in transitioning the port to Forwarding mode. Recall that each port saves a copy of the best BPDU it has seen. As long as the bridge receives a continuous stream of BPDUs every 2 seconds, the receiving bridge maintains a continuous copy of the BPDU values. However, if the device sending this best BPDU fails, a mechanism must exist to allow other bridges to take over.

QUESTION 188
Examine the Certkiller network diagram below:

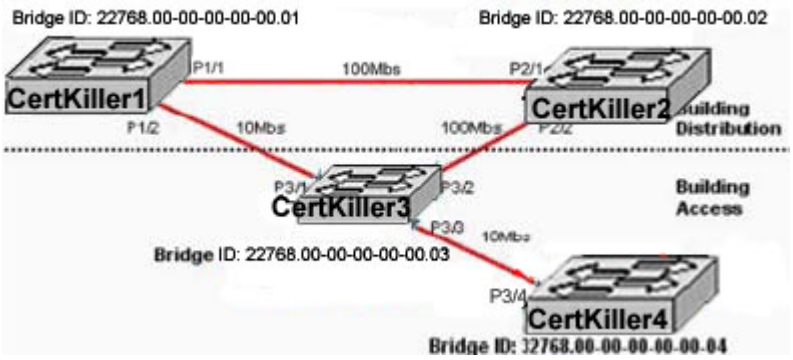

The Certkiller network administrator has recently installed the switched network using 3550s and would like to control the selection of the root bridge. Which switch should the administrator configure as the root bridge and which configuration command must the administrator enter to accomplish this?

A. Certkiller 1(config)# set spanning-tree priority 4096
B. Certkiller 2(config)# spanning-tree vlan 1 priority 4096
C. Certkiller 3(config)# spanning-tree vlan 1 priority 4096
D. Certkiller 3(config)# set spanning-tree priority 4096
E. Certkiller 2(config)# set spanning-tree priority 4096
F. Certkiller 1(config)# spanning-tree vlan 1 priority 4096
G. None of the above

Answer: B

Explanation:
To configure a Catalyst switch to become the Root Bridge, use one of the following methods:
1. Directly modify the Bridge Priority value so that a switch can be given a lower-than-default Bridge ID value to win a Root Bridge election:

Switch (config)# spanning-tree vlan vlan-id priority bridge-priority
The bridge-priority value defaults to 32,768, but you can also assign a value of 0 to 65,535. Remember that Catalyst switches run one instance of STP for each VLAN (PVST+), so the VLAN ID must always be given. You should designate an appropriate Root Bridge for each VLAN.
1. Let the switch become the Root by automatically choosing a Bridge Priority value:
Switch(config)# spanning-tree vlan vlan-id root {primary | secondary}
[diameter diameter]

QUESTION 189
The Certkiller switched LAN is displayed below:

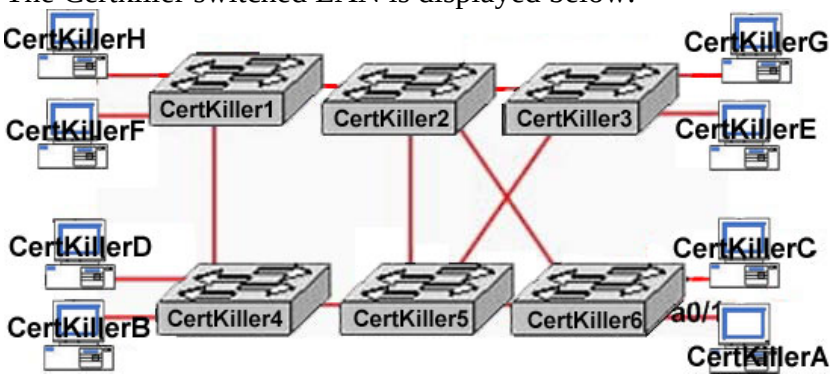

Study the exhibit carefully. The command "spanning-tree bpdufilter enable" is configured on interface Fa0/1 on switch Certkiller 6. The link between switch Certkiller 5 and Certkiller 6 fails. Will Host Certkiller A be able to reach Host Certkiller H?

A. Yes. Traffic will pass from switch Certkiller 6 to Certkiller 2 to Certkiller 1.
B. No. Traffic will pass from switch Certkiller 6 to Certkiller 2 and dead-end at Certkiller 2.
C. Fifty percent of the traffic will successfully reach Host Certkiller H, and fifty percent will dead-end at switch Certkiller 3 because of a partial spanning-tree loop.
D. No. Traffic will loop back and forth between switches Certkiller 2 and Certkiller 3.
E. No. Traffic will loop back and forth between switch Certkiller 6 and Host Certkiller A.

Answer: A

Explanation:
The BPDU filtering feature can be globally enabled on the switch or can be enabled per interface, but the feature operates with some differences.
At the global level, you can enable BPDU filtering on Port Fast-enabled ports by using the spanning-tree portfast bpdufilter default global configuration command. This command prevents ports that are in a Port Fast-operational state from sending or receiving BPDUs. The ports still send a few BPDUs at link-up before the switch begins to filter outbound BPDUs. You should globally enable BPDU filtering on a switch so that hosts connected to these ports do not receive BPDUs. If a BPDU is received on a Port Fast-enabled port, the port loses its Port Fast-operational status, and BPDU filtering is disabled.
At the interface level, you can enable BPDU filtering on any port by using the

spanning-tree bpdufilter enable interface configuration command without also enabling the Port Fast feature. This command prevents the port from sending or receiving BPDUs.

QUESTION 190
You are a network troubleshooter and you've been called into the Certkiller to manually put a switch port back into service after it was put into the error disabled state upon receipt of Spanning Tree messages. Which of the following STP features puts a switch port into an error-disabled state when it receives Spanning Tree data messages?

A. BDPU Filtering
B. Root Guard
C. BDPU Guard
D. Port Fast
E. Loop Guard
F. None of the above

Answer: C

Explanation:
Understanding BPDU Guard
The BPDU guard feature can be globally enabled on the switch or can be enabled per interface, but the feature operates with some differences.
At the global level, you can enable BPDU guard on Port Fast-enabled ports by using the spanning-tree portfast bpduguard default global configuration command. Spanning tree shuts down ports that are in a Port Fast-operational state. In a valid configuration, Port Fast-enabled ports do not receive BPDUs. Receiving a BPDU on a Port Fast-enabled port signals an invalid configuration, such as the connection of an unauthorized device, and the BPDU guard feature puts the port in the error-disabled state.
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps5206/products_configuration_guide_chapter09186a008017

QUESTION 191
Exhibit

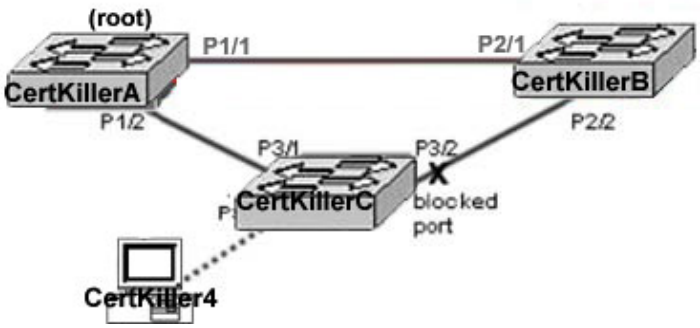

You work as a technician at Certkiller .com. Study the exhibit carefully. Spanning tree is enabled on all devices. Currently either Switch Certkiller B or Certkiller C can

serve as the root should switch Certkiller A fail. A client recently connected to Device Certkiller 4, a PC running Windows XP SP2 and switching application software, to Switch Certkiller C port P3/3. You must configure Root Guard to ensure that the Certkiller 4 PC does not assume the role of the root. All other parameters must stay the same. On which interface(s) must Root Guard be enabled?

A. P1/2
B. P2/2
C. P3/3
D. P1/1 and P1/2
E. P1/2 and P2/2
F. P1/2, P2/2 and P3/3

Answer: C

Explanation:
The Layer 2 network of a service provider (SP) can include many connections to switches that are not owned by the SP. In such a topology, the spanning tree can reconfigure itself and select a customer switch as the root switch. You can avoid this situation by enabling root guard on SP switch interfaces that connect to switches in your customer's network. If spanning-tree calculations cause an interface in the customer network to be selected as the root port, root guard then places the interface in the root-inconsistent (blocked) state to prevent the customer's switch from becoming the root switch or being in the path to the root.
If a switch outside the SP network becomes the root switch, the interface is blocked (root-inconsistent state), and spanning tree selects a new root switch. The customer's switch does not become the root switch and is not in the path to the root.
If the switch is operating in multiple spanning-tree (MST) mode, root guard forces the port to be a designated port. If a boundary port is blocked in an internal spanning-tree (IST) instance because of root guard, the port also is blocked in all MST instances. A boundary port is a port that connects to a LAN, the designated switch of which is either an 802.1D switch or a switch with a different MST region configuration.
Root guard enabled on an interface applies to all the VLANs to which the interface belongs. VLANs can be grouped and mapped to an MST instance.
If your switch is running PVST+, rapid PVST+, or MSTP, you can enable this feature by using the spanning-tree guard root interface configuration command.

QUESTION 192
Which action could have caused the following output to appear on a switch?
%SPANTREE-2-ROOTGUARDBLOCK: Port 1/1 tried to become non-designated in VLAN 77.
Moved to root-consistent state

A. The switch is configured with Loop Guard and stops receiving BPDUs.
B. The switch is configured with Port Fast and starts receiving BPDUs.
C. The switch is configured with Root Guard and starts receiving superios BPDUs.

D. The switch is configured with BackboneFast and starts receiving inferior BPDUs.

Answer: C

Explanation:
The root guard feature was developed as a means to control where candidate Root Bridges can be
connected and found on a network. Basically, a switch learns the cureent Root Bridge's Bridge ID.
If another switch advertises a superior BPDU, or one with a better Bridge ID, on a port where root
guard is enabled, the local switch will not allow the new switch to become the Root. As long as the
superior BPDUs are being received on the port, the port will be kept in the root-inconsistent STP
state. No data can be sent or received in that state, but the switch can listen to BPDUs received on
the port.
The following message is printed once root guard blocks a port.
%SPANTREE-2-ROOTGUARDBLOCK: Port 1/1 tried to become non-designated in VLAN 77. Moved to root-inconsistent state

---

QUESTION 193
While logged in to switch Certkiller 1 you see the following output:
%SPANTREE-2-ROOTGUARDBLOCK: Port 1/1 tried to become non-designated in VLAN 77.
Moved to root-inconsistent state.
Which action could have caused the following output to appear on a switch?

A. The switch is configured with Loop Guard and stops receiving BPDUs.
B. The switch is configured with PortFast and starts receiving BPDUs
C. The switch is configured with Loop Guard and stops receiving superior BPDUs
D. The switch is configured with Loop Guard and starts receiving inferior BPDUs

Answer: C

Explanation:
The loop guard is intended to provide additional protection against L2 forwarding loops (STP loops). An STP loop is created when an STP blocking port in a redundant topology erroneously transitions to forwarding state. This usually happens because one of the ports of a physically redundant topology (not necessarily the STP blocking port) stopped receiving STP BPDUs. In its operation, STP relies on continuous reception or transmission of BPDUs, depending on the port role (designated port transmits, non-designated port receives BPDUs).
When one of the ports in a physically redundant topology stops receiving BPDUs, the STP conceives the topology as loop free. Eventually, the blocking port from the alternate

or backup port becomes designated, and moves to forwarding state, thus creating a loop. With the loop guard, an additional check is made. If BPDUs are not received any more on a non-designated port and the loop guard is enabled, that port will be moved into the STP loop-inconsistent blocking state instead of moving to the listening / learning / forwarding state. Without the loop guard, the port would assume the designated port role. The port would move to STP forwarding state, and thus create a loop.
When the loop guard blocks an inconsistent port, the following message is logged.
SPANTREE-2-LOOPGUARDBLOCK: No BPDUs were received on port 3/2 in vlan 3. Moved to loop-inconsistent state.
Once the BPDU is received on a port in a loop-inconsistent STP state, the port will transition into another STP state. According to the received BPDU, this means that the recovery is automatic, and no intervention is necessary. After the recovery, the following message is logged.
SPANTREE-2-LOOPGUARDUNBLOCK: port 3/2 restored in vlan 3.
Reference:
http://www.cisco.com/en/US/tech/ CK3 89/ CK6 21/technologies_tech_note09186a0080094640.shtml#feature

---

QUESTION 194
Switch CK1 has been configured with the root guard feature. What statement is true if the spanning tree enhancement Root Guard is enabled?

A. If BPDUs are not received on a non-designated port, the port is moved into the STP loop-inconsistent blocked state
B. IF BPDUs are received on a PortFast enabled port, the port is disabled.
C. If superior BPDUs are received on a designated port, the interface is placed into the root-inconsistent blocked state.
D. If inferior BPDUs are received on a root port, all blocked ports become alternate paths to the root bride.

Answer: C

Explanation:
Root guard is configured on a per-port basis, and does not allow the port to become a STP root port. This means that the port is always STP-designated. If there is a better BPDU received on this port, root guard will put the port into root-inconsistent STP state, rather than taking the BPDU into account and electing a new STP root. Root guard needs to be enabled on all ports where the root bridge should not appear. In a way one can configure a perimeter around part of network where STP root is allowed to be located.
Reference:
http://www.cisco.com/en/US/tech/ CK3 89/ CK6 21/technologies_tech_note09186a00800ae96b.shtml

---

QUESTION 195
Study the Certkiller network below carefully. Initially, LinkA is connected and forwarding traffic. A new LinkB is then attached between Switch Certkiller 1 and Hub Certkiller 1. Which two statements are true about the possible result of attaching the second link? (Select two)

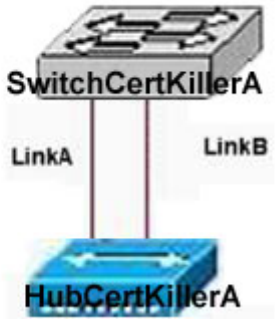

A. The switch port attached to LinkA will immediately transition to the blocking state.
B. A heavy traffic load could cause BPDU transmissions to be blocked and leave a switching loop.
C. One of the two switch ports attached to the hub will go into blocking mode when a BPDU is received.
D. The switch port attached to LinkB will not transistion to up.
E. Both switch ports attached to the hub will transition to the blocking state.

Answer: B, C

Explanation:
Load sharing divides the bandwidth supplied by parallel trunks connecting switches. To avoid loops, STP normally blocks all but one parallel link between switches. Using load sharing, divide the traffic between the links according to which VLAN the traffic belongs.
Configure load sharing on trunk ports by using STP port priorities or STP path costs. For load sharing using STP port priorities, both load-sharing links must be connected to the same switch. For load sharing using STP path costs, each load-sharing link can be connected to the same switch or to two different switches.
Load Sharing Using STP Port Priorities
When two ports on the same switch form a loop, the STP port priority setting determines which port is enabled and which port is in a blocking state. The priorities on a parallel trunk port can be set so that the port carries all the traffic for a given VLAN. The trunk port with the higher priority (lower values) for a VLAN is forwarding traffic for that VLAN. The trunk port with the lower priority (higher values) for the same VLAN remains in a Blocking state for that VLAN. One trunk port sends or receives all traffic for the VLAN.

QUESTION 196
You need to troubleshoot an issue on the Certkiller switched LAN. When you issue a command "show port 3/1" on a Certkiller switch, you observe the Giants column has a non-zero entry. What could cause this?

A. IEEE 802.10
B. Misconfigured NIC
C. User configuration

D. IEEE 802.1Q
E. None of the above

Answer: D

Explanation:
IEEE 802.1Q Protocol
The IEEE 802.1Q protocol can also carry VLAN associations over trunk links. However, this frame identification method is standardized, allowing VLAN trunks to exist and operate between equipment from multiple vendors. In particular, the IEEE 802.1Q standard defines an architecture for VLAN use, services provided with VLANs, and protocols and algorithms used to provide VLAN services.
Like Cisco ISL, IEEE 802.1Q can be used for VLAN identification with Ethernet trunks. Instead of encapsulating each frame with a VLAN ID header and trailer, 802.1Q embeds its tagging information within the Layer 2 frame. This method is referred to as single-tagging or internal tagging.
802.1Q also introduces the concept of a native VLAN on a trunk. Frames belonging to this VLAN are not encapsulated with any tagging information. In the event that an end station is connected to an 802.1Q trunk link, the end station can receive and understand only the native VLAN frames. This provides a simple way to offer full trunk encapsulation to the devices that can understand it, while giving normal access stations some inherent connectivity over the trunk.

QUESTION 197
Certkiller 1 configuration exhibit:

```
CertKiller1 # show spanning-tree interface FastEthernet 0/1 detail

    Port ip FastEthernet0/2) 01 VLAN0001 is forwarding
    Port path cost 19, Port priority 128, Port Identifier 128.1.
    Designated root has priority 32769, address 000a.4107.7400
    Designated bridge has priority 32769, address 000a.4107.7400
    Designated port id is 128.1, designated path cost 0
    Timers: message age 1, forward delay 0, hold 0
    Number of transitions to forwarding state: 1
    BPDU: sent 237, received 1
```

Certkiller 2 configuration exhibit:

```
CertKiller2 # show spanning-tree interface FastEthernet 0/1 detail

    Port ip FastEthernet0/2) 01 VLAN0001 is forwarding
    Port path cost 19, Port priority 128, Port Identifier 128.2.
    Designated root has priority 32769, address 000a.4107.7400
    Designated bridge has priority 32769, address 000a.4107.7400
    Designated port id is 128.1, designated path cost 0
    Timers: message age 1, forward delay 0, hold 0
    Number of transitions to forwarding state: 0
    BPDU: sent 1, received 242
```

Certkiller 3 configuration exhibit:

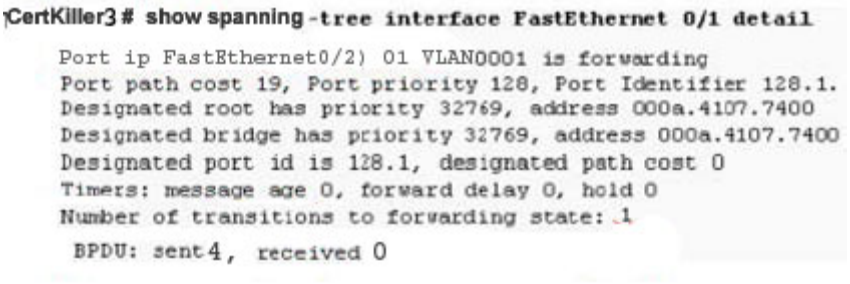

Study the exhibits carefully. Based on the information shown above, which statement is true?

A. The port on switch Certkiller 3 is forwarding and receiving BPDUs correctly.
B. The port on switch Certkiller 1 is forwarding and sending BPDUs correctly.
C. The port on switch Certkiller 1 is blocking and sending BPDUs correctly.
D. The port on switch Certkiller 2 is blocking and sending BPDUs correctly.
E. The port on switch Certkiller 2 is forwarding and receiving BPDUs correctly.
F. The port on switch Certkiller 3 is forwarding, sending, and receiving BPDUs correctly.
G. None of the above.

Answer: B

Explanation:
STP States
To participate in STP, each port of a switch must progress through several states. A port begins its life in a Disabled state, moving through several passive states and, finally, into an active state if allowed to forward traffic. The STP port states are as follows:
1. Disabled-Ports that are administratively shut down by the network administrator, or by the system due to a fault condition, are in the Disabled state. This state is special and is not part of the normal STP progression for a port.
2. Blocking-After a port initializes, it begins in the Blocking state so that no bridging loops can form. In the Blocking state, a port cannot receive or transmit data and cannot add MAC addresses to its address table. Instead, a port is allowed to receive only BPDUs so that the switch can hear from other neighboring switches. In addition, ports that are put into standby mode to remove a bridging loop enter the Blocking state.
3. Listening-The port will be moved from Blocking to Listening if the switch thinks that the port can be selected as a Root Port or Designated Port. In other words, the port is on its way to begin forwarding traffic. In the Listening state, the port still cannot send or receive data frames. However, the port is allowed to receive and send BPDUs so that it can actively participate in the Spanning Tree topology process. Here, the port is finally allowed to become a Root Port or Designated Port because the switch can advertise the port by sending BPDUs to other switches. Should the port lose its Root Port or Designated Port status, it returns to the Blocking state.
4. Learning-After a period of time called the Forward Delay in the Listening state, the port is allowed to move into the Learning state. The port still sends and receives BPDUs as before. In addition, the switch can now learn new MAC addresses to add to its address table. This gives the port an extra period of silent participation and allows the switch to

assemble at least some address table information.
5. Forwarding-After another Forward Delay period of time in the Learning state, the port is allowed to move into the Forwarding state. The port can now send and receive data frames, collect MAC addresses in its address table, and send and receive BPDUs. The port is now a fullyfunctioning switch port within the Spanning Tree topology.

QUESTION 198
The Certkiller switched LAN is shown below:

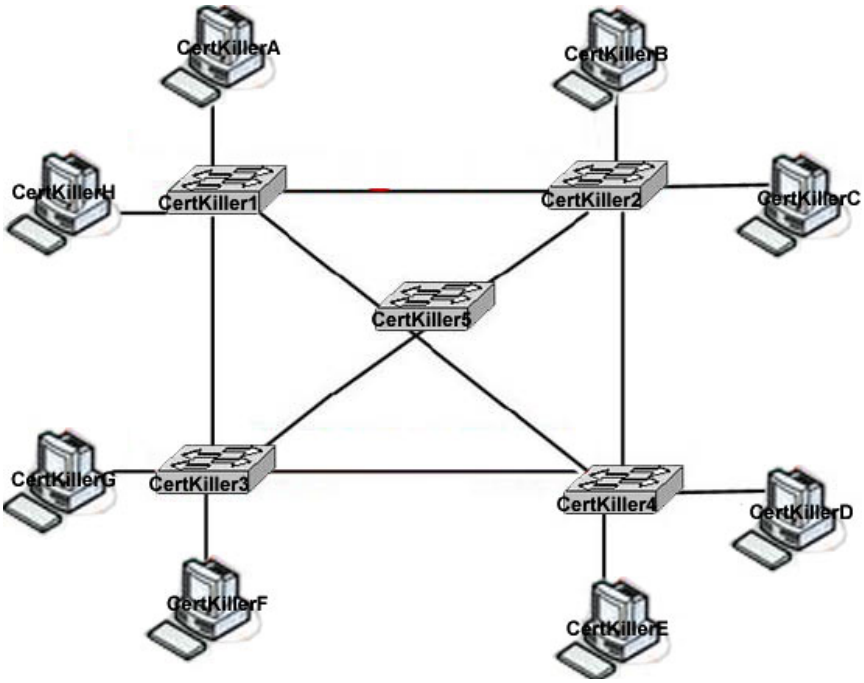

Study the exhibit above carefully. Switch Certkiller 5 is configured as the root switch for VLAN 10 but not for VLAN 20. If the STP configuration is correct, what will be true about Switch Certkiller 5?

A. All ports in VLAN 10 will be in forwarding mode and all ports in VLAN 20 will be in standby mode.
B. All ports will be in forwarding mode.
C. All ports in VLAN 10 will be in forwarding mode and all ports in VLAN 20 will be in blocking mode.
D. All ports in VLAN 10 will be in forwarding mode.
E. None of the above.

Answer: D

Explanation:
STP States
To participate in STP, each port of a switch must progress through several states. A port begins its life in a Disabled state, moving through several passive states and, finally, into an active state if allowed to forward traffic. The STP port states are as follows:
1. Disabled-Ports that are administratively shut down by the network administrator, or by

the system due to a fault condition, are in the Disabled state. This state is special and is not part of the normal STP progression for a port.
2. Blocking-After a port initializes, it begins in the Blocking state so that no bridging loops can form. In the Blocking state, a port cannot receive or transmit data and cannot add MAC addresses to its address table. Instead, a port is allowed to receive only BPDUs so that the switch can hear from other neighboring switches. In addition, ports that are put into standby mode to remove a bridging loop enter the Blocking state.
3. Listening-The port will be moved from Blocking to Listening if the switch thinks that the port can be selected as a Root Port or Designated Port. In other words, the port is on its way to begin forwarding traffic. In the Listening state, the port still cannot send or receive data frames. However, the port is allowed to receive and send BPDUs so that it can actively participate in the Spanning Tree topology process. Here, the port is finally allowed to become a Root Port or Designated Port because the switch can advertise the port by sending BPDUs to other switches. Should the port lose its Root Port or Designated Port status, it returns to the Blocking state.
4. Learning-After a period of time called the Forward Delay in the Listening state, the port is allowed to move into the Learning state. The port still sends and receives BPDUs as before. In addition, the switch can now learn new MAC addresses to add to its address table. This gives the port an extra period of silent participation and allows the switch to assemble at least some address table information.
5. Forwarding-After another Forward Delay period of time in the Learning state, the port is allowed to move into the Forwarding state. The port can now send and receive data frames, collect MAC addresses in its address table, and send and receive BPDUs. The port is now a fullyfunctioning switch port within the Spanning Tree topology.

QUESTION 199
The following output was shown on switch Certkiller 1:

```
CertKiller1#show spanning-tree vlan 200

VLAN0200
 Spanning tree enabled protocaol eee
 Root ID     Priority    32968
             Address     000c.ce29.ef00
             Cost        19
             Port        2 (FastEthernet0/2)
             Hello Time  10 sec  Max Age 20 sec  Forward Delay 30 sec

 Bridge ID   Priority    32968 (priority 32768 sys-id-ext 200)
             Address     000c.ce2a.4180
             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
Fa0/3                  Altn BLK 19        128.3    P2p
```

Based on the "show spanning-tree vlan 200" output shown in the exhibit, which two statements about the STP process for VLAN 200 are true? (Select two)

A. This switch is the root bridge for VLAN 200.
B. The maximum length of time that the BPDU information will be saved is 30 seconds.
C. BPDUs will be sent out every 10 seconds.
D. The time spent in the listening state will be 30 seconds.

E. BPDUs will be sent out every two seconds.
F. The time spent in the learning state will be 15 seconds.

Answer: C, D

Explanation:
STP operation is controlled by three timers. The Hello Time is the amount of time between the sending of Configuration BPDUs. The 802.1D standard specifies a default value of 2 seconds. This value controls Configuration BPDUs as the Root Bridge generates them. Other bridges propagate BPDUs from the Root Bridge as they are received.
If BPDUs stop arriving for the time interval ranging from 2 to 20 seconds because of a network disturbance, or if the Root Bridges stop sending periodic BPDUs during this time, the timer will expire. 2 to 20 seconds is the range between the expected receipt of a BPDU and the expiration of the Max Age time. If the outage lasts for more than 20 seconds, the default Max Age time, the bridge invalidates the saved BPDUs and begins looking for a new Root Port.
Forward Delay is the amount of time the bridge spends in the Listening and Learning states. This is a single value that controls both states. The default value of 15 seconds was originally derived assuming a maximum network size of seven bridge hops, a maximum of three lost BPDUs, and a Hello Time of 2 seconds. The Forward Delay timer also controls the bridge table age-out period after a change in the active topology.
Max Age is the STP timer that controls how long a bridge stores a BPDU before discarding it. Max Age is only an issue when the link failure is not on a directly connected link. When a failure occurs on a directly connected link, the switch knows there will not be any BPDUs coming in on that link, so Max Age is not considered in transitioning the port to Forwarding mode. Recall that each port saves a copy of the best BPDU it has seen. As long as the bridge receives a continuous stream of BPDUs every 2 seconds, the receiving bridge maintains a continuous copy of the BPDU values. However, if the device sending this best BPDU fails, a mechanism must exist to allow other bridges to take over.

QUESTION 200
Refer to the following Certkiller network exhibits:

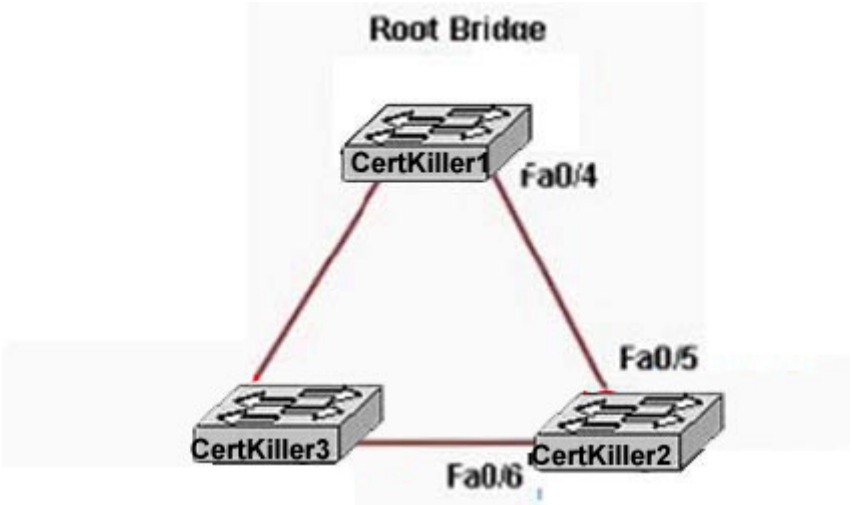

Certkiller 1 configuration exhibit:

```
CertKiller1# show running-config

interface FastEthernet0/4
switchport mode access
no ip address
duplex half
speed 100

<output omitted>
```

Certkiller 2 configuration exhibit:

```
CertKiller1# show running-config

interface FastEthernet0/
switchport mode access
no ip address
duplex full
speed 100

<output omitted>
```

Refer to the network topology exhibit and the partial configuration exhibits of switch Certkiller 1 and Certkiller 2. STP is configured on all switches in the network. Certkiller 2 receives this error message on the console port:
00:06:34: %CDP-4-DUPLEX_MISMATCH: duplex mismatch discovered on FastEthernet0/5 (not half duplex), with Certkiller 1 FastEthernet0/4 (half duplex) ,with TBA05071417(Cat6K-B) 0/4 (half duplex).
What would be the possible outcome of the problem shown in this message?

A. The root port on switch Certkiller 2 will fallback to full-duplex mode.
B. Interface Fa 0/6 on switch Certkiller 2 will transition to a forwarding state and create a bridging loop.
C. The interfaces between switches Certkiller 1 and Certkiller 2 will transition to a blocking state.
D. The root port on switch Certkiller 1 will automatically transition to full-duplex mode.
E. None of the above.

Answer: B

Explanation:
STP States
To participate in STP, each port of a switch must progress through several states. A port begins its life in a Disabled state, moving through several passive states and, finally, into an active state if allowed to forward traffic. The STP port states are as follows:
1. Disabled-Ports that are administratively shut down by the network administrator, or by the system due to a fault condition, are in the Disabled state. This state is special and is not part of the normal STP progression for a port.
2. Blocking-After a port initializes, it begins in the Blocking state so that no bridging loops can form. In the Blocking state, a port cannot receive or transmit data and cannot add MAC addresses to its address table. Instead, a port is allowed to receive only BPDUs so that the switch can hear from other neighboring switches. In addition, ports that are put into standby mode to remove a bridging loop enter the Blocking state.
3. Listening-The port will be moved from Blocking to Listening if the switch thinks that the port can be selected as a Root Port or Designated Port. In other words, the port is on its way to begin forwarding traffic. In the Listening state, the port still cannot send or receive data frames. However, the port is allowed to receive and send BPDUs so that it can actively participate in the Spanning Tree topology process. Here, the port is finally allowed to become a Root Port or Designated Port because the switch can advertise the port by sending BPDUs to other switches. Should the port lose its Root Port or Designated Port status, it returns to the Blocking state.
4. Learning-After a period of time called the Forward Delay in the Listening state, the port is allowed to move into the Learning state. The port still sends and receives BPDUs as before. In addition, the switch can now learn new MAC addresses to add to its address table. This gives the port an extra period of silent participation and allows the switch to assemble at least some address table information.
5. Forwarding-After another Forward Delay period of time in the Learning state, the port is allowed to move into the Forwarding state. The port can now send and receive data frames, collect MAC addresses in its address table, and send and receive BPDUs. The port is now a fully functioning switch port within the Spanning Tree topology.

QUESTION 201
The following "show" command was issued on a Certkiller switch:

```
CertKiller6 # show spanning-tree

VLAN0001
  Spanning tree enabled protocaol eee
  Root ID    Priority    24577
             Address     000a.b724.3c80
             Cost        8
             Port        50 (GigabitEthernet0/2)
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    28673  (priority 28672 sys-id-ext 1)
             Address     0009.e811.7280
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time 300

<Output omitted>
```

Study the exhibit carefully. Based on the output shown above, which statement is

true?

A. Switch Certkiller 6 has been configured with the "spanning-tree vlan 1 hello-time 2" global configuration command.
B. The root bridge has been configured with the "spanning-tree vlan 1 root secondary" global configuration command.
C. Switch Certkiller 6 has been configured with the "spanning-tree vlan 1 priority 24577" global configuration command.
D. Switch Certkiller 6 has been configured with the "spanning-tree vlan 1 root primary" global configuration command.
E. Switch Certkiller 6 has been configured with the "spanning-tree vlan 1 root secondary" global configuration command.
F. None of the above.

Answer: E

Explanation:
To configure a Catalyst switch to become the Root Bridge, use one of the following methods:
1. Directly modify the Bridge Priority value so that a switch can be given a lower-than-default Bridge ID value to win a Root Bridge election:
Switch (config)# spanning-tree vlan vlan-id priority bridge-priority
The bridge-priority value defaults to 32,768, but you can also assign a value of 0 to 65,535. Remember that Catalyst switches run one instance of STP for each VLAN (PVST+), so the VLAN ID must always be given. You should designate an appropriate Root Bridge for each VLAN.
1. Let the switch become the Root by automatically choosing a Bridge Priority value:
Switch(config)# spanning-tree vlan vlan-id root {primary | secondary}
[diameter diameter]
This command is actually a macro on the Catalyst that executes several other commands. The result is a more direct and automatic way to force one switch to become the Root Bridge. Actual Bridge Priorities are not given in the command. Rather, the switch modifies STP values according to the current values in use within the active network. These values are modified only once, when the macro command is issued.
Use the primary keyword to make the switch attempt to become the primary Root Bridge. This command modifies the switch's Bridge Priority value to become less than the Bridge Priority of the current Root Bridge. If the current Root Priority is more than 24,576, the local switch sets its priority to 24,576. If the current Root Priority is less than that, the local switch sets its priority to 4096 less than the current Root. For the secondary Root Bridge, the Root Priority is set to 28,672. There is no way to query or listen to the network to find another potential secondary Root, so this priority is used under the assumption that it is less than the default priorities (32,768) that might be used elsewhere.

QUESTION 202
The Certkiller switched LAN is displayed below:

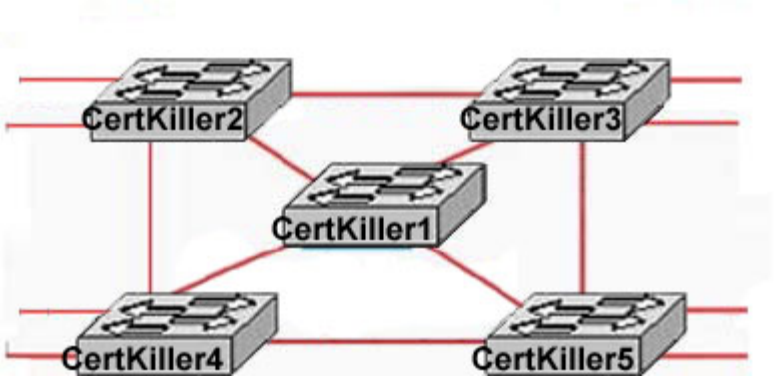

In this network, STP has been implemented. Switch Certkiller 1 is the root switch for the default VLAN. To reduce the broadcast domain, the network administrator decides to split users on the network into VLAN 2 and VLAN 10. The administrator issues the command spanning-tree vlan 2 root primary on switch Certkiller 1. What will happen as a result of this change?

A. Switch Certkiller 1 will change its spanning tree priority to become root for VLAN 2 only.
B. All ports of the root switch Certkiller 1 will remain in forwarding mode throughout the reconvergence of the spanning tree domain.
C. No other switch in the network will be able to become root as long as switch Certkiller 1 is up and running.
D. Switch Certkiller 1 will remain root for the default VLAN and will become root for VLAN 2.
E. None of the above

Answer: D

Explanation:
By default, switches with Cisco PVST and PVST+ maintain a separate spanning-tree instance for each active VLAN configured on it. A bridge ID, consisting of the switch priority and the switch MAC address, is associated with each instance. For each VLAN, the switch with the lowest bridge ID becomes the root switch for that VLAN.
To configure a switch to become the root for the specified VLAN, use the spanning-tree vlan vlan-id root primary global configuration command to modify the switch priority from the default value (32768) to a significantly lower value. When this command is entered, the switch checks the switch priority of the root switches for each VLAN. Because of the extended system ID support, the switch sets its own priority for the specified VLAN to 24576 if this value will cause this switch to become the root for the specified VLAN.
If any root switch for the specified VLAN has a switch priority lower than 24576, the switch sets its own priority for the specified VLAN to 4096 less than the lowest switch priority. 4096 is the value of the least-significant bit of a 4-bit switch priority value.

QUESTION 203
The Certkiller office has been segmented as shown below:

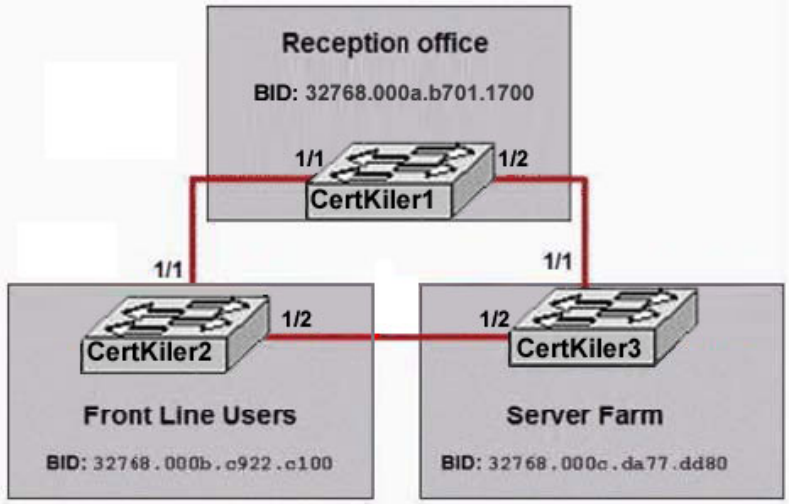

Study the exhibit above carefully. All network links are FastEthernet. Although there is complete connectivity throughout the network, Front Line users have been complaining that they experience slower network performance when accessing the server farm than the Reception office experiences. Based on the exhibit, which two statements are true? (Select two)

A. Disabling the Spanning Tree Protocol would improve network performance.
B. Changing the bridge priority of S3 to 4096 would improve network performance.
C. Changing the bridge priority of S1 to 36864 would improve network performance.
D. Changing the bridge priority of S2 to 36864 would improve network performance.
E. Changing the bridge priority of S1 to 4096 would improve network performance.
F. Upgrading the link between S2 and S3 to Gigabit Ethernet would improve performance.

Answer: B, C

Explanation:
An algorithm is a formula or set of steps for solving a particular problem. Algorithms rely on a set of rules. They have a clear beginning and end. The spanning-tree algorithm is no exception.
The spanning-tree algorithm is defined in the IEEE 802.1D standard. The parameters used by the algorithm, including the Bridge ID, are explored here. The remaining parameters, Path Cost and Port ID, will be covered in the following two topics.
The spanning-tree algorithm characterizes STP. The spanning-tree Algorithm relies on a set of parameters to make decisions. The Bridge ID (BID) is the first parameter used by the spanning-tree algorithm. The Bridge ID (BID) is used by STP to determine the center of the bridged network, known as the Root Bridge. The Bridge ID (BID) parameter is an 8-byte field consisting of an ordered pair of numbers. The first is a 2-byte decimal number called the Bridge Priority, and the second is a 6-byte (hexadecimal) MAC address. The Bridge Priority is a decimal number used to measure the preference of a bridge in the spanning-tree Algorithm. The possible values range between 0 and 65,535. The default setting is 32,768.
The MAC address in the BID is one of the MAC addresses of the switch. Each switch has

a pool of MAC addresses, one for each instance of STP, used as BIDs for the VLAN spanning-tree instances (one per VLAN). For example, Catalyst 6000 switches each have a pool of 1024 MAC addresses assigned to the supervisor module or backplane for this purpose.

QUESTION 204
The Certkiller switched LAN is displayed below:

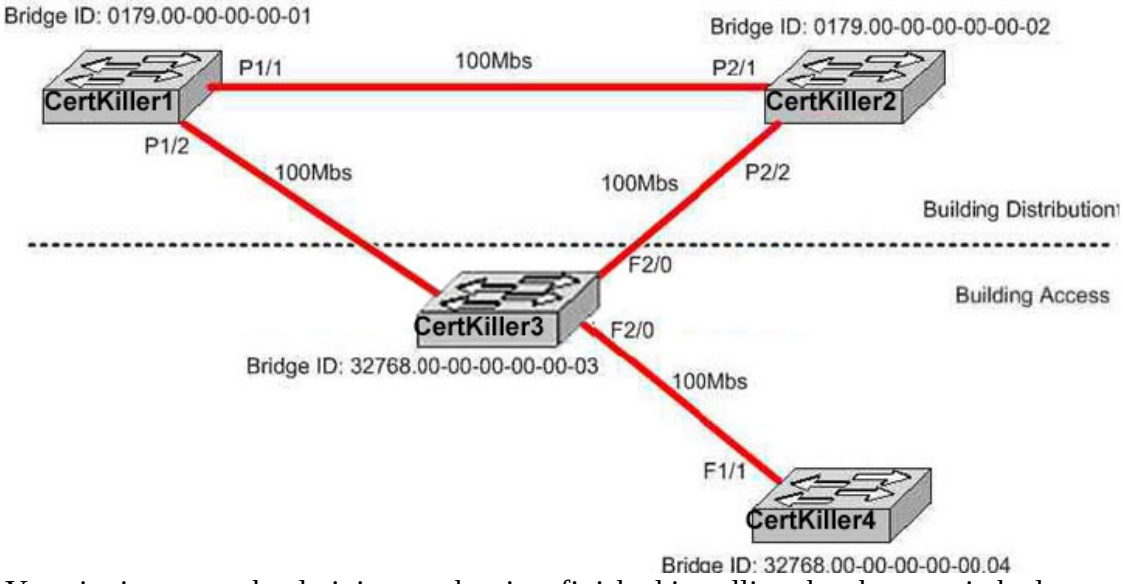

Your junior network administrator has just finished installing the above switched network using Cisco 3550s and would like to manipulate the root bridge election. Which switch should he configure as the root bridge and with which command?

A. Certkiller 1(config)# spanning-tree vlan 1 priority 4096
B. Certkiller 2(config)# set spanning-tree priority 4096
C. Certkiller 3(config)# spanning-tree vlan 1 priority 4096
D. Certkiller 1(config)# set spanning-tree priority 4096
E. Certkiller 2(config)# spanning-tree vlan 1 priority 4096
F. Certkiller 3(config)# set spanning-tree priority 4096

Answer: E

Explanation:
An algorithm is a formula or set of steps for solving a particular problem. Algorithms rely on a set of rules. They have a clear beginning and end. The spanning-tree algorithm is no exception.
The spanning-tree algorithm is defined in the IEEE 802.1D standard. The parameters used by the algorithm, including the Bridge ID, are explored here. The remaining parameters, Path Cost and Port ID, will be covered in the following two topics.
The spanning-tree algorithm characterizes STP. The spanning-tree Algorithm relies on a set of parameters to make decisions. The Bridge ID (BID) is the first parameter used by the spanning-tree algorithm. The Bridge ID (BID) is used by STP to determine the center of the bridged network, known as the Root Bridge. The Bridge ID (BID) parameter is an

8-byte field consisting of an ordered pair of numbers. The first is a 2-byte decimal number called the Bridge Priority, and the second is a 6-byte (hexadecimal) MAC address. The Bridge Priority is a decimal number used to measure the preference of a bridge in the spanning-tree Algorithm. The possible values range between 0 and 65,535. The default setting is 32,768.
The MAC address in the BID is one of the MAC addresses of the switch. Each switch has a pool of MAC addresses, one for each instance of STP, used as BIDs for the VLAN spanning-tree instances (one per VLAN). For example, Catalyst 6000 switches each have a pool of 1024 MAC addresses assigned to the supervisor module or backplane for this purpose.

QUESTION 205
Refer to the output shown on switch CK1 below:
VLAN 1 bridge priority set to 8192.
VLAN 1 bridge max aging time set to 20.
VLAN 1 bridge hello time set to 2.
VLAN 1 bridge forward delay set to 15.
Switch is now the root switch for active VLAN 1.
What command would you enter to reproduce this output? (Type in answer below)

Answer: set spantree root 1

Explanation:
According to Cisco:
The default priority for switches is 32768. This command setting means that the switch will be selected as the root switch because it has the lowest priority. This command will set the bridge priority to 8192, unless another switch on the network is already configured with a priority value less that 8192. If this is the case, the priority will be set to one less this value, ensuring that it will become the root switch.
Note: In STP, a lower bridge priority is preferred over a higher value.

QUESTION 206
Refer to the output shown on switch CK1 below:
Warning: Spantree port fast start should only be enabled on ports connected to a single host. Connecting hubs, concentrators, switches, bridges, etc. to a fast start port can cause temporary spanning tree loops. Use with caution.
Spantree ports 4/1-24 fast start enabled.
What command could you enter to reproduce this output? (Type in answer below)

Answer: set spantree portfast 4/1-24 enable

Explanation:
The output shown in this question is the result of the "set spantree portfast" command. This setting should be configured only on ports that are connected to workstations or PCs. Do not enable portfast on any port connected to another switch.

QUESTION 207
The Certkiller switched LAN is displayed in the diagram below:

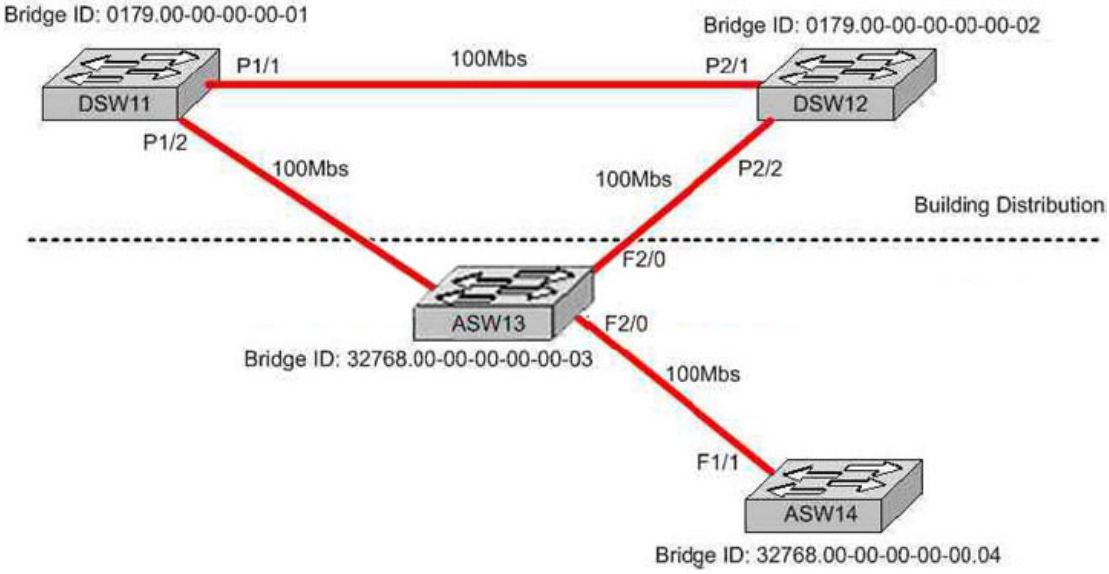

Based on the assumption that STP is enabled on all the switch devices, which of the following statements are true? (Choose two)

A. DSW11 will be elected the root bridge.
B. DSW12 will be elected the root bridge.
C. ASW13 will be elected the root bridge.
D. P1/1 will be elected the nondesignated port.
E. P2/1 will be elected the nondesignated port.
F. F3/0 will be elected the nondesignated port.

Answer: A, F
The root bridge should be placed as close to the core as possible and should be the most centrally located. By default, the switch with the lowest bridge ID will become the root bridge, assuming all other parameters are left as default. This makes DSW11 the root bridge. Also, all ports directly connected to the root bridge will become designated ports, since they are closest to the root bridge. In this case, port F3/0 will become the non-designated port.

QUESTION 208
The Certkiller switched LAN is displayed below:

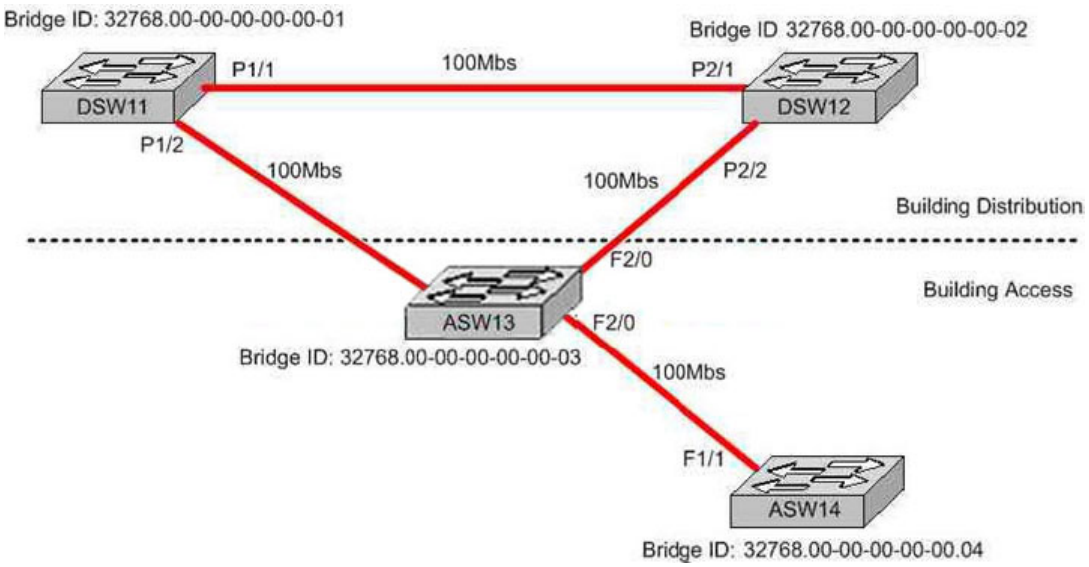

Your junior network administrator has just finished installing the above switched network using Cisco 3550s and would like to manipulate the root bridge election. Which switch should he configure as the root bridge and with which command?

A. DSW11(config)# spanning-tree vlan 1 priority 4096
B. DSW12(config)# set spanning-tree priority 4096
C. ASW13(config)# spanning-tree vlan 1 priority 4096
D. DSW11(config)# set spanning-tree priority 4096
E. DSW12(config)# spanning-tree vlan 1 priority 4096
F. ASW13(config)# set spanning-tree priority 4096

Answer: C

Explanation:
Catalyst 3550 is IOS-based switch, so it doesn't use set-based commands, the correct answer should be 'spanning-tree vlan 1 priority 4096' (Answer C).
Note:
Before configuring STP, you need to select a switch to be the root of the spanning-tree. It does not necessarily have to be the most powerful switch; it should be the most centralized switch on the network. All dataflow across the network will be from the perspective of this switch. It is also important that this switch be the least disturbed switch in the network. The backbone switches are often selected for this function, because they typically do not have end stations connected to them. They are also less likely to be disturbed during moves and changes within the network. In this case, switch ASW13 is the most centrally located switch so it should have its bridge priority lowered to become the root.
Note: In the network shown above, if no configuration changes are made, switch DSW11 will become the root by default, since it has the lowest Bridge ID.

QUESTION 209
SIMULATION

You are connected to switch Certkiller 1 as displayed in the diagram below:

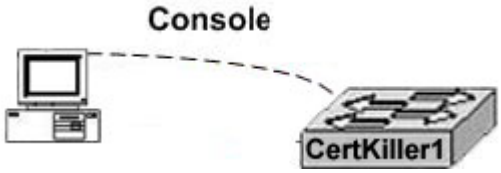

You work as a network engineer at Certkiller .com. The Certkiller .com Toronto office is installing a temporary Catalyst 3550 in an IDF to connect 24 additional users. To prevent network corruption, it is important to have the correct configuration prior to connecting to the production network. It will be necessary to ensure the switch does not participate in VTP but forwards VTP advertisements received on trunk ports.
All interfaces should transition immediately to the forwarding state of Spanning-Tree due to errors that have been experienced on office computers. Also, configure the user ports (All FastEthernet ports) so that he ports are permanently non-trunking.
You will configure FastEthernet ports 0/12 through 0/24 for users who belong to VLAN 20. Also, all VLAN and VTP configurations are to be completed in global configuration mode as VLAN database mode is being deprecated by Cisco.
You are required to accomplish the following tasks:
1. Ensure the switch does not participate in VTP but forwards VTP advertisements received on trunk ports.
2. Ensure all non-trunking interfaces (Fa0/1 to Fa0/24) transition immediately to the forwarding state of Spanning-Tree.
3. Ensure all FastEthernet interfaces are in a permanent non-trunking mode.
4. Place FastEthernet interfaces 0/12 through 0/24 in VLAN 20

Answer:
Configuration:
switch#configure terminal
switch(config)#vtp mode transparent
switch(config)#spanning-tree portfast default
switch(config)#interface range fa0/1 - 24
switch (config-if-range)#switchport mode access
switch (config-if-range)#end
switch#copy running-config startup-config
Alternative:
switch#configure terminal
switch(config)#vtp mode transparent
switch#interface range fa0/1 - 24
switch (config-if-range)#switchport mode access
switch (config-if-range)#spanning-tree portfast
switch (config-if-range)#end
switch#copy running-config startup-config
VTP:
The role of the VLAN Trunking Protocol (VTP) is to maintain VLAN configuration consistency across the entire network. VTP is a messaging protocol that uses Layer 2 trunk frames to

manage the addition, deletion, and renaming of VLANs on a network-wide basis from a centralized switch that is in the VTP server mode. VTP is responsible for synchronizing VLAN information within a VTP domain. This reduces the need to configure the same VLAN information on each switch.
VTP minimizes the possible configuration inconsistencies that arise when changes are made. These inconsistencies can result in security violations, because VLANs can crossconnect when duplicate names are used. They also could become internally disconnected when they are mapped from one LAN type to another, for example, Ethernet to ATM LANE ELANs or FDDI 802.10 VLANs. VTP provides a mapping scheme that enables seamless trunking within a network employing mixed-media technologies.
VTP provides the following benefits:
1. VLAN configuration consistency across the network
2. Mapping scheme that allows a VLAN to be trunked over mixed media
3. Accurate tracking and monitoring of VLANs
4. Dynamic reporting of added VLANs across the network
5. Plug-and-play configuration when adding new VLANs
There are three different VTP modes:
1. Server:
By default, a Catalyst switch is in the VTP server mode and in the "no management domain" state until the switch receives an advertisement for a domain over a trunk link or a VLAN management domain is configured. A switch that has been put in VTP server mode and had a domain name specified can create, modify, and delete VLANs. VTP servers can also specify other configuration parameters such as VTP version and VTP pruning for the entire VTP domain. VTP information is stored in NVRAM.
VTP servers advertise their VLAN configuration to other switches in the same VTP domain, and synchronize the VLAN configuration with other switches based on advertisements received over trunk links. When a change is made to the VLAN configuration on a VTP server, the change is propagated to all switches in the VTP domain. VTP advertisements are transmitted out all trunk connections, including ISL, IEEE 802.1Q, IEEE 802.10, and ATM LANE trunks.
2. Client:
The VTP client maintains a full list of all VLANs within the VTP domain, but it does not store the information in NVRAM. VTP clients behave the same way as VTP servers, but it is not possible to create, change, or delete VLANs on a VTP client. Any changes made must be received from a VTP server advertisement.
3. Transparent
VTP transparent switches do not participate in VTP. A VTP transparent switch does not advertise its VLAN configuration, and does not synchronize its VLAN configuration based on received advertisements. However, in VTP Version 2, transparent switches do forward VTP advertisements that the switches receive out their trunk ports. VLANs can be configured on a switch in the VTP transparent mode, but the information is local to the switch (VLAN information is not propagated to other switches) and is stored in NVRAM.
To change the VTP mode:
Switch(Config)# vtp mode <Mode>
Or
Switch#vlan database
Switch#vtp <mode>
PortFast

A prime reason for enabling PortFast is in cases where a PC boots in a period less than the 30 seconds it takes a switch to put a port into forwarding mode from disconnected state. Some NICs do not enable a link until the MAC layer software driver is actually loaded. Most operating systems try to use the network almost immediately after loading the driver, as in the case of DHCP. This can create a problem because the 30 seconds of STP delay from listening to Forwarding states begins right when the IOS begins trying to access the network. In the case of DHCP, the PC will not obtain a valid IP address from the DHCP server. This problem is common with PC Card (PCMCIA) NICs used in laptop computers. Additionally, there is a race between operating systems and CPU manufacturers. CPU manufacturers keep making the chips faster, while at the same time, operating systems keep slowing down, but the chips are speeding up at a greater rate than the operating systems are slowing down. As a result, PCs are booting faster than ever. In fact, modern machines are often finished booting and need to use the network before the STP 30-second delay is over.
Use the spanning-tree portfast global configuration command to globally enable the PortFast feature on all non-trunking ports.

QUESTION 210
If the root bridge fails, configuration BPDUs will no longer be sent. Which STP timer will have to expire before the other switches can actively restore connectivity with topology change procedure of STP?

A. hello timer
B. BPDU timer
C. Forward_delay timer
D. Max_age timer
E. Dead timer
F. Wait timer

Answer: D

Explanation:
Once a stable network topology has been established, all bridges listen for Hello BPDUs (Bridge Protocol Data Units) transmitted from the root bridge. If a bridge does not get a Hello BPDU after a predefined interval (Max Age), the bridge assumes that the link to the root bridge is down. This bridge then initiates negotiations with other bridges to reconfigure the network to re-establish a valid network topology.
Max age takes into account that the switch at the periphery of the network should not time out the root information under stable condition (that is, if the root is still alive). This is the value that max age needs to take into account the total BPDU propagation delay and the message age overestimate. As such, the formula for max age is as follows:
Max_age
= End-to-end_BPDU_propa_delay + Message_age_overestimate
= 14 + 6
= 20 sec

This explains how IEEE reaches the default recommended value for max age.
Reference: http://www.zyxel.com/support/supportnote/ves1012/app/stp.htm

---

QUESTION 211
Exhibit

Assuming that VLAN 1 and VLAN 2 traffic is enabled on the above network, what effect will the following command have when entered on port 0/2 on switch Certkiller B?
spanning-tree vlan 1 port-priority 16

A. VLAN 1 traffic will be blocked on Switch Certkiller B port 1/1.
B. VLAN 2 traffic will be blocked on Switch Certkiller B port 1/1.
C. VLAN 2 traffic will be blocked on Switch Certkiller A port 0/2.
D. VLAN 1 and 2 traffic will be blocked on Switch Certkiller A port 0/1.
E. VLAN 1 and 2 traffic will be blocked on Switch Certkiller A port 0/2.

Answer: A

Explanation:
Load Sharing Using STP Port Priorities
When two ports on the same switch form a loop, the STP port priority setting determines which port is enabled and which port is in a blocking state. The priorities on a parallel trunk port can be set so that the port carries all the traffic for a given VLAN. The trunk port with the higher priority (lower values) for a VLAN is forwarding traffic for that VLAN. The trunk port with the lower priority (higher values) for the same VLAN remains in a Blocking state for that VLAN. One trunk port sends or receives all traffic for

the VLAN.

	Command	Purpose
Step 1	`spanning-tree vlan 8 port-priority 10`	Assign the port priority of 10 for VLAN 8.
Step 2	`spanning-tree vlan 9 port-priority 10`	Assign the port priority of 10 for VLAN 9.
Step 3	`spanning-tree vlan 10 port-priority 10`	Assign the port priority of 10 for VLAN 10.
Step 4	`exit`	Return to global configuration mode.
Step 5	`interface fastehternet0/2`	Enter interface configuration mode, and define the interface to set the STP port priority.
Step 6	`spanning-tree vlan 3 port-priority 10`	Assign the port priority of 10 for VLAN 3.
Step 7	`spanning-tree vlan 4 port-priority 10`	Assign the port priority of 10 for VLAN 4.
Step 8	`spanning-tree vlan 5 port-priority 10`	Assign the port priority of 10 for VLAN 5.
Step 9	`spanning-tree vlan 6 port-priority 10`	Assign the port priority of 10 for VLAN 6.
Step 10	`end`	Return to privileged EXEC mode.

QUESTION 212
Exhibit
Certkiller 1#show spanning-tree vlan 200
VLAN200
Spanning tree enabled protocol ieee
Root ID Priority 32968
Address 000c.ce29.ef00
Cost 19
Port 2 (FastEthernet0/2)
Hello time 10 Sec Max Age 20 sec Forward Delay 30 sec
Bridge ID Priority 32968 (priority 32768 sys-id-ext 200)
Address 000c.ce2a.4180
Hello Time 2 sec Max Age 20 Sec Forward Delay 15 sec
Interface Role Sts Cost PrioNbr Type

Fa0/2 Root FWD 19 128.2 P2p
Fa0/3 Altn BLK 19 128.3 P2p
Based on the show spanning-tree vlan 200 output shown in the exhibit, which two statements about the STP process for VLAN 200 are true? (Choose two)

A. BDPUs will be sent out every two seconds.
B. The time spent in the listening state will be 30 seconds
C. The time spent in the learning state will be 15 seconds
D. The maximum length of time that the BPDU information will be saved is 30 seconds.
E. This switch is the root bridge for VLAN 200.

F. BPDUs will be sent out every 10 seconds.

Answer: B, F
Changing the Spanning Tree Protocol Timers The STP timers (hello, forward delay, and max age) are included in each BPDU. An IEEE bridge is not concerned about its local configuration of the timers value. It will consider the value of the timers contained in the BPDU that it is receiving. Effectively, that means only a timer configured on the root bridge of the STP is important. Obviously, in case you would lose the root, the new root would start to impose its local timer value to the entire network. So, even if it is not required to configure the same timer value in the entire network, it is at least mandatory to configure any timer changes on the root bridge and on the backup root bridge.

QUESTION 213
What should you do to reduce spanning-tree protocol BPDU traffic during extended periods of instability in your VLANs?

A. Combine all the VLAN spanning trees into a single spanning tree.
B. Set forward delay and max-age timers to the maximum possible values.
C. None of the choices.
D. Change the router VTP server mode.
E. Disable the root bridge

Answer: B

Explanation:
There are several STP timers, as listed below:
1. hello: the hello time is the time between each Bridge Protocol Data Unit (BPDU) that is sent on a port. This is equal to two seconds by default, but can be tuned to be between one and ten seconds.
2. forward delay: the forward delay is the time spent in the listening and learning state. This is by default equal to 15 seconds, but can be tuned to be between four and 30 seconds.
3. max age: the max age timer controls the maximum length of time a bridge port saves its configuration BPDU information. This is 20 seconds by default and can be tuned to be between six and 40 seconds.
The STP timers (hello, forward delay, and max age) are included in each BPDU. An IEEE bridge is not concerned about its local configuration of the timers value. It will consider the value of the timers contained in the BPDU that it is receiving. Effectively, that means only a timer configured on the root bridge of the STP is important. Obviously, in case you would lose the root, the new root would start to impose its local timer value to the entire network. So, even if it is not required to configure the same timer value in the entire network, it is at least mandatory to configure any timer changes on the root bridge and on the backup root bridge.
In order to reduce the number of BPDU's in the spanning tree topology, the forward delay and max-age timers should be increased. This will reduce the BPDU traffic, but it will also increase the convergence time during a topology change.

QUESTION 214
You are network consultant troubleshooting a problem at Certkiller Inc. The local technician tells you that users can't access the Domain Controllers or DHCP servers from their workstations. To top it off, they aren't seeing their Novel Login Screen and they can't access their AppleTalk network. The customer use Cisco 4000, Cisco 5000, and Cisco 6000 switches. What command could you use to resolve these problems?

A. spanning-tree portfast
B. set port connect mod/port
C. spantree start-forwarding
D. set spantree portfast mod/port enable

Answer: D

Explanation:
When the switch powers up, or when a device is connected to a port, the port normally enters the spanning tree listening state. When the forward delay timer expires, the port enters the learning state. When the forward delay timer expires a second time, the port is transitioned to the forwarding or blocking state. This delay could cause the problems described in the scenario. We remove the delay with the PortFast feature. We enable PortFast on a switch port connected to a single workstation or server with the set spantree portfast mod_num/port_num enable command.
Note: The spanning tree PortFast feature causes a port to enter the spanning tree forwarding state immediately, bypassing the listening and learning states. You can use PortFast on switch ports connected to a single workstation or server to allow those devices to connect to the network immediately, instead of waiting for the port to transition from the listening and learning states to the forwarding state.
Reference: Cisco, Configuring Spanning Tree PortFast, UplinkFast, and BackboneFast

QUESTION 215
What command should you enter if do you want to find out whether or not the Backbone Fast convergence feature of STP is enabled on switch CK1 ? (Type in answer below):

Answer: show spantree backbonefast

Explanation:
The following list various commands to use for troubleshooting Catalyst switches:
show spantree vlan_id - Shows the current state of the spanning tree for the "vlan_id" entered from the perspective of the switch on which it is entered.
show spantree summary - Provides a summary of connected spanning tree ports by VLAN.
show spantree statistics - Shows spanning tree statistical information.
show spantree backbonefast - Displays whether the spanning tree Backbone Fast

Convergence feature is enabled.
show spantree blockedports - Displays only the blocked ports.
show spantree portstate - Determines the current spanning tree state of a Token Ring port within a spanning tree.
show spantree portvlancost - Shows the path cost for the VLANs on a port.
show spantree uplinkfast - Shows the uplinkfast settings.

QUESTION 216
The Certkiller network is displayed in the diagram below:

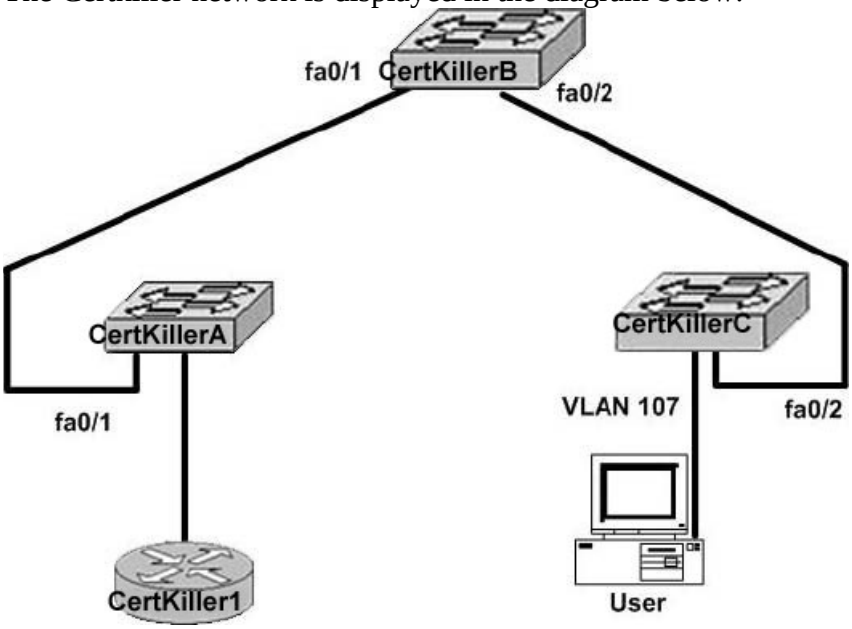

You use the following information for switch Certkiller A:
Port Mode Encapsulation Status Native VLAN
fa0/1 desirable n-802.1q trunking 5
Port VLANs is allowed on trunk
fa0/ 1 1-100, 102-1005
Port VLANs is owned and active in management domain
fa0/1 1-6. 8-100, 102-115, 197-999, 1002-1005
Port VLANs in spanning tree forwarding state and not pruned
fa0/1 1-6, 8-100, 102-105, 108-999, 1002-1005
Certkiller users in VLAN 107 complain that they are unable to gain access to the resources through the Certkiller 1 router.
What is the cause of this problem?

A. VLAN 107 is not configured on the trunk.
B. VLAN 107 does not exist on switch Certkiller A.
C. VTP is pruning VLAN 107.
D. Spanning tree is not enabled on VLAN 107.
E. None of the above

Answer: C

Explanation:
In this example, VLAN 7, 101, 106, and 107 are being pruned. VLAN 107 is being pruned incorrectly in this case. By disabling VTP pruning, VLAN 107 should be able to once again gain access to the network resources.
Incorrect Answers:
A, B: Based on the output shown above, VLAN 107 is known and active within the management domain. Therefore, it must have been configured and the VLAN is indeed allowed to traverse the trunk. Only VLAN 101 has been configured to not pass along this trunk.
D: By default, STP is enabled on all VLANs.

QUESTION 217
Which of the following commands would you enter if you wanted to display spanning tree statistical information?

A. show spantree backbonefast
B. show spantree statistics
C. show spantree uplinkfast
D. show spantree blockedports
E. show spantree portstate
F. show spantree portvlancost

Answer: B

Explanation:
The command 'show spantree statistics' is the correct IOS command to show spanning tree statistical information and is obviously the correct answer choice.
The following list various commands to use for troubleshooting Catalyst switches:
show spantree vlan_id - Shows the current state of the spanning tree for the "vlan_id" entered from the perspective of the switch on which it is entered.
show spantree summary - Provides a summary of connected spanning tree ports by VLAN.
show spantree statistics - Shows spanning tree statistical information.
show spantree backbonefast - Displays whether the spanning tree Backbone Fast Convergence feature is enabled.
show spantree blockedports - Displays only the blocked ports.
show spantree portstate - Determines the current spanning tree state of a Token Ring port within a spanning tree.
show spantree portvlancost - Shows the path cost for the VLANs on a port.
show spantree uplinkfast - Shows the uplinkfast settings.

QUESTION 218
Is the following statement True or False?
The "show spanning-tree" command only shows information about ports with their red or amber lights on.

A. True
B. There is not enough information to determine
C. False

Answer: C

Explanation:
The show spanning-tree command only displays information for ports with an active link (green light is on). If these conditions are not met, you can issue a show running-configuration command to confirm the configuration.

QUESTION 219
Is the following statement True or False?
For optimal performance you should manually select the root switch.

A. False
B. True
C. There is not enough information to determine

Answer: B

Explanation:
The selection of the root switch for a particular VLAN is very important. You can choose it, or you can let the switches decide on their own. The second option is risky because there may be sub-optimal paths in your network if the root selection process is not controlled by you.

QUESTION 220
On switch CK1 the following output was shown:
VLAN1 is executing the ieee compatible Spanning Tree protocol
Bridge Identifier has priority 32768, address 0030.94fc.0a00
Configured hello time 2, max age 20, forward delay 15
Current root has priority 32768, address 0001.6445.4400
Root port is 323 (FastEthernet6/3), cost of root path is 19
Topology change flag not set, detected flag not set
Number of topology changes 2 last change occurred 00:02:19 ago
from FastEthernet6/1
Times: hold 1, topology change 35, notification 2
hello 2, max age 20, forward delay 15
Timers:hello 0, topology change 0, notification 0, aging 300
Port 323 (FastEthernet6/3) of VLAN1 is forwarding
Port path cost 19, Port priority 128, Port Identifier 129.67.
Designated root has priority 32768, address 0001.6445.4400
Designated bridge has priority 32768, address 0001.6445.4400
Designated port id is 129.67, designated path cost 0
Timers:message age 2, forward delay 0, hold 0

Number of transitions to forwarding state:1
BPDU:sent 3, received 91
Which command could you use to reproduce the above output (Type in answer below)

Answer: show spanning-tree vlan 1

Explanation:
This example shows how to display spanning tree information for a specific VLAN:
Switch# showspanning-treevlan1
VLAN1isexecutingtheieeecompatibleSpanningTreeprotocol
BridgeIdentifierhaspriority32768,address0030.94fc.0a00
Configuredhellotime2,maxage20,forwarddelay15
Wearetherootofthespanningtree
Topologychangeflagnotset,detectedflagnotset
Numberoftopologychanges5lastchangeoccurred01:50:47ago
fromFastEthernet6/16
Times:hold1,topologychange35,notification2
hello2,maxage20,forwarddelay15
Timers:hello0,topologychange0,notification0,aging300
Port335(FastEthernet6/15)ofVLAN1isforwarding
Portpathcost19,Portpriority128,PortIdentifier129.79.
Designatedroothaspriority32768,address0030.94fc.0a00
Designatedbridgehaspriority32768,address0030.94fc.0a00
Designatedportidis129.79,designatedpathcost0
Timers:messageage0,forwarddelay0,hold0
Numberoftransitionstoforwardingstate:1
BPDU:sent6127,received0
Switch#
Reference:
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/12_1_12/cmdref/show1.htm#30158

QUESTION 221
What command would you enter if you wanted to display the current state of the spanning tree for the "vlan_id" entered from the perspective of the switch on which it is entered?

A. show spantree id vlan_id
B. show spantree vlan_id state
C. show spantree vlan_id
D. show spantree state vlan_id
E. show spantree vlan vlan_id

Answer: C

Explanation:

Commands to Use for Verifying the Configuration is Working:
show spantree vlan_id - Shows the current state of the spanning tree for the "vlan_id" entered from the perspective of the switch on which it is entered.

QUESTION 222
You want to load balance traffic across your LAN. Which of the methods below are NOT the valid ways to configure load sharing with trunk ports? (Select all that apply)

A. using STP vector metrics
B. using ISL VLAN
C. using STP path costs
D. using STP port priorities
E. using STP SID

Answer: A, B, E

Explanation:
Load sharing divides the bandwidth supplied by parallel trunks connecting switches. To avoid loops, Spanning-Tree Protocol (STP) normally blocks all but one parallel link between switches. With load sharing, you divide the traffic between the links according to which VLAN the traffic belongs to. There are two ways to configure load sharing by using trunk ports: using STP port priorities or using STP path costs.
Incorrect Answers:
C: If you configure load sharing using STP path costs, each load-sharing link can be connected to the same switch or to two different switches.
D: If you configure load sharing using STP port priorities, both load-sharing links must be connected to the same switch.

QUESTION 223
What switch characteristic(s), aside from the MAC address, determines which switch will become the root bridge (Select all that apply)?

A. IP address
B. The port cost
C. Path cost
D. Priority number
E. The port ID

Answer: D

Explanation:
When you configure a switch as the secondary root, the spanning-tree bridge priority is modified from the default value (32768) to 16384 so that the switch is likely to become the root for the specified VLANs if the primary root switch fails (assuming the other switches in the network use the default bridge priority of 32768). The MAC address is

also used in the determination as a tie-breaker if two switches have the same priority value.
Note: In STP, lower is better, meaning that the lower bridge priority is preferred over a higher value.

QUESTION 224
In a CLI based switch, what command will display the information comparable to the IOS command "show span"? (Type in answer below)

Answer: show spantree

Explanation:
Use the show spantree command to display spanning-tree information for a VLAN.
show spantree [vlan | mod_num/port_num] [active]vlan (Optional) Number of the VLAN. If the VLAN number is not specified, the default is VLAN 1.
mod_num (Optional) Number of the module.
port_num (Optional) Number of the port on the module.
active (Optional) Keyword that specifies to display only the active ports.

QUESTION 225
You are a network troubleshooter, and you've arrived at a jobsite to troubleshoot a Catalyst 5000 switch. After talking with the system administrator you come to suspect that the Root Bridge for VLAN 1 is incorrect. Which command would you enter at the CLI to determine VLAN 1's root bridge?

A. show span 1
B. show spantree
C. show bridge vlan 1
D. show spantree root bridge
E. None of the above

Answer: B

Explanation: By default the show spantree command displays the STP information for VLAN 1. The bridge ID, MAC address, and timers are displayed.
Sample output:
Certkiller > (enable) show spantree
VLAN 1
Spanning tree enabled
Spanning tree type ieee
Designated Root 00-d1-22-24-56-00
<Rest of output deleted>
The Designated Root value in the output is the MAC address of the root bridge.

QUESTION 226
Which command would you enter to display the blocked ports on a spanning tree

environment? (Type in answer below)

Answer: show spantree blockedports

Explanation:
Use the show spantree blockedports command to display only the blocked ports.
show spantree blockedports [vlan_num]vlan_num (Optional) Number of the VLAN.
The following list various commands to use for troubleshooting Catalyst switches:
show spantree vlan_id - Shows the current state of the spanning tree for the "vlan_id" entered from the perspective of the switch on which it is entered.
show spantree summary - Provides a summary of connected spanning tree ports by VLAN.
show spantree statistics - Shows spanning tree statistical information.
show spantree backbonefast - Displays whether the spanning tree Backbone Fast Convergence feature is enabled.
show spantree blockedports - Displays only the blocked ports.
show spantree portstate - Determines the current spanning tree state of a Token Ring port within a spanning tree.
show spantree portvlancost - Shows the path cost for the VLANs on a port.
show spantree uplinkfast - Shows the uplinkfast settings.

QUESTION 227
You have a congested Ethernet network with only one Root Bridge on the Certkiller network. What can you do to reduce BPDU traffic on this network?

A. Remove redundant links between switches
B. Decrease the MaxAger timer on all non-Root Bridges
C. Increase the BPDU Hello timer only on the Root Bridge
D. Increase the Path Cost on the Designated Port on all non-Root Bridges

Answer: C

Explanation:
There are several STP timers, as listed below:
1. hello: the hello time is the time between each Bridge Protocol Data Unit (BPDU) that is sent on a port. This is equal to two seconds by default, but can be tuned to be between one and ten seconds.
2. forward delay: the forward delay is the time spent in the listening and learning state. This is by default equal to 15 seconds, but can be tuned to be between four and 30 seconds.
3. max age: the max age timer controls the maximum length of time a bridge port saves its configuration BPDU information. This is 20 seconds by default and can be tuned to be between six and 40 seconds.
The STP timers (hello, forward delay, and max age) are included in each BPDU. An IEEE bridge is not concerned about its local configuration of the timers value.
To reduce BPDU traffic on this network, increase these timers.

Incorrect Answers:
A: Redundant links are not used when STP is in use. STP will block these redundant links automatically.
B: This will be counterproductive, as it will increase the number of BPDU's
D: This will have no effect on the BPDU traffic.

QUESTION 228
Exhibit

Assuming that VLAN 1 and VLAN 2 traffic is enabled on the above network, what effect will the following command have when entered on port 0/2 on switch Certkiller A?
spanning-tree vlan 1 port-priority 16

A. VLAN 1 traffic will be blocked on Switch Certkiller B port 1/1.
B. VLAN 2 traffic will be blocked on Switch Certkiller B port 1/1.
C. VLAN 2 traffic will be blocked on Switch Certkiller A port 0/2.
D. VLAN 1 and 2 traffic will be blocked on Switch Certkiller A port 0/1.
E. VLAN 1 and 2 traffic will be blocked on Switch Certkiller A port 0/2.

Answer: A

Explanation:
Load Sharing Using STP Port Priorities
When two ports on the same switch form a loop, the STP port priority setting determines which port is enabled and which port is in a blocking state. The priorities on a parallel trunk port can be set so that the port carries all the traffic for a given VLAN. The trunk port with the higher priority (lower values) for a VLAN is forwarding traffic for that VLAN. The trunk port with the lower priority (higher values) for the same VLAN remains in a Blocking state for that VLAN. One trunk port sends or receives all traffic for

the VLAN.

	Command	Purpose
Step 1	spanning-tree vlan 8 port-priority 10	Assign the port priority of 10 for VLAN 8.
Step 2	spanning-tree vlan 9 port-priority 10	Assign the port priority of 10 for VLAN 9.
Step 3	spanning-tree vlan 10 port-priority 10	Assign the port priority of 10 for VLAN 10.
Step 4	exit	Return to global configuration mode.
Step 5	interface fastehternet0/2	Enter interface configuration mode, and define the interface to set the STP port priority.
Step 6	spanning-tree vlan 3 port-priority 10	Assign the port priority of 10 for VLAN 3.
Step 7	spanning-tree vlan 4 port-priority 10	Assign the port priority of 10 for VLAN 4.
Step 8	spanning-tree vlan 5 port-priority 10	Assign the port priority of 10 for VLAN 5.
Step 9	spanning-tree vlan 6 port-priority 10	Assign the port priority of 10 for VLAN 6.
Step 10	end	Return to privileged EXEC mode.

QUESTION 229
You need to troubleshoot a port aggregation issue with a Certkiller switch. Study the exhibit below carefully:

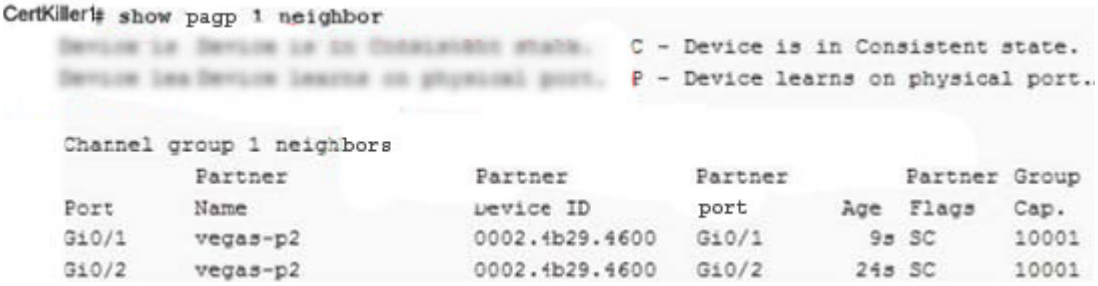

Based on the information shown above, which statement is true about the display of the command "show pagp 1 neighbor" command?

A. STP packets are sent out the Gi0/1 interface only.
B. CDP packets are sent out the Gi0/1 interface only.
C. STP packets are sent out both the Gi0/1 and Gi0/2 interfaces.
D. CDP packets are sent out the Gi0/2 interface only.
E. None of the above

Answer: A

Explanation:
To provide automatic EtherChannel configuration and negotiation between switches, Cisco developed the Port Aggregation Protocol (PAgP). PAgP packets are exchanged

between switches over EtherChannel-capable ports. The identification of neighbors and port group capabilities are learned and compared with local switch capabilities. Ports that have the same neighbor device ID and port group capability are bundled together as a bidirectional, point-to-point EtherChannel link.
PAgP forms an EtherChannel only on ports that are configured for either identical static VLANs or trunking. PAgP also dynamically modifies parameters of the EtherChannel if one of the bundled ports is modified. For example, if the VLAN, speed, or duplex mode of a port in an established bundle is changed, PAgP changes that parameter for all ports in the bundle.
PAgP can be configured in active mode ("desirable"), where a switch actively asks a far-end switch to negotiate an EtherChannel, or in passive mode ("auto," the default), where a switch negotiates an EtherChannel only if the far-end initiates it.

Display Function	Command Syntax
Current EtherChannel status of each member port	**show etherchannel summary** **show etherchannel port**
Timestamps of EtherChannel changes	**show etherchannel port-channel**
Detailed status about each EtherChannel component	**show etherchannel detail**
Load balancing hashing algorithm	**show etherchannel load-balance**
Load balancing port index used by hashing algorithm	**show etherchannel port-channel**
EtherChannel Neighbors on each port	**show {pagp | lacp} neighbor**
LACP System ID	**show lacp sys-id**

QUESTION 230
A new Certkiller switch was configured as shown below:
Study the exhibit carefully. What does the command "channel-group 1 mode desirable" do?

```
interface FastEthernet 0/13
channel-group 1 mode desirable
```

A. It enables PAgP only if a PAgP device is detected
B. It enables PAgP unconditionally
C. It enables LACP only if a LACP device is detected
D. It enables Etherchannel only
E. It enables LACP unconditionally
F. None of the above

Answer: B

Explanation:

To configure switch ports for PAgP negotiation (the default), use the following commands:
Switch(config)# interface type mod/num
Switch(config-if)# channel-protocol pagp
Switch(config-if)# channel-group number mode {on | auto | desirable}
On all IOS-based Catalyst models (3550, 4500, and 6500), you can select between PAgP and LACP as a channel negotiation protocol. The Catalyst 2950, however, offers only PAgP, so the channelprotocol command is not available. Each interface that will be included in a single EtherChannel bundle must be assigned to the same unique channel group number (1 to 64). Channel negotiation must be set to on (unconditionally channel; no PAgP negotiation), auto (passively listen and wait to be asked), or desirable (actively ask).
By default, PAgP operates in "silent" mode with the desirable and auto modes, and allows ports to be added to an EtherChannel even if the other end of the link is silent and never transmits PAgP packets. This might seem to go against the idea of PAgP, where two endpoints negotiate a channel. However, this allows a switch to form an EtherChannel with a device, such as a file server or a network analyzer, that doesn't participate in PAgP. Then, what's the point of running PAgP? Because links should be added to the EtherChannel bundle as PAgP would normally do. In the case of a network analyzer connected to the far end, you might also want to see the PAgP packets generated by the switch, as if you were using a normal PAgP EtherChannel.

QUESTION 231
The following output was seen on a Certkiller switch:

```
CertKiller1# show etherchannel summary

Flap:   D [illegible]          P - in port-channel
        I - [illegible]        S - [illegible]
        R - Layer3             S - Layer2
        u - unsuitable for bundling
        U - port-channel in use
        d - default port

Group   Port-channel      Ports
------+------------+-------------------------------
1       Po1(SU)           Fa0/1(P) Fa0/2(P)
2       Po2(SU)           Fa0/3(P) Fa0/4(P)


 TestKing1# show etherchannel brief

Channel-group listing:
----------------------

Group: 1
----------
Group state = L2
Ports: 2    Maxports = 8
Port-channels: 1 Max Port-channels = 1

Group: 2
----------
Group state = L2
Ports: 2    Maxports = 8
Port-channels: 1 Max Port-channels = 1
```

Study the exhibit above carefully. On the basis of the information that is generated by the two different show commands, which two EtherChannel statements are true? (Select two)

A. Interface Port-Channels 1 and 2 have been assigned IP addresses with the ip address commands.
B. Port-Channels 1 and 2 are capable of combining up to 8 FastEthernet ports to provide full-duplex bandwidth of up to 16 Gbps between a switch and another switch or host.
C. Interfaces FastEthernet 0/3 and 0/4 have been configured with the no switchport command.
D. Switch Certkiller 1 has been configured with a Layer 3 EtherChannel.
E. Interfaces FastEthernet 0/1 and 0/2 have been configured with the channel-group 1 mode desirable command.
F. Port-Channels 1 and 2 are providing two 400 Mbps EtherChannels.

Answer: E, F

Explanation:
To configure switch ports for PAgP negotiation (the default), use the following commands:
Switch(config)# interface type mod/num
Switch(config-if)# channel-protocol pagp
Switch(config-if)# channel-group number mode {on | auto | desirable}
On all IOS-based Catalyst models (3550, 4500, and 6500), you can select between PAgP and LACP as a channel negotiation protocol. The Catalyst 2950, however, offers only PAgP, so the channelprotocol command is not available. Each interface that will be included in a single EtherChannel bundle must be assigned to the same unique channel group number (1 to 64). Channel negotiation must be set to on (unconditionally channel; no PAgP negotiation), auto (passively listen and wait to be asked), or desirable (actively ask).
By default, PAgP operates in "silent" mode with the desirable and auto modes, and allows ports to be added to an EtherChannel even if the other end of the link is silent and never transmits PAgP packets. This might seem to go against the idea of PAgP, where two endpoints negotiate a channel. However, this allows a switch to form an EtherChannel with a device, such as a file server or a network analyzer, that doesn't participate in PAgP. Then, what's the point of running PAgP? Because links should be added to the EtherChannel bundle as PAgP would normally do. In the case of a network analyzer connected to the far end, you might also want to see the PAgP packets generated by the switch, as if you were using a normal PAgP EtherChannel.

QUESTION 232
Two Certkiller core switches are connected via a channel group using LACP. Which three statements are true of the Link Aggregation Control Protocol (LACP)? (Select three)

A. LACP packets are sent with the command channel-group 1 mode desirable.

B. Standby interfaces should be configured with a higher priority.
C. Standby interfaces should be configured with a lower priority.
D. LACP packets are sent with the command channel-group 1 mode active.
E. LACP is used to connect to non-Cisco devices.

Answer: B, D, E

Explanation:
LACP is a standards-based alternative to PAgP, defined in IEEE 802.3ad (also known as IEEE 802.3 Clause 43, "Link Aggregation"). LACP packets are exchanged between switches over EtherChannelcapable ports. Like PAgP, the identification of neighbors and port group capabilities is learned and compared with local switch capabilities. However, LACP also assigns roles to the EtherChannel's endpoints.
The switch with the lowest system priority (a 2-byte priority value followed by a 6-byte switch MAC address) is allowed to make decisions about what ports are actively participating in the EtherChannel at a given time.
Ports are selected and become active according to their port priority value (a 2-byte priority followed by a 2-byte port number), where a low value indicates a higher priority. A set of up to 16 potential links can be defined for each EtherChannel. Through LACP, a switch selects up to eight of these having the lowest port priorities as active EtherChannel links at any given time. The other links are placed in a standby state and will be enabled in the EtherChannel if one of the active links goes down.
Like PAgP, LACP can be configured in active mode ("active"), where a switch actively asks a farend switch to negotiate an EtherChannel, or in passive mode ("passive"), where a switch negotiates an EtherChannel only if the far-end initiates it.
To configure switch ports for LACP negotiation, use the following commands:
Switch(config)# lacp system-priority priority
Switch(config)# interface type mod/num
Switch(config-if)# channel-protocol lacp
Switch(config-if)# channel-group number mode {on | passive | active}
Switch(config-if)# lacp port-priority priority
First, the switch should have its LACP system priority defined (1 to 65,535, default 32,768). If desired, one switch should be assigned a lower system priority than the other so that it can make decisions about the EtherChannel's makeup. Otherwise, both switches will have the same system priority (32,768), and the one with the lower MAC address will become the decision-maker.

QUESTION 233
Two Certkiller core switches are connected as shown below:

Configuration exhibit:

```
CertKiller1# conf t
CertKiller1(config)#interface range gigabitethernet3/1 -2
CertKiller1(config-if)# channel-group 5 mode active
```

Study the exhibits carefully. LACP has been configured on Certkiller 1 as shown. Which is the correct command set to configure LACP on Certkiller 2?

A. Certkiller 2# configure terminal
Certkiller 2(config)# interface range gigabitethernet3/1 -2
Certkiller 2(config-if)# channel-group 5 mode desirable
B. Certkiller 2# configure terminal
Certkiller 2(config)# interface range gigabitethernet3/1 -2
Certkiller 2(config-if)# channel-group 5 mode auto
C. Certkiller 2# configure terminal
Certkiller 2(config)# interface range gigabitethernet3/1 -2
Certkiller 2(config-if)# channel-group 5 mode on
D. Certkiller 2# configure terminal
Certkiller 2(config)# interface range gigabitethernet3/1 -2
Certkiller 2(config-if)# channel-group 5 mode passive
E. None of the above

Answer: D

Explanation:
LACP is a standards-based alternative to PAgP, defined in IEEE 802.3ad (also known as IEEE 802.3 Clause 43, "Link Aggregation"). LACP packets are exchanged between switches over EtherChannelcapable ports. Like PAgP, the identification of neighbors and port group capabilities is learned and compared with local switch capabilities. However, LACP also assigns roles to the EtherChannel's endpoints.
The switch with the lowest system priority (a 2-byte priority value followed by a 6-byte switch MAC address) is allowed to make decisions about what ports are actively participating in the EtherChannel at a given time.
Ports are selected and become active according to their port priority value (a 2-byte priority followed by a 2-byte port number), where a low value indicates a higher priority. A set of up to 16 potential links can be defined for each EtherChannel. Through LACP, a switch selects up to eight of these having the lowest port priorities as active EtherChannel links at any given time. The other links are placed in a standby state and will be enabled in the EtherChannel if one of the active links goes down.
Like PAgP, LACP can be configured in active mode ("active"), where a switch actively asks a farend switch to negotiate an EtherChannel, or in passive mode ("passive"), where a switch negotiates an EtherChannel only if the far-end initiates it.
To configure switch ports for LACP negotiation, use the following commands:
Switch(config)# lacp system-priority priority
Switch(config)# interface type mod/num
Switch(config-if)# channel-protocol lacp
Switch(config-if)# channel-group number mode {on | passive | active}
Switch(config-if)# lacp port-priority priority

First, the switch should have its LACP system priority defined (1 to 65,535, default 32,768). If desired, one switch should be assigned a lower system priority than the other so that it can make decisions about the EtherChannel's makeup. Otherwise, both switches will have the same system priority (32,768), and the one with the lower MAC address will become the decision-maker.

QUESTION 234
Switch CK1 and CK2 are connected via an ethernet channel as shown below:

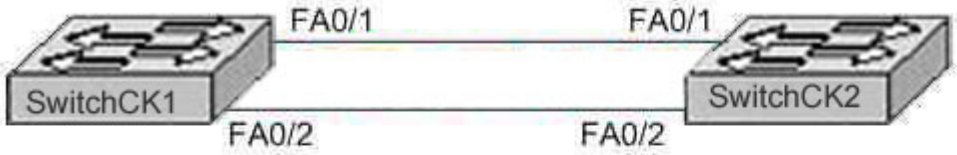

Switch CK1:

```
interface port channel 1
switchport
switchport access vlan 10
interface fastethernet 0/1
channel-group 1 mode passive
interface fastethernet 0/2
channel-group 1 mode passive
```

Switch CK2:

```
interface port-channel 1
switchport
switchport access vlan 10
interface fastethernet 0/1
channel-group 1 mode passive
interface fastethernet 0/2
channel-group 1 mode passive
```

In accordance with the above configuration, which of the statements below is correct?

A. PAgP is correctly configured and the EtherChannel will form.
B. LACP is correctly configured and the EtherChannel will form.
C. One switch must be in LACP Active mode for the EtherChannel to form.
D. Only one switch must be in the On mode and the other in the LACP Passive mode for Etherchannel to form.
E. Each physical port in the EtherChannel must have the command switchport access vlan 10 for the EtherChannel to form.

Answer: C

Explanation:
Link Aggregation Control Protocol (LACP) is part of an IEEE specification (802.3ad) that allows you to bundle several physical ports together to form a single logical channel. LACP allows a switch to negotiate an automatic bundle by sending LACP packets to the peer. It performs a similar function as Port Aggregation Protocol (PAgP) with Cisco EtherChannel.
To start automatic EtherChannel configuration with LACP, configure at least one end of the link to active mode to initiate channelling, because ports in passive mode passively respond to initiation and never initiate the sending of LACP packets.

QUESTION 235
What does an EtherChannel port do if a VLAN range doesn't match its port list?

A. The ports will form EtherChannel if they are set to auto mode.
B. The ports will form EtherChannel if they are all set to the same trunk type.
C. The ports will not form an EtherChannel.
D. The ports will form an EtherChannel if the mode is set to on.

Answer: C

Explanation:
An EtherChannel supports the same allowed range of VLANs on all the ports in a trunking EtherChannel. If the allowed range of VLANs is not the same for a port list, the ports do not form an EtherChannel even when set to the auto or desirable mode with the set port channel command.
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a00800eb

---

QUESTION 236
Which command would you enter on your Catalyst 2900XL switch if you wanted to enable an EtherChannel bundle?

A. Port group
B. Set port channel on
C. Port etherchannel enable
D. Set etherchannel port enable

Answer: A
Explanation
Under the interface command you have to indicate the syntax as port group (group number) for each interface you want to bundle in the etherchannel.
The port group command is used to enable an etherchannel bundle on a Catalyst 2900XL switch.
Incorrect Answers:
B: The set port channel command is not used on a Catalyst 2900XL switch. It used on other switches, such as Cisco 5000 series however.
C, D: These are invalid configuration commands.

QUESTION 237
You have a Catalyst 5000 and you've just configured an Etherchannel bundle. If one of the links were to fail, how long will it take for traffic to be rerouted to a new link?

A. one minute
B. a few seconds
C. a few milliseconds

D. not until appropriate commands are entered

Answer: C

Explanation:
If a link is lost in an EtherChannel network, traffic is rerouted to one of the other links in just a few milliseconds. This rerouting is automatic, and the time it takes for traffic to get re-routed is normally not noticeable by end users.

QUESTION 238
You are configuring a switching solution and you want to take advantage of the Fast EtherChannel ports. When configuring FastEthernet ports, which precautions can you take to avoid configuration problems which can cause the ports to be automatically disabled? (Select two)

A. Allow some ports in a channel to be partly disabled.
B. Configure ALL the ports in a channel as dynamic.
C. Assign all ports in a channel to the same VLAN
D. Allow some ports in a channel to be disabled.
E. Allow all ports in a channel to be disabled.
F. Configure all ports in a channel to operate at the same speed but in different duplex modes
G. Assign all ports in a channel to the same VLAN or configure them as trunk ports.

Answer: C, G

Explanation:
Cisco's Fast EtherChannel technology builds upon standards based 802.3 full duplex Fast Ethernet to provide network managers a reliable high speed solution for the campus network backbone. Fast EtherChannel provides bandwidth scalability within the campus by providing increments from 200 Mbps to 800 Mbps with multi-gigabit capacity in the future. Fast EtherChannel technology not only solves the immediate problem of scaling bandwidth within the network backbone today, but also paves the path for an evolution to standards-based Gigabit Ethernet and beyond, because Fast EtherChannel technology can be applied to support Gigabit EtherChannel.
In order for a channel to function properly, the aggregated links should be in the same VLAN or the links should be assigned as a trunk. In addition, all links should have identical speed and duplex settings.

QUESTION 239
You have just configured an EtherChannel bundle on switch CK1 and it is now operational on a trunk. Which of the following could cause the disabling of the ports in this bundle? (Select two)

A. Disabling port security
B. Excessive errors on one port

C. Changing the VLAN mode to dynamic
D. Changing the speed attribute of one port in the bundle.

Answer: C, D

Explanation:
C: Do not configure the ports in an EtherChannel as dynamic VLAN ports. It could adversely affect switch performance.
D: All ports in an EtherChannel should be configured to operate at the same speed and duplex mode (full or half duplex).
Reference: Cisco, Configuring Fast EtherChannel and Gigabit EtherChannel

---

QUESTION 240
Switch CK1 is a Catalyst 5000 switch. Which of the following set commands would you use to enable Fast EtherChannel on this switch?

A. "set channel fast"
B. "set port channel"
C. "set link channel"
D. "set etherchannel"
E. None of the above

Answer: B

Explanation:
In order to configure ports on a switch to belong to an etherchannel bundle, use the "set port channel" configuration command. For example: set port channel 3/1-2 will configure ports 3/1 and 3/2 to belong to the channel.
Reference: Cisco Application Note, Understanding and Designing Networks Using Fast EtherChannel Technology

---

QUESTION 241
Cisco switches use a logical operation to determine which links to send EtherChannle traffic. What kind of logical operation is it?

A. OR
B. AND
C. XOR
D. NAND

Answer: C

Explanation:
EtherChannel performs the XOR operation, which works like this:
A B C
0 XOR 0 -> 0

0 XOR 1 -> 1
1 XOR 0 -> 1
1 XOR 1 -> 0

QUESTION 242
Which of the following commands would you enter if you wanted to find out whether or not switch CK1 is capable of supporting EtherChannel?

A. show trunk
B. show interface
C. show port channel
D. show port capabilities

Answer: D

Explanation:
The show port capabilities command will show you the capabilities of the modules and ports in a switch. For example, it will display the type, speed, and duplex.
Incorrect Answers:
A: Trunking information does not include the required information.
B: Interface configuration information does not include the required information.
C: Use the show port channel command to display EtherChannel information for a specific module or port. It will not show the capabilities of the switch of the switch however.

QUESTION 243
Switch A and Switch B are configured for FastEthernet Channel as shown in the diagram below:

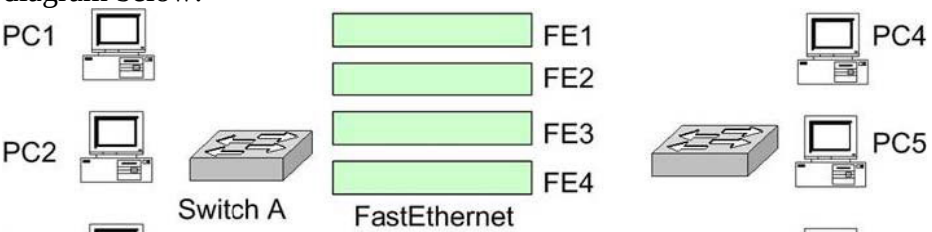

Assuming that Fast EtherChannel was set properly set up; if FE1 were to fail, what would happen with the traffic flow between PC1 & PC5?

A. Traffic is transferred to FE2.
B. Traffic is transferred to FE4.
C. PC1 to PC5 traffic is distributed over the remaining links.
D. The session is disconnected while spanning tree rebuilds.

Answer: C

Explanation:
If a port within an EtherChannel fails, traffic previously carried over the failed port switches to the remaining ports within the EtherChannel.
Note: Fast/Gigabit EtherChannel allow high-speed redundant links in a spanning tree by allowing dual parallel links to be treated as though they were one link. If a link is lost in a Fast/Gigabit EtherChannel network, traffic rerouted to one of the other links in just a few milliseconds.
Reference: Configuring Fast EtherChannel and Gigabit EtherChannel
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/rel7_1/config/channel.htm

---

QUESTION 244
You are configuring a switching solution and you want to take advantage of the Fast EtherChannel ports. When configuring FastEthernet ports, which precautions can you take to avoid configuration problems which can cause the ports to be automatically disabled? (Select two)

A. Allow some ports in a channel to be partly disabled.
B. Configure ALL the ports in a channel as dynamic.
C. Configure all ports in a channel to operate at the same speed and duplex mode
D. Assign all ports in a channel to the same VLAN or configure them as trunk ports.

Answer: C, D

Explanation:
Cisco's Fast EtherChannel technology builds upon standards based 802.3 full duplex Fast Ethernet to provide network managers a reliable high speed solution for the campus network backbone. Fast EtherChannel provides bandwidth scalability within the campus by providing increments from 200 Mbps to 800 Mbps with multi-gigabit capacity in the future. Fast EtherChannel technology not only solves the immediate problem of scaling bandwidth within the network backbone today, but also paves the path for an evolution to standards-based Gigabit Ethernet and beyond, because Fast EtherChannel technology can be applied to support Gigabit EtherChannel.
In order for a channel to function properly, the aggregated links should be in the same VLAN or the links should be assigned as a trunk. In addition, all links should have identical speed and duplex settings.

QUESTION 245
The Certkiller network is displayed in the following network topology exhibit:

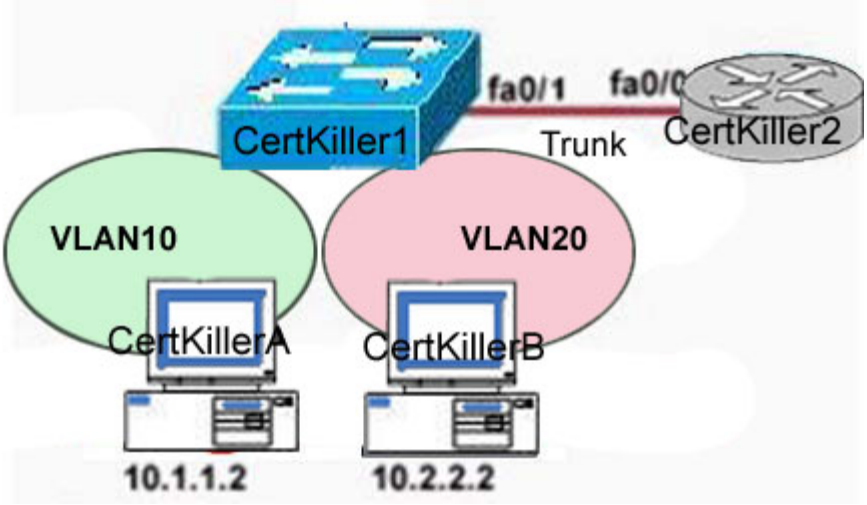

Router configuration exhibit:

```
CertKiller2# show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP EX-ZISRP external 0, OSPE, IA - OSPF inter area
       N - EIGRP EX-ZISRP external 0, OSPE  NSSA external type 2
       E1 - OSPF external type 1, E2- OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - S-IS level-2, ia - IS-IS inter area
       * - c directly connected FastEthernet - ODR
       P - p directly connected FastEthern
Gateway of last resort is not set
     10.1.0.0/24 is subnetted, 2 subnets
C       10.1.1.0 is directly connected, FastEthernet0/0.1
C       10.2.2.0 is directly connected, FastEthernet0/0.2
```

Based on the network diagram and routing table output in the exhibit, which of these statements is true?

A. Although interVLAN routing is not enabled, both workstations will have connectivity to each other.
B. Although interVLAN routing is enabled, the workstations will not have connectivity to each other.
C. InterVLAN routing has been configured properly, and the workstations have connectivity to each other.
D. InterVLAN routing will not occur since no routing protocol has been configured.
E. None of the above.

Answer: C

Explanation:
A Layer 2 network can also exist as a VLAN inside one or more switches. VLANs are essentially isolated from each other so that packets in one VLAN cannot cross into another VLAN.
To transport packets between VLANs, you must use a Layer 3 device. Traditionally, this has been a router's function. The router must have a physical or logical connection to each VLAN so that it can forward packets between them. This is known as interVLAN routing. InterVLAN routing can be performed by an external router that connects to each of the VLANs on a switch. Separate physical connections can be used, or the router can

access each of the VLANs through a single trunk link.
The Switch Port which is connected with Router should be trunk link, You need to configure like:
Switch(config)#interface fa 0/1
Switch(config-if)#switchport mode trunk
Switch(config-if)#switchport trunk encapsulation dot1q
In Router you need to configure like:
Router(config)#interface fa 0/0
Router(config-if)#description VLAN 1
Router(config-if)#ip address 192.168.10.1 255.255.255.0
Router(config)#interface fa 0/0.10
Router(config-subif)#description Management VLAN 10
Router(config-subif)#encapsulation dot1q 10
Router(config-subif)#ip address 192.168.91.1 255.255.255.0
Router(config)#interface fa 0/0.20
Router(config-subif)#description Engineering VLAN 20
Router(config-subif)#encapsulation dot1q 20
Router(config-subif)#ip address 192.168.20.1 255.255.255.0

QUESTION 246
SIMULATION
Network topology exhibit:

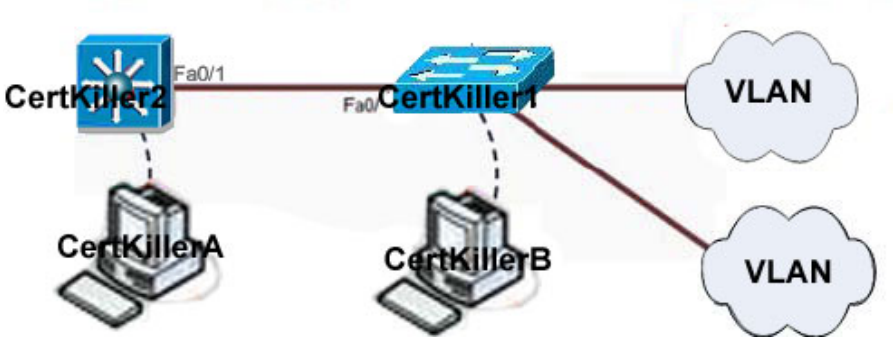

You work as a network engineer at Certkiller .com. Certkiller .com is a large international company with offices on all continents. You work at the New York headquarter. Two layer 3 switches has recently been purchased. Cisco personnel has installed the switch named Certkiller 1 as a distribution layer switch. The switch Certkiller 2 has been cabled to it. Certkiller 2 is used as an access-layer switch.
Your boss at Certkiller .com, Miss Certkiller, has asked you to configure the new switches as follows:
* configure VTP to share VLAN information from Certkiller 1 to the access-layer devices.
* configure interVLAN routing on Certkiller 1 to route traffic between the different VLANs that are configured on the access layer switches
* you are not required to make the specific VLAN port assignments on Certkiller 2
* all VLAN and VTP configuration are to be completed in the global configuration mode (as the Cisco technician has configured the switches for VLAN database mode)

* use the following VTP and VLAN information:
VTP Domain name : Certkiller
VLAN Ids 20 31
IP Addresses 172.16.71.1/24 172.16.132.1/24

Answer:

Explanation:
The information of the question
These are your specific tasks:
1. Configure the VTP information with the distribution layer switch Certkiller 1 as the VTP server
2. Configure the VTP information with the access layer switch Certkiller 2 as a VTP client
3. Configure VLANs on the distribution layer switch Certkiller 1
4. Configure inter-VLAN routing on the distribution layer switch Certkiller 1
5. Specific VLAN port assignments will be made as users are added to the access layer switches in the future.
6. All VLANs and VTP configurations are to completed in the global configuration To configure the switch click on the host icon that is connected to the switch be way of a serial console cable.
vtp server configuration:
Certkiller 1#conf t
Certkiller 1(config)#vtp mode server
Certkiller 1(config)#vtp domain Certkiller
Certkiller 1(config)#vlan 20
Certkiller 1(config)#vlan 31
Certkiller 1(config)#int vlan 20
Certkiller 1(if-config)#ip add 172.64.20.1 255.255.255.0
Certkiller 1(if-config)#no shut
Certkiller 1(if-config)#int vlan 31
Certkiller 1(if-config)#ip add 192.162.31.1 255.255.255.0
Certkiller 1(if-config)#no shut
Certkiller 1(if-config)#exit
Certkiller 1#ip routing
Certkiller 1#copy run start
vtp client configuration:
Certkiller 2#conf t
Certkiller 2(config)#vtp mode client
Certkiller 2(config)#vtp domain CISCO
Certkiller 2#copy run start

QUESTION 247
The Certkiller lab network is shown in the diagram below:

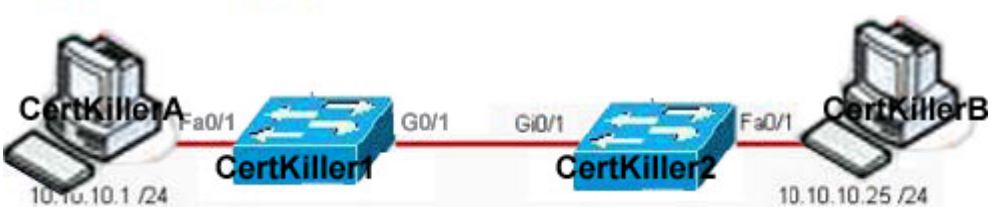

Configuration exhibit Certkiller 1:

```
hostname CertKiller1
interface gigabitethernet 0/1
switchport mode access
switchport access vlan 2

interface fastethernet 0/1
switchport mode access
switchport access vlan 2
```

Configuration exhibit Certkiller 2:

```
hostname CertKiller2
interface gigabitethernet 0/1
switchport mode access
switchport access vlan 5

interface fastethernet 0/1
switchport mode access
switchport access vlan 5
```

Study the exhibits carefully. Both host stations are part of the same subnet but are in different VLANs. On the basis of the information presented in the exhibit, which statement is true about an attempt to ping from host to host?

A. The ping command will be successful without any further configuration changes.
B. A trunk port will need to be configured on the link between Certkiller 1 and Certkiller 2 for the ping command to be successful.
C. A Layer 3 device is needed for the ping command to be successful.
D. The two different hosts will need to be in the same VLAN in order for the ping command to be successful.
E. None of the above.

Answer: C

Explanation:
To transport packets between VLANs, you must use a Layer 3 device. Traditionally, this has been a router's function. The router must have a physical or logical connection to each VLAN so that it can forward packets between them. This is known as interVLAN routing. InterVLAN routing can be performed by an external router that connects to each of the VLANs on a switch. Separate physical connections can be used, or the router can access each of the VLANs through a single trunk link.
See Example on This Figure

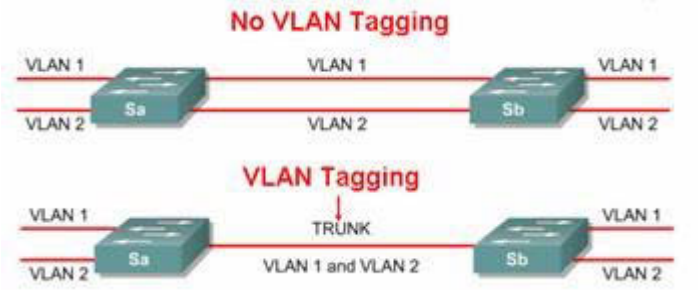

In Exhibit two hosts belongs to two different VLANs so it needs Layer 3 device with Trunk link to transport packets between VLANs.

QUESTION 248
An SVI has been configured on a Certkiller device. Which two statements are true about a switched virtual interface (SVI)? (Select two)

A. An SVI is created for the default VLAN (VLAN1) to permit remote switch administration by default.
B. Multiple SVIs can be associated with a VLAN.
C. SVI is another name for a routed port.
D. An SVI is created by entering the no switchport command in interface configuration mode.
E. An SVI provides a default gateway for a VLAN.

Answer: A, E

Explanation:
On a multilayer switch, you can also enable Layer 3 functionality for an entire VLAN on the switch. This allows a network address to be assigned to a logical interface-that of the VLAN itself. This is useful when the switch has many ports assigned to a common VLAN, and routing is needed in and out of that VLAN.
The logical Layer 3 interface is known as an SVI. However, when it is configured, it uses the much more intuitive interface name vlan vlan-id, as if the VLAN itself is a physical interface. First, define or identify the VLAN interface, and then assign any Layer 3 functionality to it with the following configuration commands:
Switch(config)# interface vlan vlan-id
Switch(config-if)# ip address ip-address mask [secondary]
The VLAN must be defined and active on the switch before the SVI can be used. Make sure the new VLAN interface is also enabled with the no shutdown interface configuration command.

QUESTION 249
Certkiller has implemented numerous multilayer switches that utilize FIB tables. Which statement is true about the Forward Information Base (FIB) table?

A. The FIB is derived from the IP routing table and is optimized for maximum lookup throughput.

B. The FIB table is derived from the Address Resolution Protocol table, and it contains Layer 2 rewrite (MAC) information for the next hop.
C. When the FIB table is full, a wildcard entry redirects traffic to the Layer 3 engine.
D. The FIB lookup is based on the Layer 2 destination MAC address.
E. None of the above

Answer: A

Explanation:
The Layer 3 engine (essentially a router) maintains routing information, whether from static routes or dynamic routing protocols. Basically, the routing table is reformatted into an ordered list with the most specific route first, for each IP destination subnet in the table. The new format is called a Forwarding Information Base (FIB) and contains routing or forwarding information that the network prefix can reference.
In other words, a route to 10.1.0.0/16 might be contained in the FIB, along with routes to 10.1.1.0/24 and 10.1.1.128/25, if those exist. Notice that these examples are increasingly more specific subnets. In the FIB, these would be ordered with the most specific, or longest match, first, followed by less specific subnets. When the switch receives a packet, it can easily examine the destination address and find the longest match entry in the FIB. The FIB also contains the next-hop address for each entry. When a longest match entry is found in the FIB, the Layer 3 next-hop address is found, too.

QUESTION 250
The Certkiller network needs to pass traffic between VLANs. Which device should be used to accomplish this?

A. Hub
B. Switch
C. Router
D. Bridge
E. None of the above

Answer: C

Explanation:
A VLAN is a virtual LAN contained within a switch, so for it to pass information into a different VLAN within the same switch it has to leave that switch and re-enter via a router. VLANs contain local traffic only, so in order to reach users in another VLAN the traffic must go through a router or a layer 3 routing processor.

QUESTION 251
You are configuring a Cisco multilayer switch for the Certkiller network. Which command would you use to configure a port to act as a routed interface?

A. ip routing
B. switchport mode trunk

C. no switchport
D. switchport trunk native vlan 1
E. None of the above

Answer: C

Explanation:
Physical switch ports can also operate as Layer 3 interfaces, where a Layer 3 network address is assigned and routing can occur. Figure 13-2 shows an example of this. By default, all switch ports on the Catalyst 6500 (native IOS) platform operate in the Layer 3 mode. For Layer 3 functionality, you must explicitly configure switch ports with the following command sequence:
Switch(config)# interface type mod/num
Switch(config-if)# no switchport
Switch(config-if)# ip address ip-address mask [secondary]
The no switchport command takes the port out of Layer 2 operation. You can then assign a network address to the port, as you would to a router interface.

QUESTION 252
Inter-VLAN routing has been implemented in the Certkiller network. In VLAN routing, what are some of the disadvantages of designing a router-on-stick configuration? (Select three)

A. InterVLAN routing cannot be filtered by the router.
B. The router becomes a single point of failure for the network.
C. Routers will not route STP BPDUs.
D. There is a possibility of inadequate bandwidth for each VLAN.
E. Additional overhead on the router can occur.
F. NetFlow Switching is required for InterVLAN accounting.

Answer: B, D, E

Explanation:
A router connected to a switch via a single trunk link is better known as router-on-stick or even a one armed router. Since there's only one router, if that router were to go down there'd be no backup. Since there's only one router, that router would have to handle all the bandwidth of every VLAN so there's a chance it could be overloaded, as with the overhead problems of being responsible for too much.
Because traffic routed between the VLANs traverse a single physical port, there is the potential to not provide for enough bandwidth for a VLAN at any given time.
Inter-VLAN routing also does indeed require additional configuration, management, and overhead.
Incorrect Answers:
A: This is not true since routers can indeed filter traffic that is routed between the VLAN subinterfaces.
C: This is not an advantage. Since BPDU's are local to the VLAN, there is generally no

need to route this traffic between the VLANs.
F: This does not apply as a disadvantage to inter-VLAN routing.

QUESTION 253
Which of the following could be used to provide a Layer 3 data path between separate VLANs? (Choose two.)

A. VLAN trunking
B. An external router
C. An internal route processor
D. VLAN capable bridge
E. EtherChannel

Answer: B, C

Explanation:
To transport packets between VLANs, you must use a Layer 3 device. Traditionally, this has been a
router's function. The router must have a physical or logical connection to each VLAN so that it can
forward packets between them. This is known as interVLAN routing.
InterVLAN routing can be performed by an external router that connects to each of the VLANs on
a switch. Separate physical connections can be used, or the router can access each of the VLANs
through a single trunk link.

QUESTION 254
You want to optimize the speed of the switches in the Certkiller network using CEF. Which option correctly identifies the Cisco IOS switching methods in descending order from the fastest method to the slowest method?

A. fast switching, process switching, distributed CEF (dCEF), CEF
B. process switching, distributed CEF (dCEF), CEF, fast switching
C. process switching, CEF, distributed CEF (dCEF), fast switching
D. process switching, fast switching, distributed CEF (dCEF), CEF
E. distributed CEF (dCEF), CEF, fast switching, process switching
F. CEF, distributed CEF (dCEF), fast switching, process switching
G. None of the above

Answer: E

Explanation:
Cisco Express Forwarding (CEF) is a multilayer-switching technology that allows for increased scalability and performance to meet the requirements of large enterprise networks. CEF has evolved to accommodate the traffic patterns realized by modern

networks. These networks are characterized by an increasing number of short-duration flows. Shorter flows are common in environments with a high degree of web-based activity, or other highly interactive types of traffic.
CEF uses these strategies to expediently switch data packets to their destinations. It caches information generated by the Layer 3 routing engine. CEF caches routing information in one table (the FIB) and caches Layer 2 next-hop addresses for all FIB entries in an adjacency table. Because CEF maintains multiple tables for forwarding information, parallel paths can exist and enable CEF to load balance per packet. CEF operates in one of two modes.
Central CEFmode: The CEF FIB and adjacency tables reside on the route processor, and the route processor performs the express forwarding. Use this CEF mode when line cards are not available for CEF switching, or when features are not compatible with distributed CEF.
Distributed Cisco Express Forwarding (dCEF) mode: dCEF is supported on only Cisco Catalyst 6500 switches. When dCEF is enabled, line cards maintain identical copies of the FIB and adjacency tables. The line cards can perform the express forwarding by themselves, relieving the main processor of involvement in the switching operation. dCEF uses an interprocess communications (IPC) mechanism to ensure synchronization of FIBs and adjacency tables on the route processor and line cards.
The Cisco term "punt" describes the action of sending a packet "down" to the next-fastest switching level. This list defines the order of preferred Cisco IOS switching methods, from fastest to slowest.
1. Distributed CEF
2. CEF
3. Fast switching
4. Process switching

QUESTION 255
The following output was seen on a Certkiller device as shown below:

```
CertKiller1# show ip cef ethernet0/0 172.19.233.33 detail

IP CEF with switching (Table [illegible] Version 136808)
45800 routes, 8 unresolved routes (0 old, 8 new) 45800
leaves, 2868 nodes, 8444360 bytes,
136808 inserts, 91008 invalidations 1 load sharing
elements, 208 bytes, 1 references 1 CEF
resets, 1 revisions of existing leaves refcounts: 527343
[illegible]. 465638 [illegible]

172.19.233.33/32, version 417, cached adjacency
172.19.233.33 0 packets, 0 bytes,
Adjacency-prefix
via 172.19.233.33, Ethernet0/0, 0 dependencies
next hop 172.19.233.33, Ethernet0/0
valid cached adjacency
```

Study the exhibit carefully. For what purpose is the command "show ip cef" used?

A. To display CEF-based MLS lookups

B. To display ARP throttling
C. To display TCAM matches
D. To display rewritten IP unicast packets
E. To display entries in the Forwarding Information Base (FIB)
F. To display ARP resolution packets
G. None of the above

Answer: E

Explanation:
Cisco Express Forwarding (CEF) is a multilayer-switching technology that allows for increased scalability and performance to meet the requirements of large enterprise networks. CEF has evolved to accommodate the traffic patterns realized by modern networks. These networks are characterized by an increasing number of short-duration flows. Shorter flows are common in environments with a high degree of web-based activity, or other highly interactive types of traffic.
The contents of the FIB table can be viewed by issuing the command show ip cef detail from the MSFC2. The command show ip cef can be used to view the contents of the CEF adjacency table from the MSFC2. The command show ip cef summary provides a brief overview of the CEF process. It shows information such as the total number of adjacencies and routes.
Example of show ip cef detail

```
MSFC2#sh ip cef detail
  IP CEF with switching (Table Version 477965)
  445 routes, 0 reresolve, 0 unresolved (0 old, 0 new)
  446 leaves, 76 nodes, 132560 bytes, 477966 inserts,
477520 invalidations
  0 load sharing elements, 0 bytes, 0 references
  1 CEF resets, 2 revisions of existing leaves
  refcounts:  15824 leaf, 15038 node
  Default 192.35.86.0/24
0.0.0.0/32, version 0, receive
10.1.0.0/16, version 121980, attached, connected
0 packets, 0 bytes
  via Vlan10, 0 dependencies
    valid glean adjacency
```

QUESTION 256
The following output was displayed on a Certkiller switch.

```
CertKiller1 # show ip pef 192.168.150.0
  192.168.180.0/24, version 214,cached adjacency 192.168.199.3
  0 packets, 0 bytes
  via 192.168.199.3, VLAN 199, 0 dependencies
  next-hop 192.168.199.3, VLAN 199
  valid cached adjacency

CertKiller1# show adjacency detail | begin 192.168.199.3
  IP VLAN 199 192.168.199.3(7)
  0 packets, 0 bytes
  003071506800
  .....
  ...
  .
```

Study the exhibit carefully. The Certkiller administrator is verifying that a CEF FIB entry exists to destination network 192.168.150.0. Given the output generated by the "show ip cef" and "show adjacency detail" commands, which three statements are true? (Select three)

A. The "valid cached adjacency" entry indicates that CEF will put all packets going to such an adjacency to the next best switching mode.
B. The number 003071506800 is the MAC address of the source IP address.
C. There is a valid CEF entry for the destination network 192.168.150.0.
D. The number 003071506800 is the MAC address of the 192.168.199.3 next hop IP address.
E. The counters (0 packets, 0 bytes) indicate a problem with the 192.168.199.3 next hop IP address.
F. There is an adjacency for the 192.168.199.3 next hop IP address.

Answer: C, D, F

Explanation:
The adjacency table contents are fundamentally a function of the ARP process, whereby Layer 2 addresses are mapped to corresponding Layer 3 addresses. When the router issues an ARP request, a corresponding reply is received, and a host entry is added to the adjacency table to reflect this. In addition, the router can also glean next hop routers from routing updates and make entries in the adjacency table to reflect this. This lets the router build the next hop rewrite information necessary for Layer 3 packet forwarding. By having this data already stored in a table, CEF can perform highly efficient and consistent forwarding, because no discovery process is required. The command
show ip cef is used to view the contents of the CEF adjacency table from the MSFC2.
The command show ip cef summary gives a brief overview of the CEF process. It shows information such as the total number of adjacencies and routes.

```
Switch#show adjacency gigabitethernet 9/5 detail
Protocol  Interface                 Address
IP        GigabitEthernet9/5        172.20.53.206(11)
                                    504 packets, 6110 bytes
                                    00605C865B82
                                    000164F83FA50800
                                    ARP       03:49:31
```

Each time an adjacency entry is created, a Layer 2 data link layer header for that adjacent node is precomputed and stored in the adjacency table. This information is subsequently used for encapsulation during CEF switching of packets. Output from the command show adjacency detail displays the content of the information to be used during this Layer 2 encapsulation. Verify that the header information is displayed as would be expected during Layer 2 operations, not using precomputed encapsulation from the adjacency table. Adjacency statistics are updated approximately every 60 seconds. Also, the show cef drops command will display an indication of packets that are being dropped due to adjacencies that are either incomplete or nonexistent. There are two known reasons for incomplete or nonexistent adjacencies. The router cannot use ARP successfully for the next-hop interface. After a clear ip arp or a clear adjacency command, the router marks the adjacency as incomplete, and then it fails to clear the entry.
The symptoms of an incomplete adjacency include random packet drops during a ping test. Use the debug ip cef command to view CEF drops caused by an incomplete adjacency.

QUESTION 257
Which process plays a major role in the creation of the CEF adjacency table?

A. Address Resolution Protocol (ARP)
B. PDU header rewrite
C. NetFlow Switching
D. hello packet exchange
E. None of the above

Answer: A

Explanation:
The adjacency table information is built from the ARP table. As a next-hop address receives a valid ARP entry, the adjacency table is updated. If an ARP entry does not exist, the FIB entry is marked as "CEF glean." This means that the Layer 3 forwarding engine can't forward the packet in hardware, due to the missing Layer 2 next-hop address. The packet is sent to the Layer 3 engine so that it can generate an ARP request and receive an ARP reply. This is known as the "CEF glean" state, where the Layer 3 engine must glean the next-hop destination's MAC address.
During the time that a FIB entry is in the CEF glean state waiting for the ARP resolution, subsequent packets to that host are immediately dropped so that the input queues do not fill and the Layer 3 engine does not become too busy worrying about the need for duplicate ARP requests. This is called ARP throttling or throttling adjacency. If an ARP reply is not received in two seconds, the throttling is released so that another ARP request can be triggered. Otherwise, after an ARP reply is received, the throttling is released, the FIB entry can be completed, and packets can be forwarded completely in hardware.

QUESTION 258
The following output was seen on a Certkiller switch:

```
Switch# show ip def 100.168.150.0

192.168.150.0,24, version 290, cached adjecency 192.168.199.3
0 packets, 0 bytes
via 192.168.199.3, VLAN 199, 0 dependencies
next-hop 192.168.199.3, VLAN 199
valid cached adjacency

Switch# show adjacency detail | begin 192.168.199.3

IP VLAN 199 192.168.199.3(7)
0 packets, 0 bytes
003071006340
.....
...
.
```

Refer to the exhibit. An administrator is verifying that a CEF FIB entry exists to destination network 192.168.150.0. Given the output generated by the show ip cef and show adjacency detail commands, which three statements are true? (Choose three.)

A. There is a valid CEF entry for the destination network 192.168.150.0.
B. The "valid cached adjacency" entry indicates that CEF will put all packets going to such an adjacency to the next best switching mode.
C. The counters (0 packets, 0 bytes) indicate a problem with the 192.168.199.3 next hop IP address.
D. There is an adjacency for the 192.168.199.3 next hop IP address.
E. The number 003071506800 is the MAC address of the 192.168.199.3 next hop IP address.
F. The number 003071506800 is the MAC address of the source IP address.

Answer: A, D, E

Explanation:
The adjacency table contents are fundamentally a function of the ARP process, whereby Layer 2 addresses are mapped to corresponding Layer 3 addresses. When the router issues an ARP request, a corresponding reply is received, and a host entry is added to the adjacency table to reflect this. In addition, the router can also glean next hop routers from routing updates and make entries in the adjacency table to reflect this. This lets the router build the next hop rewrite information necessary for Layer 3 packet forwarding. By having this data already stored in a table, CEF can perform highly efficient and consistent forwarding, because no discovery process is required. The command show ip cef is used to view the contents of the CEF adjacency table from the MSFC2. The command show ip cef summary gives a brief overview of the CEF process. It shows information such as the total number of adjacencies and routes.

```
Switch#show adjacency gigabitethernet 9/5 detail
Protocol  Interface                 Address
IP        GigabitEthernet9/5        172.20.53.206(11)
                                    504 packets, 6110 bytes
                                    00605C865B82
                                    000164F83FA50800
                                    ARP        03:49:31
```

Each time an adjacency entry is created, a Layer 2 data link layer header for that adjacent node is precomputed and stored in the adjacency table. This information is subsequently used for encapsulation during CEF switching of packets. Output from the command show adjacency detail displays the content of the information to be used during this Layer 2 encapsulation. Verify that the header information is displayed as would be expected during Layer 2 operations, not using precomputed encapsulation from the adjacency table. Adjacency statistics are updated approximately every 60 seconds. Also, the show cef drops command will display an indication of packets that are being dropped due to adjacencies that are either incomplete or nonexistent. There are two known reasons for incomplete or nonexistent adjacencies. The router cannot use ARP successfully for the next-hop interface. After a clear ip arp or a clear adjacency command, the router marks the adjacency as incomplete, and then it fails to clear the entry.
The symptoms of an incomplete adjacency include random packet drops during a ping test. Use the debug ip cef command to view CEF drops caused by an incomplete adjacency.

QUESTION 259
Network topology exhibit:

Configuration exhibit:

```
CertKillerSwitchg# show ip route
Default gateway is not set

Home          Hostway          LastUse    LocalUser Interface
ICMP redirect created

CertKillerSwitchg# sho vlan brief

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Fa0/1, Fa0/2, Fa0/3, Fa0/4
                                                Fa0/5, Fa0/6, Fa0/7, Fa0/8
                                                Fa0/9, Fa0/20, Fa0/21, Fa0/22
                                                Fa0/23, Fa0/24, Gi0/1, Gi0/2
10   EAIN                             active    Fa0/10, Fa0/11, Fa0/12, Fa0/13
                                                Fa0/14
15   ADMIN                            active    Fa0/15, Fa0/16, Fa0/17, Fa0/18
                                                Fa0/19

<Output Omitted>

CertKillerSwitchg#
```

Study the Certkiller exhibits shown above. Host Certkiller A and Host Certkiller B are

connected to the Catalyst 3550 switch and have been assigned to their respective VLANs. The rest of the 3550 configuration is the default configuration. Host Certkiller A is able to ping its default gateway, 10.10.10.1, but is unable to ping Host Certkiller B. Given the output displayed in the exhibit, which statement is true?

A. A separate router is required to support interVLAN routing.
B. VTP must be configured to support interVLAN routing.
C. The global configuration command "ip routing" must be configured on the Certkiller Switch switch.
D. HSRP must be configured on Certkiller Switch.
E. VLANs 10 and 15 must be created in the VLAN database mode.
F. Interface VLAN 10 must be configured on the Certkiller Switch switch.
G. None of the above

Answer: C

Explanation:
To transport packets between VLANs, you must use a Layer 3 device. Traditionally, this has been a router's function. The router must have a physical or logical connection to each VLAN so that it can forward packets between them. This is known as interVLAN routing.
Multilayer switches can perform both Layer 2 switching and interVLAN routing, as appropriate. Layer 2 switching occurs between interfaces that are assigned to Layer 2 VLANs or Layer 2 trunks. Layer 3 switching can occur between any type of interface, as long as the interface can have a Layer 3 address assigned to it.
Switch(config)#ip routing command enables the routing on Layer 3 Swtich

---

QUESTION 260
Network topology exhibit:

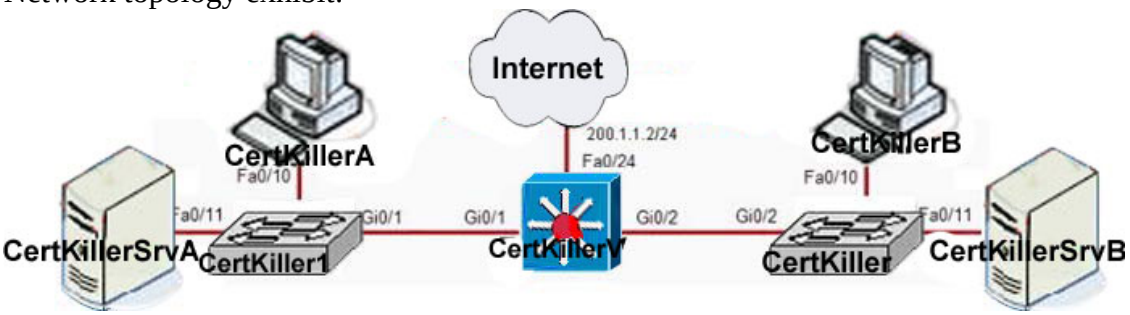

Configuration exhibit:

CertKillerV# show ip route

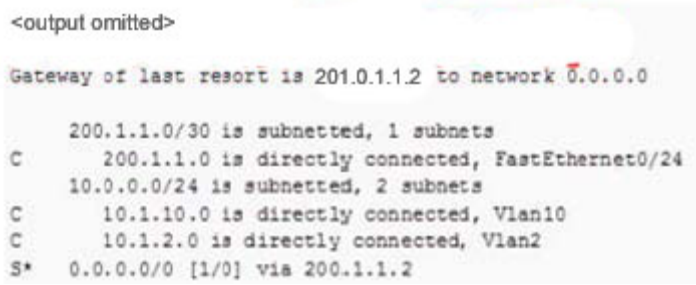

Study the exhibits above carefully. VLAN2, VLAN3, and VLAN10 are configured on the switch Certkiller V. Host computers are on VLAN 2 (10.1.2.0), servers are on VLAN 3 (10.1.3.0), and the management VLAN is on VLAN10 (10.1.10.0). Hosts are able to ping each other but are unable to reach the servers. Based on the information above, which configuration solution could rectify the problem?

A. Assign an IP address of 10.1.3.1/24 to VLAN3.
B. Configure default gateways to IP address 10.1.2.1 on each host.
C. Enable IP routing on the switch Certkiller V.
D. Configure a default route that points toward network 200.1.1.0/24.
E. Configure default gateways to IP address 200.1.1.2 on each host.
F. Configure default gateways to IP address 10.1.10.1 on each host.
G. None of the above

Answer: A

Explanation:
Although a routed port is configured for connectivity with an external router, Inter-VLAN routing would most likely be achieved through the use of a virtual interface.
Example:
To route between VLANs 10 and 20 which have been configured on the multilayer switch use the following configuration:
RouteSwitch(config)#interface vlan 10
RouteSwitch(config-if)#ip address 10.0.10.1 255.255.255.0
RouteSwitch(config)#interface vlan 20
RouteSwitch(config-if)#ip address 10.0.20.1 255.255.255.0

QUESTION 261
Part of the configuration file of a Certkiller switch is shown below:

Study the exhibit above carefully. PCs in VLAN 2 are not able to communicate with PCs in VLAN 3. What could be the cause?

A. VTP is not configured correctly on the interfaces.
B. The command "mls rp ip" must be disabled to enable the routing.
C. The command "mls rp management-interface" is missing.
D. IP routing is not enabled.
E. None of the above.

Answer: D

Explanation:
To transport packets between VLANs, you must use a Layer 3 device. Traditionally, this has been a router's function. The router must have a physical or logical connection to each VLAN so that it can forward packets between them. This is known as interVLAN routing.
Multilayer switches can perform both Layer 2 switching and interVLAN routing, as appropriate. Layer 2 switching occurs between interfaces that are assigned to Layer 2 VLANs or Layer 2 trunks. Layer 3 switching can occur between any type of interface, as long as the interface can have a Layer 3 address assigned to it.
The first step in troubleshooting Inter-VLAN routing is to ensure that routing is actually enabled using the show ip route command. If no entries are seen in the routing table then IP routing needs to be enabled with the command:
Switch(config)#ip routing

QUESTION 262
You have just purchased a new Cisco 3550 switch running the enhanced IOS and need configure it to be installed in a high availability network. On a 3550 EMI switch, which three types of interfaces can be used to configure HSRP? (Select three)

A. SVI interface
B. Access port
C. EtherChannel port channel

D. Loopback interface
E. Routed port
F. BVI interface

Answer: A, C, E

Explanation:
HSRP is a Cisco-proprietary protocol developed to allow several routers (or multilayer switches) to appear as a single gateway IP address. RFC 2281 describes this protocol in more detail.
Basically, each of the routers that provides redundancy for a given gateway address is assigned to a common HSRP group. One router is elected as the primary, or active, HSRP router; another is elected as the standby HSRP router; and all the others remain in the listen HSRP state. The routers exchange HSRP hello messages at regular intervals so they can remain aware of each other's existence and that of the active router.
An HSRP group can be assigned an arbitrary group number, from 0 to 255. If you configure HSRP groups on several VLAN interfaces, it can be handy to make the group number the same as the VLAN number. However, most Catalyst switches support only up to 16 unique HSRP group numbers. If you have more than 16 VLANs, you will quickly run out of group numbers. An alternative is to make the group number the same (that is, 1) for every VLAN interface. This is perfectly valid because the HSRP groups are only locally significant on an interface. In other words, HSRP Group 1 on interface VLAN 10 is unique and independent from HSRP Group 1 on interface VLAN 11. HSRP can be configured on SVI, Etherchannel and Routed port.

QUESTION 263
You need to configure two Certkiller routers for high availability. Which protocol enables a group of routers to form a single virtual router and use the real IP address of a router as the gateway address?

A. HSRP
B. IRDP
C. Proxy ARP
D. GLBP
E. VRRP
F. None of the above

Answer: E

Explanation:
The Virtual Router Redundancy Protocol (VRRP) is a standards-based alternative to HSRP, defined in IETF standard RFC 2338. VRRP is so similar to HSRP that you need to learn only slightly different terminology and a couple of slight functional differences.
1. VRRP provides one redundant gateway address from a group of routers. The active router is called the master router, while all others are in the backup state. The master router is the one with the highest router priority in the VRRP group.

2. VRRP group numbers range from 0 to 255; router priorities range from 1 to 254 (254 is the highest; 100 is the default).
3. The virtual router MAC address is of the form 0000.5e00.01xx, where xx is a two-digit hex VRRP group number.
4. VRRP advertisements are sent at 1-second intervals. Backup routers can optionally learn the advertisement interval from the master router.
5. By default, all VRRP routers are configured to preempt the current master router, if their priorities are greater.
6. VRRP has no mechanism for tracking interfaces to allow more capable routers to take over the master role.

QUESTION 264
You want to implement a high availability network using two Cisco routers. You also want to ensure that the stat of the WAN interfaces can be tracked in doing so. Which router redundancy protocol cannot be configured for interface tracking?

A. GLBP
B. HSRP
C. RPR
D. VRRP
E. SLB
F. RPR+
G. None of the above

Answer: D

Explanation:
The Virtual Router Redundancy Protocol (VRRP) is a standards-based alternative to HSRP, defined in IETF standard RFC 2338. VRRP is so similar to HSRP that you need to learn only slightly different terminology and a couple of slight functional differences.
1. VRRP provides one redundant gateway address from a group of routers. The active router is called the master router, while all others are in the backup state. The master router is the one with the highest router priority in the VRRP group.
2. VRRP group numbers range from 0 to 255; router priorities range from 1 to 254 (254 is the highest; 100 is the default).
3. The virtual router MAC address is of the form 0000.5e00.01xx, where xx is a two-digit hex VRRP group number.
4. VRRP advertisements are sent at 1-second intervals. Backup routers can optionally learn the advertisement interval from the master router.
5. By default, all VRRP routers are configured to preempt the current master router, if their priorities are greater.
6. VRRP has no mechanism for tracking interfaces to allow more capable routers to take over the master role.

QUESTION 265
Certkiller uses GLBP to provide for router redundancy in the network. Which

describes the default load balancing scheme used by the Gateway Load Balancing Protocol (GLBP)?

A. Per host basis using a strict priority scheme
B. Per session using a round-robin scheme
C. Per session using a strict priority scheme
D. Per GLBP group using a strict priority scheme
E. Per host basis using a round-robin scheme
F. Per GLBP group using a round-robin scheme

Answer: E

Explanation:
To provide a virtual router, multiple switches (routers) are assigned to a common GLBP group. Rather than having just one active router performing forwarding for the virtual router address, all routers in the group can participate and offer load balancing by forwarding a portion of the overall traffic. The advantage is that none of the clients have to be pointed toward a specific gateway address-they can all have the same default gateway set to the virtual router IP address. The load balancing is provided completely through the use of virtual router MAC addresses in ARP replies returned to the clients. As a client sends an ARP request looking for the virtual router address, GLBP sends back an ARP reply with the virtual MAC address of a selected router in the group. The result is that all clients use the same gateway address but have differing MAC addresses for it.

QUESTION 266
You want to implement router redundancy in the Certkiller network using the best method available. Which protocol specified by RFC 2281 provides network redundancy for IP networks, ensuring that user traffic immediately and transparently recovers from first-hop failures in network edge devices or access circuits?

A. ICMP
B. IRDP
C. HSRP
D. STP
E. None of the above

Answer: C

Explanation:
HSRP is a Cisco-proprietary protocol developed to allow several routers (or multilayer switches) to appear as a single gateway address. RFC 2281 describes this protocol in more detail. Basically, each of the routers that provides redundancy for a given gateway address is assigned to a common HSRP group. One router is elected as the primary, or active, HSRP router, another is elected as the standby HSRP router, and all the others remain in the listen HSRP state. The routers exchange HSRP hello messages at regular

intervals, so they can remain aware of each other's existence, as well as that of the active router.
An HSRP group can be assigned an arbitrary group number, from 0 to 255. If you configure HSRP groups on several VLAN interfaces, it can be handy to make the group number the same as the VLAN number. However, most Catalyst switches support only up to 16 unique HSRP group numbers. If you have more than 16 VLANs, you will quickly run out of group numbers. An alternative is to make the group number the same (that is, 1) for every VLAN interface. This is perfectly valid because the HSRP groups are only locally significant on an interface. HSRP Group 1 on interface VLAN 10 is unique from HSRP Group 1 on interface VLAN 11.

QUESTION 267
Routers CK1 and CK2 use HSRP so that one router will back up the other should there be a failure. Which two statements are true about the Hot Standby Router Protocol (HSRP)? (Select two)

A. Load sharing with HSRP is achieved by creating multiple subinterfaces on the HSRP routers.
B. Routers configured for HSRP can belong to multiple groups and multiple VLANs.
C. Load sharing with HSRP is achieved by creating HSRP groups on the HSRP routers.
D. All routers configured for HSRP load balancing must be configured with the same priority.
E. Routers configured for HSRP must belong to only one group per HSRP interface.

Answer: B, C

Explanation:
HSRP is a Cisco-proprietary protocol developed to allow several routers (or multilayer switches) to appear as a single gateway address. RFC 2281 describes this protocol in more detail. Basically, each of the routers that provides redundancy for a given gateway address is assigned to a common HSRP group. One router is elected as the primary, or active, HSRP router, another is elected as the standby HSRP router, and all the others remain in the listen HSRP state. The routers exchange HSRP hello messages at regular intervals, so they can remain aware of each other's existence, as well as that of the active router.
An HSRP group can be assigned an arbitrary group number, from 0 to 255. If you configure HSRP groups on several VLAN interfaces, it can be handy to make the group number the same as the VLAN number. However, most Catalyst switches support only up to 16 unique HSRP group numbers. If you have more than 16 VLANs, you will quickly run out of group numbers. An alternative is to make the group number the same (that is, 1) for every VLAN interface. This is perfectly valid because the HSRP groups are only locally significant on an interface. HSRP Group 1 on interface VLAN 10 is unique from HSRP Group 1 on interface VLAN 11.

QUESTION 268
You need to decide on the best router redundancy protocol to use in the Certkiller

network. Which two statements are true about HSRP, VRRP, and GLBP? (Select two)

A. GLBP and VRRP allow for MD5 authentication, whereas HSRP does not.
B. HSRP allows for multiple upstream active links being simultaneously used, whereas GLBP does not.
C. GLBP allows for router load balancing of traffic from a network segment without the different host IP configurations required to achieve the same results with HSRP.
D. Unlike HSRP and VRRP, GLBP allows automatic selection and simultaneous use of multiple available gateways.
E. GLBP allows for router load balancing of traffic from a network segment by utilizing the creation of multiple standby groups.

Answer: C, D

Explanation:
1. GLBP
To provide a virtual router, multiple switches (routers) are assigned to a common GLBP group. Rather than having just one active router performing forwarding for the virtual router address, all routers in the group can participate and offer load balancing by forwarding a portion of the overall traffic.
2. VRRP
The Virtual Router Redundancy Protocol (VRRP) is a standards-based alternative to HSRP, defined in IETF standard RFC 2338. VRRP is so similar to HSRP that you need to learn only slightly different terminology and a couple of slight functional differences.
1. VRRP provides one redundant gateway address from a group of routers. The active router is called the master router, while all others are in the backup state. The master router is the one with the highest router priority in the VRRP group.
2. VRRP group numbers range from 0 to 255; router priorities range from 1 to 254 (254 is the highest; 100 is the default).
3. The virtual router MAC address is of the form 0000.5e00.01xx, where xx is a two-digit hex VRRP group number.
4. VRRP advertisements are sent at 1-second intervals. Backup routers can optionally learn the advertisement interval from the master router.
5. By default, all VRRP routers are configured to preempt the current master router, if their priorities are greater.
6. VRRP has no mechanism for tracking interfaces to allow more capable routers to take over the master role.
3. HSRP
HSRP is a Cisco-proprietary protocol developed to allow several routers (or multilayer switches) to appear as a single gateway address. RFC 2281 describes this protocol in more detail. Basically, each of the routers that provides redundancy for a given gateway address is assigned to a common HSRP group. One router is elected as the primary, or active, HSRP router, another is elected as the standby HSRP router, and all the others remain in the listen HSRP state. The routers exchange HSRP hello messages at regular

intervals, so they can remain aware of each other's existence, as well as that of the active router.

QUESTION 269
You need to implement a high availability design in the Certkiller routed network. Which protocol allows for the automatic selection and simultaneous use of multiple available gateways as well as automatic failover between those gateways?

A. VRRP
B. GLBP
C. IRDP
D. HSRP
E. None of the above

Answer: B

Explanation:
To provide a virtual router, multiple switches (routers) are assigned to a common GLBP group. Rather than having just one active router performing forwarding for the virtual router address, all routers in the group can participate and offer load balancing by forwarding a portion of the overall traffic. The advantage is that none of the clients have to be pointed toward a specific gateway address-they can all have the same default gateway set to the virtual router IP address. The load balancing is provided completely through the use of virtual router MAC addresses in ARP replies returned to the clients. As a client sends an ARP request looking for the virtual router address, GLBP sends back an ARP reply with the virtual MAC address of a selected router in the group. The result is that all clients use the same gateway address but have differing MAC addresses for it.

QUESTION 270
CK1 and CK2 exchange HSRP between each other in the Certkiller network. In which three HSRP states do routers send hello messages? (Select three)

A. Learn
B. Speak
C. Standby
D. Listen
E. Active
F. Remove

Answer: B, C, E

Explanation:
When HSRP is configured on an interface, the router progresses through a series of states before becoming active. This forces a router to listen for others in a group and see where it fits into the pecking order. The HSRP state sequence is Disabled, Init, Listen, Speak, Standby, and, finally, Active.

Only the standby (second highest priority) router monitors the hello messages from the active router. By default, hellos are sent every 3 seconds. If hellos are missed for the duration of the holdtime timer (default 10 seconds, or 3 times the hello timer), the active router is presumed down. The standby router is then clear to assume the active role. If other routers are sitting in the Listen state, the next-highest priority router is allowed to become the new standby router.

QUESTION 271
HSRP is being set up between two Certkiller devices. In what three states is it possible for an HSRP router to be in? (Select three)

A. Standby
B. Established
C. Active
D. Idle
E. Backup
F. Init

Answer: A, C, F

Explanation:
With HSRP, a set of routers work together to present the illusion of a single virtual router to the hosts on the LAN. This set is known as an HSRP group or a standby group. A single router elected from the group is responsible for forwarding the packets that hosts send to the virtual router. This router is known as the Active router. Another router is elected as the Standby router. In the event that the Active router fails, the Standby assumes the packet-forwarding duties of the Active router. Although an arbitrary number of routers may run HSRP, only the Active router forwards the packets sent to the virtual router. Before a router becomes the active or standby router, it will be in the Init (initial) state.
Reference:
http://www.cisco.com/en/US/tech/ CK6 48/ CK3 62/technologies_tech_note09186a0080094a91.shtml

---

QUESTION 272
To protect against first-hop router failure, four protocols were developed to ensure IP routing redundancy. Which of the following are they? (Select four)

A. HSRP
B. IRDP
C. ICMP
D. VRRP
E. MSTP
F. GLBP

Answer: A, B, D, F

Explanation:
A: HSRP is the Hot Standby Routing Protocol, which is the Cisco proprietary method for automatic failover and provides for redundant default gateways for hosts.
B: Some newer IP hosts use ICMP Router Discovery Protocol (IRDP) (RFC 1256) to find a new router when a route becomes unavailable. A host that runs IRDP listens for hello multicast messages from its configured router and uses an alternate router when it no longer receives those hello messages.
D: VRRP is the Virtual Router Redundancy Protocol, which is similar in many ways to HSRP. One key difference is that VRRP is standards based, where HSRP is Cisco developed.
F: Gateway Load Balancing Protocol (GLBP) protects data traffic from a failed router or circuit, like Hot Standby Router Protocol (HSRP) and Virtual Router Redundancy Protocol (VRRP), while allowing packet load sharing between a group of redundant routers.

QUESTION 273
HSRP has been configured between two Certkiller devices. What kind of message does an HSRP configured router send out every 3 seconds? (Select all that apply)

A. Retire
B. Coup
C. Resign
D. Send
E. Hello

Answer: E

Explanation:
Hello-The hello message conveys to other HSRP routers the router's HSRP priority and state information. By default, an HSRP router sends hello messages every three seconds.
Incorrect Answers:
A, D: These messages are not used by HSRP.
B: Coup-When a standby router assumes the function of the active router, it sends a coup message. This message is used by HSRP, but it is not sent our every 3 seconds.
C: Resign-A router that is the active router sends this message when it is about to shut down or when a router that has a higher priority sends a hello message. This message is only sent before it resigns, not every 3 seconds.

QUESTION 274
HSRP has been configured between two Certkiller devices. Which of the following describe reasons for deploying HSRP? (Select all that apply)

A. HSRP provides redundancy and fault tolerance
B. HSRP allows one router to automatically assume the function of the second router if the second router fails
C. HSRP allows one router to automatically assume the function of the second router if

the second router starts
D. HSRP provides redundancy and load balancing

Answer: A, B, D

Explanation:
One way to achieve near-100 percent network uptime is to use HSRP, which provides network redundancy for IP networks, ensuring that user traffic immediately and transparently recovers from first hop failures in network edge devices or access circuits. By sharing an IP address and a MAC (Layer 2) address, two or more routers can act as a single "virtual" router. The members of the virtual router group continually exchange status messages. This way, one router can assume the routing responsibility of another, should it go out of commission for either planned or unplanned reasons. Hosts continue to forward IP packets to a consistent IP and MAC address, and the changeover of devices doing the routing is transparent.
Through the use of multiple HSRP standby groups, traffic can be load balanced between the HSRP routers. For example, users on one VLAN could use one router as the primary HSRP router, and users on another VLAN can use the other HSRP router as the primary.

QUESTION 275
Which one of the statements below correctly describes the Virtual Router Redundancy Protocol (VRRP)?

A. A VRRP group has one active and one or more standby virtual routers.
B. A VRRP group has one master and one or more backup virtual routers.
C. A VRRP group has one active and one or more standby virtual routers.
D. A VRRP group has one master and one redundant virtual router.

Answer: B

Explanation:
The Virtual Router Redundancy Protocol (VRRP) feature can solve the static configuration problem. VRRP enables a group of routers to form a single virtual router. The LAN clients can then be configured with the virtual router as their default gateway. The virtual router, representing a group of routers, is also known as a VRRP group.
In a topology where multiple virtual routers are configured on a router interface, the interface can act as a master for one virtual router and as a backup for one or more virtual routers.
Reference:
http://www.cisco.com/en/US/products/sw/iosswrel/ps1612/products_feature_guide09186a0080080a60.html

---

QUESTION 276
Which type of scheme describes the default operation of Gateway Load Balancing Protocol (GLBP)?

A. per host using a round robin scheme

B. per host using a strict priority scheme
C. per session using a round robin scheme
D. per session using a strict priority scheme
E. per GLBP group using a round robin scheme
F. per GLBP group using a strict priority scheme

Answer: A

Explanation:
The Gateway Load Balancing Protocol feature provides automatic router backup for IP hosts configured with a single default gateway on an IEEE 802.3 LAN. Multiple first hop routers on the LAN combine to offer a single virtual first hop IP router while sharing the IP packet forwarding load. Other routers on the LAN may act as redundant GLBP routers that will become active if any of the existing forwarding routers fail.
GLBP performs a similar, but not identical, function for the user as the HSRP and the VRRP. HSRP and VRRP protocols allow multiple routers to participate in a virtual router group configured with a virtual IP address. One member is elected to be the active router to forward packets sent to the virtual IP address for the group. The other routers in the group are redundant until the active router fails. These standby routers have unused bandwidth that the protocol is not using. Although multiple virtual router groups can be configured for the same set of routers, the hosts must be configured for different default gateways, which results in an extra administrative burden. GLBP provides load balancing over multiple routers (gateways) using a single virtual IP address and multiple virtual MAC addresses. Each host is configured with the same virtual IP address, and all routers in the virtual router group participate in forwarding packets. In this way, per host load balancing is achieved using a round robin mechanism.

QUESTION 277
DRAG DROP
Match the HSRP states on the left with the correct definition on the right.

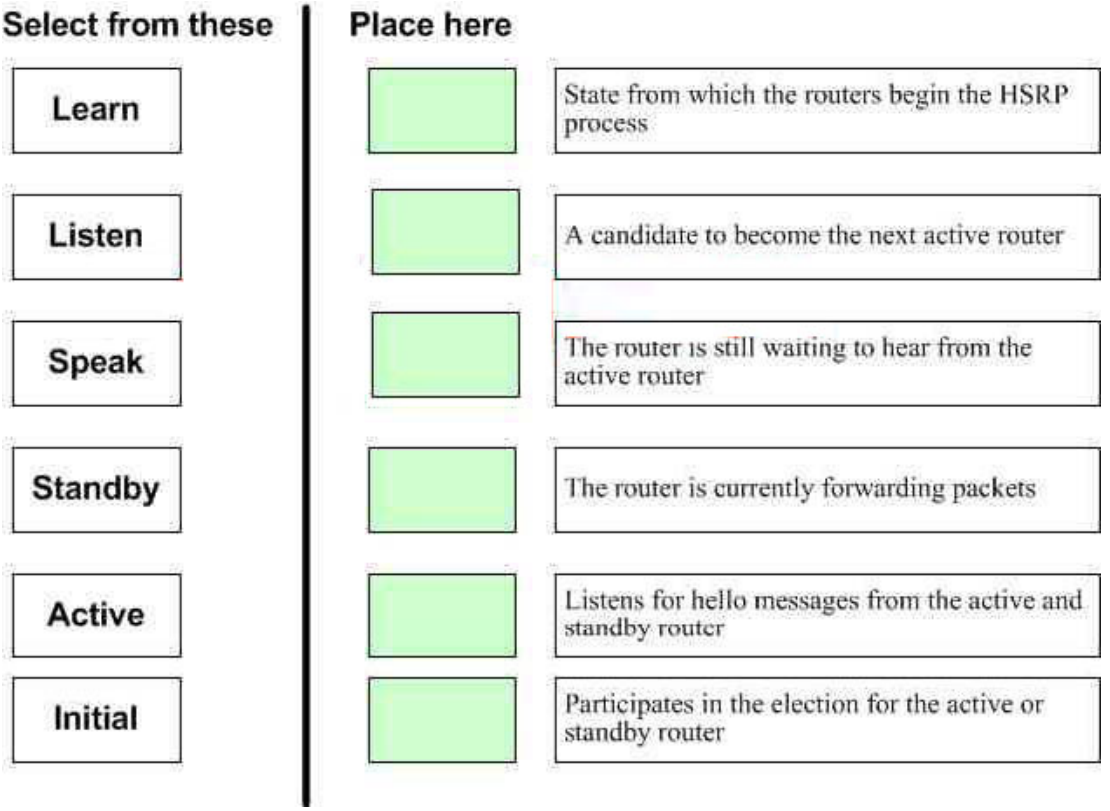

Answer:

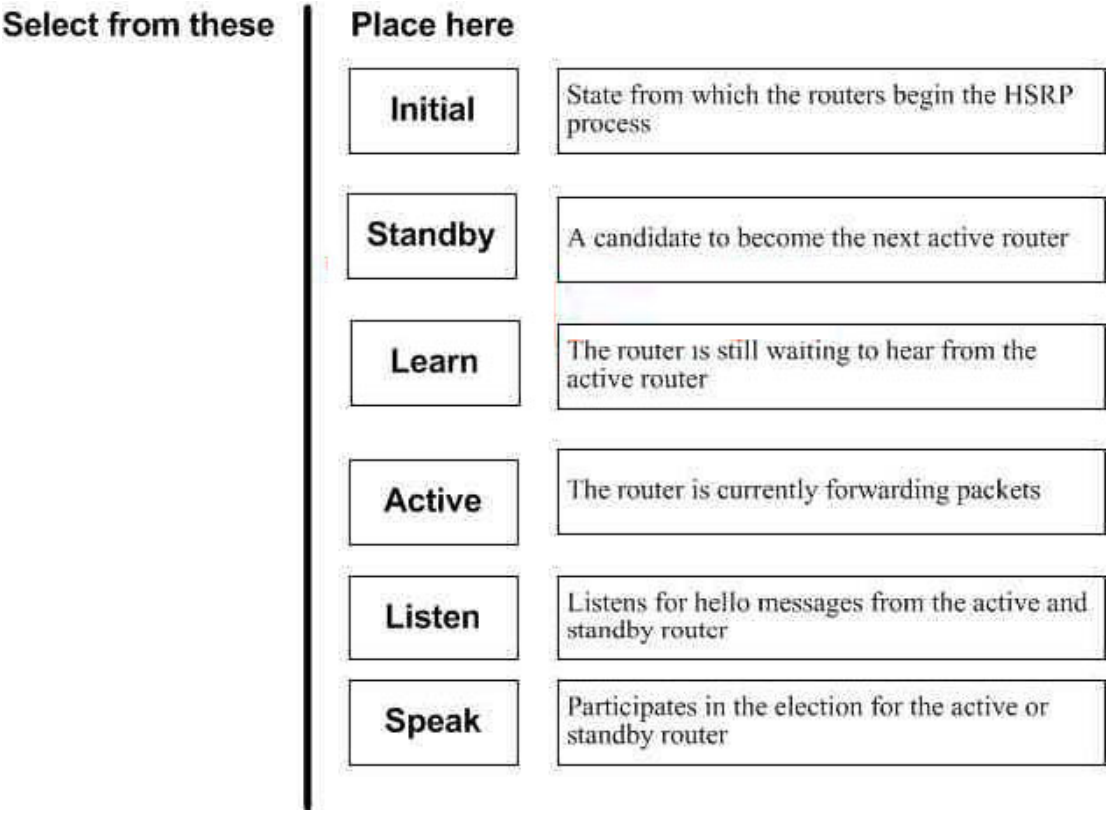

Explanation:
HSRP defines six states in which an HSRP-enabled router can exist:

1. Initial - This is the state from which the routers begin the HSRP process. This state indicates that HSRP is not running. It is entered via a configuration change or when an interface first comes up.
2. Learn - The router has not determined the virtual IP address, and has not yet seen an authenticated hello message from the active router. In this state the router is still waiting to hear from the active router.
3. Listen - The router knows the virtual IP address, but is neither the active router nor the standby router. It listens for hello messages from those routers. Routers other than the active and standby router remain in the listen state.
4. Speak - The router sends periodic hello messages and is actively participating in the election of the active or standby router. A router cannot enter Speak state unless it has the virtual IP address.
5. Standby - The router is a candidate to become the next active router and sends periodic hello messages. Excluding transient conditions, there must be at most one router in the group in Standby state.
6. Active - The router is currently forwarding packets that are sent to the group virtual MAC address. The router sends periodic hello messages. Excluding transient conditions, there must be at most one router in Active state in the HSRP group.

QUESTION 278
In the hardware address 0000.0c07.ac0av, what does 07.ac represent?

A. HSRP well-known physical MAC address
B. Vendor code
C. HSRP router number
D. HSRP group number
E. HSRP well-known virtual MAC address

Answer: E

Explanation:
HSRP code (HSRP well-known virtual MAC address) - The fact that the MAC address is for an HSRP virtual router is indicated in the next two bytes of the address. The HSRP code is always 07.ac. The HSRP protocol uses a virtual MAC address, which always contains the 07.ac numerical value.
Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 268

QUESTION 279
The Certkiller network needs to enhance the reliability of the network. Which of the following protocols provides network redundancy for IP networks, ensuring that user traffic immediately and transparently recovers from first-hop failures in network edge devices or access circuits, as defined by RFC 2281?

A. STP
B. IRDP
C. ICMP

D. HSRP
E. None of the above

Answer: D

Explanation:
HSRP is defined in RFC 2281. The Hot Standby Router Protocol, HSRP, provides a mechanism which is designed to support non-disruptive failover of IP traffic in certain circumstances. In particular, the protocol protects against the failure of the first hop router when the source host cannot learn the IP address of the first hop router dynamically. The protocol is designed for use over multi-access, multicast or broadcast capable LANs (e.g., Ethernet). HSRP is not intended as a replacement for existing dynamic router discovery mechanisms and those protocols should be used instead whenever possible. A large class of legacy host implementations that do not support dynamic discovery are capable of configuring a default router. HSRP provides failover services to those hosts.
Reference: http://www.faqs.org/rfcs/rfc2281.html

QUESTION 280
Which of the following protocols enables a group of routers to form a single virtual router, and then use the real IP address of a router as the gateway address, as defined in RFC 2338?

A. Proxy ARP
B. HSRP
C. IRDP
D. VRRP
E. GLBP

Answer: D

Explanation:
The Virtual Router Redundancy Protocol (VRRP) feature enables a group of routers to form a single virtual router. The LAN clients can then be configured with the virtual router as their default gateway. The virtual router, representing a group of routers, is also known as a VRRP group. VRRP is defined in RFC 2338.
Reference: http://www.faqs.org/rfcs/rfc2338.html

QUESTION 281
Routers can send hello messages in three HSRP states. Which ones are they? (Select three)

A. standby
B. learn
C. listen
D. speak

E. active

Answer: A, D, E

Explanation:
The various HSRP states are described below:
Listen: The router knows the virtual IP address, but is neither the active router nor the standby router. It listens for hello messages from those routers.
Speak: The router sends periodic hello messages, and is actively participating in the election of the active and/or standby router. A router cannot enter speak state unless it has the virtual IP address.
Standby: The router is a candidate to become the next active router, and sends periodic hello messages. Excluding transient conditions, there would be at most one router in the group in standby state.
Active: The router is currently forwarding packets that are sent to the group's virtual MAC address. The router sends periodic hello messages. Excluding transient conditions, there must be at most one router in active state in the group.
Initial: This is the starting state, and indicates that HSRP is not running. This state is entered via a configuration change, or when an interface first comes up.
Learn: The router has not determined the virtual IP address, and has not yet seen an authenticated hello message from the active router. In this state, the router is still waiting to hear from the active router.

QUESTION 282
Two Certkiller routers are configured for HSRP. Cisco's Hot Standby Routing Protocol (HSRP) can provide automatic router backup over which networks?

A. Ethernet and FDDI
B. Ethernet, FDDI and Token Ring LANs
C. Token Ring LANs only
D. VINES and IPX only
E. Ethernet and Token Ring LANs

Answer: B

Explanation:
Cisco's Hot Standby Routing Protocol (HSRP) provides automatic router backup when you configure it on Cisco routers that run the Internet Protocol (IP) over Ethernet, Fiber Distributed Date Interface (FDDI), and Token Ring local-area networks (LANs). HSRP is compatible with Novell's Internetwork Packet Exchange (IPX), AppleTalk, and Banyan VINES, and it is compatible with DECnet and Xerox Network Systems (XNS) in certain configurations.

QUESTION 283
Which router redundancy protocol cannot be configured for interface tracking?

A. HSRP
B. GLBP
C. VRRP
D. SLB
E. RPR
F. RPR+

Answer: C

Explanation:
The Virtual Router Redundancy Protocol (VRRP) is a standards-based alternative to HSRP, defined in IETF standard RFC 2338. VRRP is so similar to HSRP that you need to learn only slightly different terminology and a couple of slight functional differences.
1. VRRP provides one redundant gateway address from a group of routers. The active router is called the master router, while all others are in the backup state. The master router is the one with the highest router priority in the VRRP group.
2. VRRP group numbers range from 0 to 255; router priorities range from 1 to 254 (254 is the highest; 100 is the default).
3. The virtual router MAC address is of the form 0000.5e00.01xx, where xx is a two-digit hex VRRP group number.
4. VRRP advertisements are sent at 1-second intervals. Backup routers can optionally learn the advertisement interval from the master router.
5. By default, all VRRP routers are configured to preempt the current master router, if their priorities are greater.
6. VRRP has no mechanism for tracking interfaces to allow more capable routers to take over the master role.

QUESTION 284
What protocol specified by RFC 1256 will allow an enabled IP host a new router when a router becomes unavailable?

A. IRDP
B. SNMP
C. HSRP
D. VRRP

Answer: A

Explanation:
Some newer IP hosts use ICMP Router Discovery Protocol (IRDP) (RFC 1256) to find a new router when a route becomes unavailable. A host that runs IRDP listens for hello multicast messages from its configured router and uses an alternate router when it no longer receives those hello messages. The default timer values of IRDP mean that it's not suitable for detection of failure of the first hop. The default advertisement rate is once every 7 to 10 minutes, and the default lifetime is 30 minutes.

QUESTION 285
The Certkiller network topology is shown in the following exhibit:

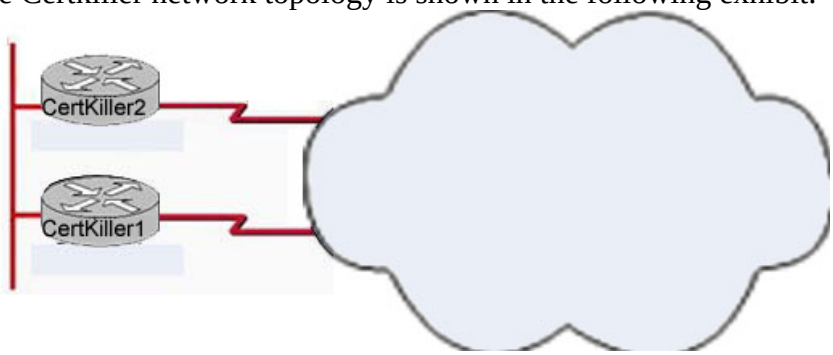

Configuration exhibit Certkiller 1:

```
Hostname CertKiller1
!
Interface fastethernet 0/0
ip address 10.10.10.2 255.255.255.0
standby 1 ip 10.10.10.10 255.255.255.0
standby 1 track serial 0/0
```

Configuration exhibit Certkiller 2:

```
Hostname CertKiller1
!
Interface fastethernet 0/0
ip address 10.10.10.1 255.255.255.0
standby 1 ip 10.10.10.10 255.255.255.0
standby 1 priority 105
standby 1 track serial 0/0
```

Which command will need to be added to Certkiller 1 to ensure that it will take over if serial 0/0 on Certkiller 2 fails?

A. standby 1 track 10.10.10.1
B. standby 1 preempt
C. standby 1 track fastethernet 0/0
D. standby 1 priority 130
E. None of the above

Answer: B

Explanation:
You can configure a router to preempt or immediately take over the active role if its priority is the highest at any time. Use the following interface configuration command to allow preemption:
Switch(config-if)# standby group preempt [delay seconds]
By default, the router can preempt another immediately, without delay. You can use the delay keyword to force it to wait for seconds before becoming active. This is usually done if there are routing protocols that need time to converge.

QUESTION 286
Study the below Certkiller exhibit carefully. The Gateway Load Balancing Protocol

has been configured on routers Certkiller 1 and Certkiller 2, and hosts Certkiller A and Certkiller B have been configured as shown. Which statement can be derived from the exhibit?

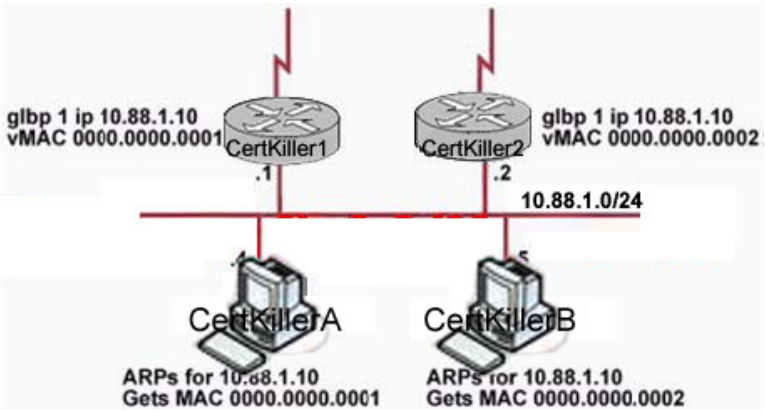

A. The GLBP host-dependent, load-balancing mode has been configured.
B. The host Certkiller A default gateway has been configured as 10.88.1.10/24.
C. The GLBP weighted load balancing mode has been configured.
D. The host Certkiller A default gateway has been configured as 10.88.1.4/24.
E. The host Certkiller A default gateway has been configured as 10.88.1.1/24.
F. The GLBP round-robin, load-balancing mode has been configured.
G. None of the above

Answer: B

Explanation:
To provide a virtual router, multiple switches (routers) are assigned to a common GLBP group. Rather than having just one active router performing forwarding for the virtual router address, all routers in the group can participate and offer load balancing by forwarding a portion of the overall traffic.
The advantage is that none of the clients have to be pointed toward a specific gateway address-they can all have the same default gateway set to the virtual router IP address. The load balancing is provided completely through the use of virtual router MAC addresses in ARP replies returned to the clients. As a client sends an ARP request looking for the virtual router address, GLBP sends back an ARP reply with the virtual MAC address of a selected router in the group. The result is that all clients use the same gateway address but have differing MAC addresses for it.

QUESTION 287
The Certkiller WAN is shown below:

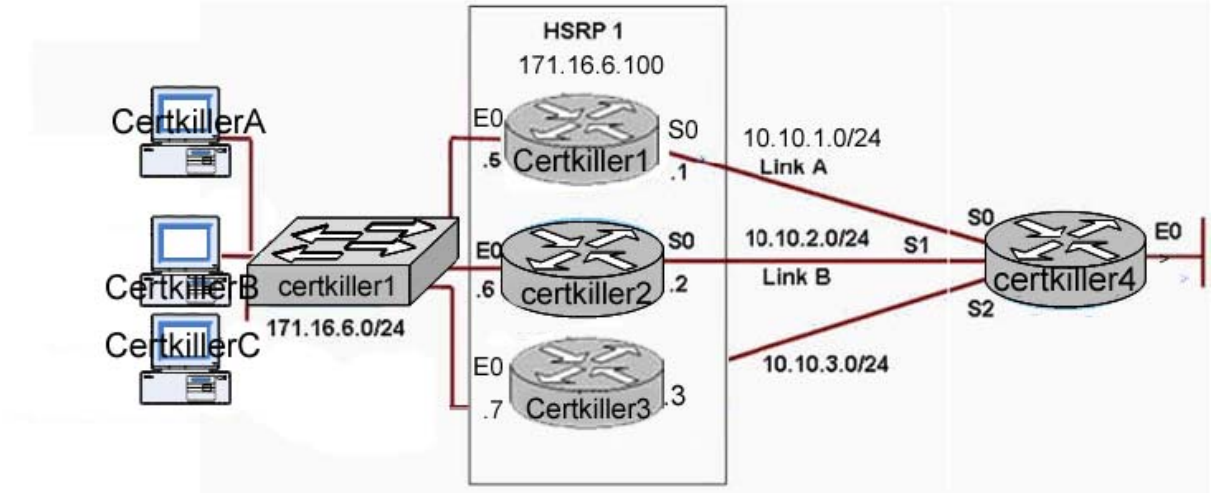

Certkiller 1 configuration exhibit:

```
Certkiller1 # show running-config
!
interface Ethernet0
 ip address 171.16.6.5 255.255.255.0
 standby 1 ip 171.16.6.100
 standby 1 priority 105
 standby 1 preempt
 Standby 1 track serial10 10

interface serial0
 ip address 10.10.1.1 255.255.255.0

 <Output omitted>
```

Certkiller 2 configuration exhibit:

```
Certkiller2# show running-config
!
interface Ethernet0
 ip address 171.16.6.6 255.255.255.0
 standby 1 ip 171.16.6.100
 standby 1 preempt
 standby 1 track serial10 10
!
interface serial9
 ip address 10.10.2.2 255.255.255.0
!
 <Output omitted>
```

Study the exhibits shown above. HSRP has been configured and Link A is the primary route to router Certkiller 4. When Link A fails, router Certkiller 2 (Link B) becomes the active router. Which router will assume the active role when Link A becomes operational again?

A. The standby router Certkiller 2 will remain active and will forward the active role to router Certkiller 1 only in the event of Link B failure.

B. The third member of the HSRP group, router Certkiller 3, will take over the active role only in event of router Certkiller 2 failure.
C. The primary router Certkiller 1 will reassume the active role when it comes back online.
D. The standby router Certkiller 2 will remain active and will forward the active role to router Certkiller 1 only in the event of its own failure.
E. None of the above.

Answer: C

Explanation:
HSRP election is based on a priority value (0 to 255) that is configured on each router in the group. By default, the priority is 100. The router with the highest priority value (255 is highest) becomes the active router for the group. If all router priorities are equal or set to the default value, the router with the highest IP address on the HSRP interface becomes the active router. To set the priority, use the following interface configuration command:
Switch(config-if)# standby group priority priority
When HSRP is configured on an interface, the router progresses through a series of states before becoming active. This forces a router to listen for others in a group and see where it fits into the pecking order. The HSRP state sequence is Disabled, Init, Listen, Speak, Standby, and, finally, Active.
Only the standby (second highest priority) router monitors the hello messages from the active router. By default, hellos are sent every 3 seconds. If hellos are missed for the duration of the holdtime timer (default 10 seconds, or 3 times the hello timer), the active router is presumed down. The standby router is then clear to assume the active role. If other routers are sitting in the Listen state, the next-highest priority router is allowed to become the new standby router.

QUESTION 288
What three tasks must a network administrator perform to properly configure Hot Standby Routing Protocol (HSRP)? (Select three)

A. Define the encapsulation type.
B. Define the standby router.
C. Define the standby IP address.
D. Enable the standby priority.

Answer: B, C, D

Explanation:
Three of the required configuration commands needed for enabling HSRP is to define the standby routing process, define the HSRP IP address, and configure the HSRP priority.
Configuring HSRP:
* Configuring an interface to participate in an HSRP standby group
* Assigning HSRP standby priority

* Configuring HSRP standby pre-empt
* Configuring HSRP over trunk links
* Configuring hello message timers
* HSRP interface tracking
* Displaying the status of HSRP

Incorrect Answers:
A: There are no encapsulation options for enabling HSRP.
Reference: Building Cisco Multilayer Switched Networks (Cisco Press) page 272

QUESTION 289
You want to allow Router CK1 to immediately become the active router if its priority is highest than the active router fails. What command would you use if you wanted to configure this?

A. en standby preempt
B. standby preempt enable
C. standby preempt
D. hot standby preempt

Answer: C

Explanation:
The HSRP preemption feature enables the router with highest priority to immediately become the Active router. Priority is determined first by the priority value that you configure, and then by the IP address. In each case a higher value is of greater priority. When a higher priority router preempts a lower priority router, it sends a coup message. When a lower priority active router receives a coup message or hello message from a higher priority active router, it changes to the speak state and sends a resign message. To configure preemption, use the "standby standby-number preempt" command.

QUESTION 290
The Certkiller network is using two routers with HSRP for their Internet access as shown below:

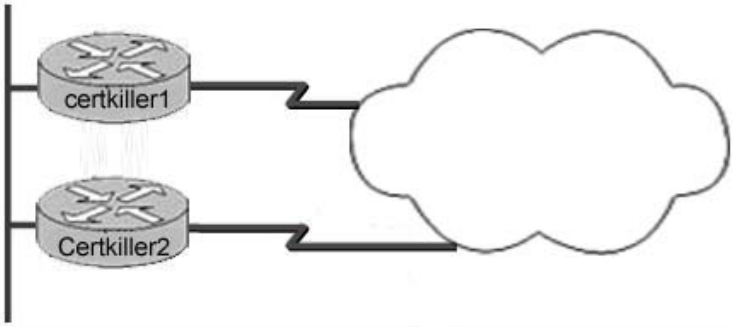


```
Hostname  Certkiller1
!
Interface fastethernet 0/0
ip address 10.10.10.1 255.255.255.0
standby 1 ip 10.10.10.10 255.255.255.0
standby 1 priority 105
standby 1 track serial 0/0
```

```
Hostname Certkiller2
!
Interface fastethernet 0/0
ip address 10.10.10.2 255.255.255.0
standby 1 ip 10.10.10.10 255.255.255.0
standby 1 track serial 0/0
```

Which command will need to be added to Certkiller 2 to ensure that it will take over if serial 0/0 on Certkiller 1 fails?

A. standby 1 preempt
B. standby 1 track 10.10.10.1
C. standby 1 priority 130
D. standby 1 track fastethernet 0/0
E. None of the above

Answer: A

Explanation:
When this command is configured, the router is configured to preempt, which means that when the local router has a Hot Standby priority higher than the current active router, the local router should attempt to assume control as the active router. If the hsrp preempt command is not configured, the local router assumes control as the active router only if it receives information indicating that no router currently is in the active state (acting as the designated router).
In this example, Certkiller 1 was properly configured to lower its own HSRP priority when the serial 0 interface goes down. However, even if this happens, router Certkiller 2 will not become the active router unless it is configured to pre-empt, or take over, if it suddenly has a higher priority than Certkiller 1.

QUESTION 291
On a 3550 EMI switch, which three types of interfaces can be used to configure HSRP? (Select three)

A. Loopback interface
B. SVI interface
C. Routed port
D. Access port

E. EtherChannel port channel
F. BVI interface

Answer: B, C, E

Explanation:
This Hot Standby Router Protocol (HSRP) provides routing redundancy for routing IP traffic without being dependent on the availability of any single router. To use this feature, you must have the enhanced multilayer software image installed on your switch. All Catalyst 3550 Gigabit Ethernet switches ship with the enhanced multilayer software image (EMI) installed. Catalyst 3550 Fast Ethernet switches can be shipped with either the standard multilayer software image (SMI) or EMI pre-installed. You can order the Enhanced Multilayer Software Image Upgrade kit to upgrade Catalyst 3550 Fast Ethernet switches from the SMI to the EMI.
Only routed interfaces that provide access to hosts can be configured for HSRP. These interfaces include: routed Ethernet, routed fast Ethernet, routed Gigabit Ethernet, SVI, and EtherChannel.
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps646/products_configuration_guide_chapter09186a00800c9

QUESTION 292
Exhibit:

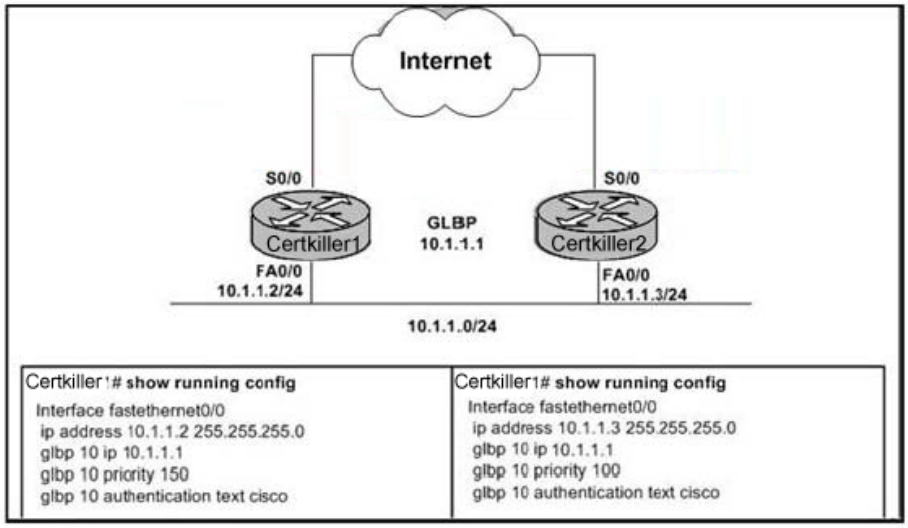

Refer to the exhibit. Which GLBP device hosts the virtual MAC addresses?

A. Certkiller 1
B. Certkiller 2
C. AVG
D. AVF

Answer: D

Explanation:
Active Virtual Gateway
The trick behind this load balancing lies in the GLBP group. One router is elected the active virtual
gateway (AVG). This router has the highest priority value, or the highest IP address in the group, if
there is no highest priority. The AVG answers all ARP requests for the virtual router address. Which
MAC address it returns depends upon which load-balancing algorithm it is configured to use. In any
event, the virtual MAC address supported by one of the routers in the group is returned.
The AVG also assigns the necessary virtual MAC addresses to each of the routers participating in
the GLBP group. Up to four virtual MAC addresses can be used in any group. Each of these routers
is referred to as an active virtual forwarder (AVF), forwarding traffic received on its virtual MAC
address. Other routers in the group serve as backup or secondary virtual forwarders, in case the AVF
fails. The AVG also assigns secondary roles.
Assign the GLBP priority to a router with the following interface configuration command:
Switch(config-if)# glbp group priority level
GLBP group numbers range from 0 to 1023. The router priority can be 1 to 255 (255 is the highest
priority), defaulting to 100.
Active Virtual Forwarder
GLBP uses a weighting function to determine which router becomes the AVF for a virtual MAC
address in a group. Each router begins with a maximum weight value (1 to 254). As specific interfaces
go down, the weight is decreased by a configured amount. GLBP uses thresholds to determine
when a router can and cannot be the AVF. If the weight falls below the lower threshold, the router
must give up its AVF role. When the weight rises above the upper threshold, the router can resume
its AVF role.
By default, a router receives a maximum weight of 100. If you want to make a dynamic weighting
adjustment, GLBP must know which interfaces to track and how to adjust the weight. You must first
define an interface as a tracked object with the following global configuration command:
Switch(config)# track object-number interface type mod/num {line-protocol | ip routing}

QUESTION 293
The following output was seen on switch Certkiller A:

Based on the output shown above, which two statements are true about the output from the show "standby vlan 50" command? (Select two)

A. Hosts using the default gateway address of 192.168.1.1 will have their traffic sent to 192.168.1.11 even after Certkiller A becomes available again.
B. Hosts using the default gateway address of 192.168.1.2 will have their traffic sent to Certkiller A.
C. The command standby 1 preempt was added to Certkiller A.
D. Certkiller A is load sharing traffic in VLAN 50.

Answer: C, D

Explanation:
HSRP uses a priority scheme to determine which HSRP-configured router is to be the default active router. To configure a router as the active router, you assign it a priority that is higher than the priority of all the other HSRP-configured routers. The default priority is 100, so if you configure just one router to have a higher priority, that router will be the default active router.
HSRP works by the exchange of multicast messages that advertise priority among HSRP-configured routers. When the active router fails to send a hello message within a configurable period of time, the standby router with the highest priority becomes the active router. The transition of packetforwarding
functions between routers is completely transparent to all hosts on the network.
HSRP-configured routers exchange three types of multicast messages:
Hello-The hello message conveys to other HSRP routers the router's HSRP priority and state information. By default, an HSRP router sends hello messages every three seconds.
Coup-When a standby router assumes the function of the active router, it sends a coup message.
Resign-A router that is the active router sends this message when it is about to shut down or when a router that has a higher priority sends a hello message.
At any time, HSRP-configured routers are in one of the following states:

Active-The router is performing packet-transfer functions.
Standby-The router is prepared to assume packet-transfer functions if the active router fails.
Speaking and listening-The router is sending and receiving hello messages.
Listening-The router is receiving hello messages.
The standby preempt interface configuration command allows the router to become the active router when its priority is higher than all other HSRP-configured routers in this Hot Standby group. The configurations of both routers include this command so that each router can be the standby router for the other router. The 1 indicates that this command applies to Hot Standby group 1. If you do not use the standby preempt command in the configuration for a router, that router cannot become the active router.

QUESTION 294
The Certkiller WAN is shown below:

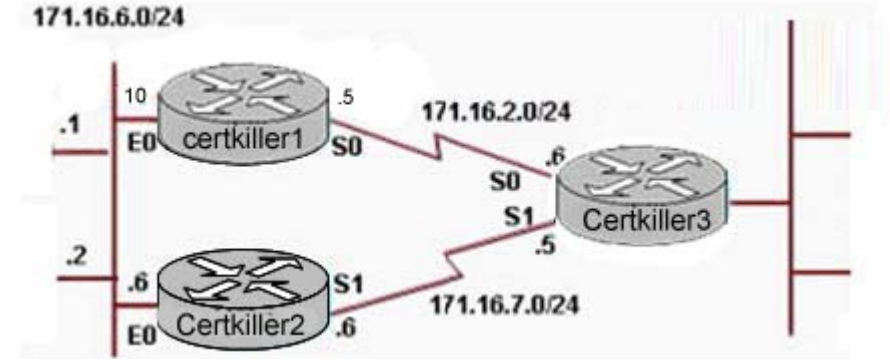

Certkiller 1 configuration exhibit:

```
Certkiller1#show run
!
interface Ethernet0
ip address 171.16.6.5 255.255.255.0
no ip redirects
standby 1 ip 171.16.6.100
standby 1 priority 115
standby 1 preempt
standby 1 track serial0
```

Certkiller 2 configuration exhibit:

```
Certkiller2#show run
!
interface Ethernet0
ip address 171.16.6.6 255.255.255.0
no ip redirects
standby 1 ip
standby 1 preempt
standby 1 track serial1
```

Refer to the exhibit and the partial configuration on routers Certkiller 1 and Certkiller 2. Hot Standby Routing Protocol (HSRP) is configured on the network to provide network redundancy for the IP traffic. The network administrator noticed that Certkiller 2 does not became active when the Certkiller 1 serial0 interface goes down. What should be changed in the configuration to fix the problem?

A. Certkiller 2 should be configured with a HSRP virtual address.
B. Certkiller 2 should be configured with a standby priority of 100.

C. The Serial0 interface on router Certkiller 1 should be configured with a decrement value of 20.
D. The Serial1 interface on router Certkiller 2 should be configured with a decrement value of 20.
E. None of the above.

Answer: C

Explanation:
You can configure a router to preempt or immediately take over the active role if its priority is the highest at any time. Use the following interface configuration command to allow preemption:
Switch(config-if)# standby group preempt [delay seconds]
By default, the router can preempt another immediately, without delay. You can use the delay keyword to force it to wait for seconds before becoming active. This is usually done if there are routing protocols that need time to converge.

QUESTION 295
The following debug output was seen on a Certkiller device:

```
*Mar  1 00:16:43.095: %LINK-3-UPDOWN: Interface Vlan11, changed state to up
*Mar  1 00:16:43.099: SB: Vl11 Interface up
*Mar  1 00:16:43.099: SB11: Vl11 Init: a/HSRP enabled
*Mar  1 00:16:43.099: SB11: Vl11 Init -> Listen
*Mar  1 00:16:43.295: SB11: Vl11 Hello in 172.16.11.112 Active pri 50 ip 172.16.11.115
*Mar  1 00:16:43.295: SB11: Vl11 Active router is 172.16.11.112
*Mar  1 00:16:43.295: SB11: Vl11 Listen: h/Hello rcvd from lower pri Active router (50/172.16.11.112)
*Mar  1 00:16:43.295: SB11: Vl11 Active router is local, was 172.16.11.112
*Mar  1 00:16:43.299: %STANDBY-6-STATECHANGE: Vlan11 Group 11 state Listen -> Active
*Mar  1 00:16:43.295: SB11: Vl11 Active router is 11.111 Active pri 100 ip 172.16.11.115
*Mar  1 00:16:43.299: %STANDBY-6-STATECHANGE: Vlan1.112 Speak pri 50 ip 172.16.11.115
*Mar  1 00:16:43.299: SB11: Vl11 Hello out 172.16.11.111 Active pri 100 ip 172.16.11.115
*Mar  1 00:16:43.299: %STANDBY-6-STATECHANGE: Vlan 1.112 Speak pri 50 ip 172.16.11.115
```

Based upon the debug output that is shown, which three statements about HSRP are true? (Select three)

A. The priority of the router with IP address 172.16.11.112 is preferred over the router with IP address 172.16.11.111.
B. The IP address 172.16.11.115 is the virtual HSRP IP address.
C. The router with IP address 172.16.11.112 has nonpreempt configured.
D. The router with IP address 172.16.11.112 is using default HSRP priority.
E. The router with IP address 172.16.11.111 has preempt configured.
F. The final active router is the router with IP address 172.16.11.111.

Answer: B, E, F

Explanation:
Each router in an HSRP group has its own unique IP address assigned to an interface. This address is used for all routing protocol and management traffic initiated by or destined to the router. In addition, each router has a common gateway IP address, the

virtual router address, that is kept alive by HSRP. This address is also referred to as the HSRP address or the standby address. Clients can point to that virtual router address as their default gateway, knowing that a router always keeps that address active. Keep in mind that the actual interface address and the virtual (standby) address must be configured to be in the same IP subnet. You can assign the HSRP address with the following interface command:
Switch(config-if)# standby group ip ip-address [secondary]
When HSRP is used on an interface that has secondary IP addresses, you can add the secondarykeyword so that HSRP can provide a redundant secondary gateway address.
You can configure a router to preempt or immediately take over the active role if its priority is the highest at any time. Use the following interface configuration command to allow preemption:
Switch(config-if)# standby group preempt [delay seconds]
By default, the router can preempt another immediately, without delay. You can use the delay keyword to force it to wait for seconds before becoming active. This is usually done if there are routing protocols that need time to converge.

QUESTION 296
The following debug output was seen on a Certkiller device:

```
*Mar  1 00:12:16.871: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip 172.16.11.115
*Mar  1 00:12:16.871: SB11: Vl11 Active router is 172.16.11.112
*Mar 1.00:12:18.619%Link-3-UPDOWN: Interface Vlan11, changed state to up
*Mar 1.00:12:18.623:SB.VIII  Interface up
*Mar  1 00:12:18.623: SB11: Vl11 Init: a/HSRP enabled
*Mar  1 00:12:18.623: SB11: Vl11 Init -> Listen
*Mar  1 00:12:19.619: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan11, changed state to up
*Mar  1 00:12:19.819: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip 172.16.11.115
*Mar  1 00:12:19.819: SB11: Vl11 Listen: h/Hello rcvd from lower pri Active router (50/172.16.11.112)
*Mar  1 00:12:22.815: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip 172.16.11.115
*Mar  1 00:12:22.815: SB11: Vl11 Listen: h/Hello rcvd from lower pri Active router
*Mar  1 00:12:25.683: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip 172.16.11.115
*Mar  1 00:12:25.683: SB11: Vl11 Listen: h/Hello rcvd from lower pri Active router (50/172.16.11.112)
*Mar  1 00:12:28.623: SB11: Vl11 Listen: d/Standby timer expired (unknown)
*Mar  1 00:12:28.623: SB11: Vl11 Listen -> Speak
*Mar  1 00:12:28.623: SB11: Vl11 Hello  out 172.16.11.111 Speak  pri 100 ip 172.16.11.115
*Mar 1.00:12.28.659: SB11:VIII Hello  in  172.16.11.112 Active  pri 50 ip 172.16.11.115
*Mar 1.00:12.28.659: SB11:VIII Speak h/Hello rcvd from lower pri Active router (50/172.16.11.112)
*Mar  1 00:12:31.539: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip 172.16.11.115
*Mar  1 00:12:31.539: SB11: Vl11 Speak h/Hello rcvd from lower pri Active router (50/172.16.11.112)
*Mar  1 00:12:31.575: SB11: Vl11 Hello  out 172.16.11.111 Speak  pri 100 ip 172.16.11.115
*Mar  1 00:12:34.491: SB11: Vl11 Hello  in  172.16.11.112 Active  pri 50 ip 172.16.11.115
```

Based on the information shown above, what can be determined about the HSRP relationship from the displayed debug output?

A. Router 172.16.11.112 will be the active router because its HSRP priority is preferred over router 172.16.11.111.
B. The IP address 172.16.11.112 is the virtual HSRP router IP address.
C. The preempt feature is not enabled on the 172.16.11.111 router.
D. Router 172.16.11.111 will be the active router because its HSRP priority is preferred over router 172.16.11.112.
E. The nonpreempt feature is enabled on the 172.16.11.112 router.
F. The IP address 172.16.11.111 is the virtual HSRP router IP address.

G. None of the above

Answer: C

Explanation:
The standby preempt interface configuration command allows the router to become the active router when its priority is higher than all other HSRP-configured routers in this Hot Standby group. The configurations of both routers include this command so that each router can be the standby router for the other router. The 1 indicates that this command applies to Hot Standby group 1. If you do not use the standby preempt command in the configuration for a router, that router cannot become the active router.

QUESTION 297
Certkiller 1 configuration exhibit:

```
Certkiller 1# show running-config interface GigabitEthernet 0/0

content configuration:205 bytes
!
Interface GigabiteEthernet
ip address 10.1.7.5 255.255.255.0
duple autp
Speed auto
media-type rj45
negotiation auto
glbp 7 ip 10.1.7.1
glbp 7 timers msec 250 msec 750
glbp 7 priority 150
```

Certkiller 1 configuration exhibit:

```
Certkiller1#show running-config interface GigabitEthernet 0/0

current configuration: 174 bytes
!
interface Serial Ethernet/0
ip address 10.1.7.6 255.255 255.0
duplex auto
speed auto
media-type rj45
negotiation auto
glbp 7 timers msec 250 msec 750
```

Study the exhibits carefully. What statement is true based upon the configuration of router Certkiller 1 and router Certkiller 2?

A. Router Certkiller 2 will become the active virtual gateway.
B. Router Certkiller 1 will become the active virtual gateway.
C. The hello and hold timers are incompatible with OSPF type 5 LSAs.
D. The hello and hold timers are incompatible with multi-homed BGP.
E. Router Certkiller 1 will become the master for Virtual Router 1, and router Certkiller 2 will become the backup for Virtual Router 2.
F. Router Certkiller 2 will become the master for Virtual Router 1, and router Certkiller 1 will become the backup for Virtual Router 2.
G. None of the above

Answer: B

Explanation:
HSRP election is based on a priority value (0 to 255) that is configured on each router in the group. By default, the priority is 100. The router with the highest priority value (255 is highest) becomes the active router for the group. If all router priorities are equal or set to the default value, the router with the highest IP address on the HSRP interface becomes the active router. To set the priority, use the following interface configuration command:
Switch(config-if)# standby group priority priority
When HSRP is configured on an interface, the router progresses through a series of states before becoming active. This forces a router to listen for others in a group and see where it fits into the pecking order. The HSRP state sequence is Disabled, Init, Listen, Speak, Standby, and, finally, Active.

QUESTION 298
The "show standby" command was placed on a Certkiller device as shown below:

```
Certkiller1# 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
    security virtual ip address 10.1.0.21
  Active virtual MAC address is 0004.4d82.7981
    Local virtual MAC address is 0004.4d82.7981 (bia)
  Hello time 4 sec, hold time 12 sec
    Next ip sent in 1.412 12 secs
  prempt run end min delay      50 sec, sync delay 40 sec
  Active router is local
  Standby router is 10.1.0.6, priority 75 (expires in 9.184 sec)
  Priority 95 (configured 120)
    Tracking 2 objects, 0 up
      Down Interface Ethernet0/2, pri 15
      Down Interface Ethernet0/3
  IP redundancy name is "HSRP1", advertisement interval is 34 sec
```

Examine the router output displayed in the exhibit above. Which two items are correct? (Select two)

A. When Ethernet 0/3 of Router Certkiller 1 comes back up, the priority will become 105.
B. The local IP address of Router Certkiller 1 is 10.1.0.20.
C. The local IP address of Router Certkiller 1 is 10.1.0.6.
D. Router Certkiller 1 will assume the active state if its priority is the highest.
E. If Ethernet 0/2 goes down, the standby router will take over.

Answer: A, D

Explanation:
HSRP election is based on a priority value (0 to 255) that is configured on each router in the group. By default, the priority is 100. The router with the highest priority value (255 is highest) becomes the active router for the group. If all router priorities are equal or set to the default value, the router with the highest IP address on the HSRP interface becomes the active router. To set the priority, use the following interface configuration command:

Switch(config-if)# standby group priority priority
When HSRP is configured on an interface, the router progresses through a series of states before becoming active. This forces a router to listen for others in a group and see where it fits into the pecking order. The HSRP state sequence is Disabled, Init, Listen, Speak, Standby, and, finally, Active.

QUESTION 299
While troubleshooting an issue in the Certkiller network, the following output was seen:

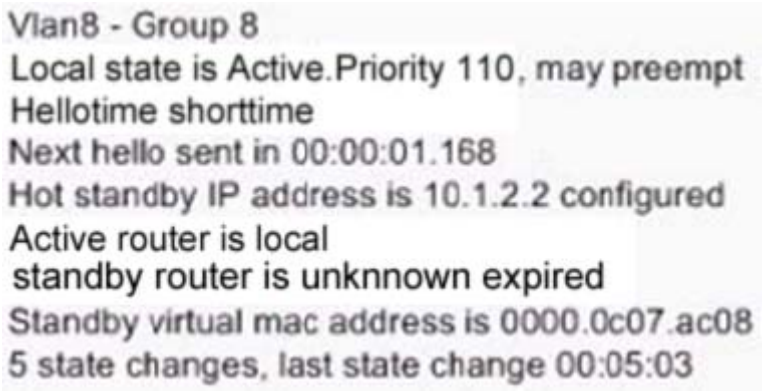

Based on the information shown above, which of the following would be the most likely cause of the exhibited output? (Select two)

A. Transport layer issues
B. HSRP misconfiguration
C. Physical layer issues
D. Spanning tree issues
E. VRRP misconfiguration
F. Application layer issues

Answer: B, C

Explanation:
Each router in an HSRP group has its own unique IP address assigned to an interface. This address is used for all routing protocol and management traffic initiated by or destined to the router. In addition, each router has a common gateway IP address, the virtual router address, that is kept alive by HSRP. This address is also referred to as the HSRP address or the standby address. Clients can point to that virtual router address as their default gateway, knowing that a router always keeps that address active. Keep in mind that the actual interface address and the virtual (standby) address must be configured to be in the same IP subnet. You can assign the HSRP address with the following interface command:
Switch(config-if)# standby group ip ip-address [secondary]
When HSRP is used on an interface that has secondary IP addresses, you can add the secondary keyword so that HSRP can provide a redundant secondary gateway address.

QUESTION 300
Exhibit

```
*Mar  1 00:16:43.095: %LINK-3-UPDOWN: Interface Vlan11, changed state to up
*Mar  1 00:16:43.099: SB:  VI11 Interface up
*Mar  1 00:16:43.099: SB11: VI11 Init: a/HSRP enabled
*Mar  1 00:16:43.099: SB11: VI11 Init -> Listen
*Mar  1 00:16:43.295: SB11: VI11 Hello  in  172.16 11.112 Active pri 50 ip 172.10.11.115
*Mar  1 00:16:43.295: SB11: VI11 Active router is 172.16.11.112
*Mar  1 00:16:43.295: SB11: VI11 Listen:  n/Hello rcvd from lower pri Active router (50/172.16.11.112)
*Mar  1 00:16:43.295: SB11: VI11 Active router is local, was 172.16.11.112
*Mar  1 00:16:43.299: %STANDBY-6-STATECHANGE: Vlan11 Group 11 state Listen -> Active
*Mar  1 00:16:43.299: SB11: VI11 Hello  out 172.16.11.111 Active  pri 100 ip 172.16.11.115
*Mar  1 00:16:43.303: SB11: VI11 Hello  in  172.16.11.112 Speak  pri 50 ip 172.16.11.115
*Mar  1 00:16:46.207: SB11: VI11 Hello  out 172.16.11.111 Active  pri 100 ip 172.16.11.115
*Mar  1 00:16:49.095: SB11: VI11 Hello  in  172.16.11.112 Speak  pri 50 ip 172.16.11.115
```

Based on the debug output shown in the exhibit, which three statements about HSRP are true? Select three.

A. The final active router is 172.16.11.111
B. The 172.16.11.111 router has preempt configured.
C. The 172.16.11.112 router has a more preferred priority than the 172.16.11.111 router does.
D. 172.16.1.115 is the virtual HSRP IP address.
E. The 172.16.11.112 router has nonpreempt configured.
F. The 172.16.11.112 router is using default HSRP priority.

Answer: A, B, D

Explanation:
Each router in an HSRP group has its own unique IP address assigned to an interface. This address is used for all routing protocol and management traffic initiated by or destined to the router. In addition, each router has a common gateway IP address, the virtual router address, that is kept alive by HSRP. This address is also referred to as the HSRP address or the standby address. Clients can point to that virtual router address as their default gateway, knowing that a router always keeps that address active. Keep in mind that the actual interface address and the virtual (standby) address must be configured to be in the same IP subnet. You can assign the HSRP address with the following interface command:
Switch(config-if)# standby group ip ip-address [secondary]
When HSRP is used on an interface that has secondary IP addresses, you can add the secondarykeyword so that HSRP can provide a redundant secondary gateway address.
You can configure a router to preempt or immediately take over the active role if its priority is the highest at any time. Use the following interface configuration command to allow preemption:
Switch(config-if)# standby group preempt [delay seconds]
By default, the router can preempt another immediately, without delay. You can use the delay keyword to force it to wait for seconds before becoming active. This is usually done if there are routing protocols that need time to converge.

QUESTION 301
Exhibit

```
CertkillerA# show standby

Ethernet0/1
  State is Active
    2 state changes, last state change 00:30:59
  Virtual IP address is 10.1.0.20
    Secondary virtual IP address 10.1.0.21
  Active virtual MAC address is 0004.4d82.7981
    Local virtual MAC address is 0004.4d82.7981 (bia)
  Hello time 4 sec, hold time 12 sec
    Next hello sent in 1.412 secs
  Preemption enabled, min delay 50 sec, sync delay 40 sec
  Active router is local
  Standby router is 10.1.0.6, priority 75 (expires in 9.184 sec)
  Priority 95 (configured 120)
    Tracking 2 objects, 0 up
      Down Interface Ethernet0/2, pri 15
      Down Interface Ethernet0/3
  IP redundancy name is "HSRP1", advertisement interval is 34 sec
```

Study the router output displayed in the exhibit.
Which two items are correct? Select two.

A. Certkiller A will assume the active state if its priority is the highest.
B. If Ethernet 0/2 goes down, the standby router will take over.
C. When Ethernet 0/3 of Certkiller A comes back up, the priority will become 105.
D. The local IP address of Certkiller A is 10.1.0.6.
E. The local IP address of Certkiller A is 10.1.0.20.

Answer: A, C

Explanation:
HSRP election is based on a priority value (0 to 255) that is configured on each router in the group. By default, the priority is 100. The router with the highest priority value (255 is highest) becomes the active router for the group. If all router priorities are equal or set to the default value, the router with the highest IP address on the HSRP interface becomes the active router. To set the priority, use the following interface configuration command:
Switch(config-if)# standby group priority priority
When HSRP is configured on an interface, the router progresses through a series of states before becoming active. This forces a router to listen for others in a group and see where it fits into the pecking order. The HSRP state sequence is Disabled, Init, Listen, Speak, Standby, and, finally, Active.

QUESTION 302
Exhibit

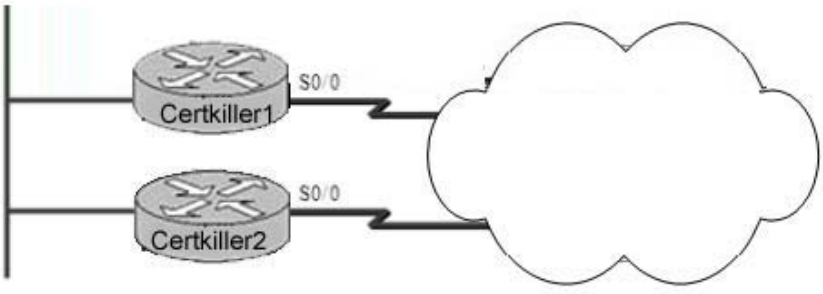

HostName Certkiller1	HostName Certkiller2
... interface fastethernet 0/0 ip address 172.16.15.1 255.255.255.0 standby 1 ip 172.16.15.20 255.255.255.0	... interface fastethernet 0/0 ip address 172.16.15.2 255.255.255.0 standby 1 ip 172.16.15.20 255.255.255.0 standby 1 preempt

Which command will ensure that Certkiller 2 will be the primary router for traffic using the gateway address of 172.16.15.20?

A. On Certkiller 2 add the command standby 1 priority 80
B. On Certkiller 1 add the command standby 1 priority 110
C. On Certkiller 1 add the command standby 1 priority 80
D. On Certkiller 2 remove the command standby 1 preempt

Answer: C

Explanation:
HSRP election is based on a priority value (0 to 255) that is configured on each router in the group. By default, the priority is 100. The router with the highest priority value (255 is highest) becomes the active router for the group. If all router priorities are equal or set to the default value, the router with the highest IP address on the HSRP interface becomes the active router. To set the priority, use the following interface configuration command:
Switch(config-if)# standby group priority priority
When HSRP is configured on an interface, the router progresses through a series of states before becoming active. This forces a router to listen for others in a group and see where it fits into the pecking order. The HSRP state sequence is Disabled, Init, Listen, Speak, Standby, and, finally, Active.

QUESTION 303
Exhibit

```
*Mar  1 00:12:16.871: SB11: VI11 Hello  in  172.16.11.112 Active  pri 50 ip 172.16.11.115
*Mar  1 00:12:16.871: SB11: VI11 Active router is 172.16.11.112
*Mar  1 00:12:18.619: %LINK-3-UPDOWN:  Interface Vlan11, changed state to up
*Mar  1 00:12:18.623: SB:  VI11 Interface up
*Mar  1 00:12:18.623: SB11: VI11 Init: a/HSRP enabled
*Mar  1 00:12:18.623: SB11: VI11 Init -> Listen
*Mar  1 00:12:19.619: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan11, changed state to up
*Mar  1 00:12:19.819: SB11: VI11 Hello  in  172.16.11.112 Active  pri 50 ip 172.16.11.115
*Mar  1 00:12:19.819: SB11: VI11 Listen: h/Hello rcvd from lower pri Active router (50/172.16.11.112)
*Mar  1 00:12:22.815: SB11: VI11 Hello  in  172.16.11.112 Active  pri 50 ip 172.16.11.115
*Mar  1 00:12:22.815: SB11: VI11 Listen: h/Hello rcvd from lower pri Active router
*Mar  1 00:12:25.683: SB11: VI11 Hello  in  172.16.11.112 Active  pri 50 ip 172.16.11.115
*Mar  1 00:12:25.683: SB11: VI11 Listen: h/Hello rcvd from lower pri Active router (50/172.16.11.112)
*Mar  1 00:12:28.623: SB11: VI11 Listen: d/standby time expire (unknown)
*Mar  1 00:12:28.623: SB11: VI11 Listen -> Speak
*Mar  1 00:12:28.623: SB11: VI11 Hello  out 172.16.11.111 Speak  pri 100 ip 172.16.11.115
*Mar  1 00:12:28.659: SB11: VI11 Hello  in  172.16.11.112 Active  pri 50 ip 172.16.11.115
*Mar  1 00:12:28.659: SB11: VI11 Speak: h/Hello rcvd from lower pri Active router (50/172.16.11.112)
*Mar  1 00:12:31.539: SB11: VI11 Hello  in  172.16.11.112 Active  pri 50 ip 172.16.11.115
*Mar  1 00:12:31.539: SB11: VI11 Speak: h/Hello rcvd from lower pri Active router (50/172.16.11.112)
*Mar  1 00:12:31.575: SB11: VI11 Hello  out 172.16.11.111 Speak  pri 100 ip 172.16.11.115
*Mar  1 00:12:34.491: SB11: VI11 Hello  in  172.16.11.112 Active  pri 50 ip 172.16.11.115
```

What can be determined about the HSRP relationship from the displayed debug output?

A. The pre-empt feature is not enable don the 172.16.11.111 router.
B. The nonpreempt feature is enabled on the 172.16.11.112 router.
C. Router 172.16.11.111 will be the activate router because its HSRP priority is preferred over router 172.16.11.112.
D. Router 172.16.11.111 will be the activate router because its HSRP priority is preferred over router 172.16.11.111.
E. The IP address 172.16.11.111 is the virtual HSRP router IP address.
F. The IP address 172.16.11.112 is the virtual HSRP router IP address.

Answer: A

Explanation:
The standby preempt interface configuration command allows the router to become the active router when its priority is higher than all other HSRP-configured routers in this Hot Standby group. The configurations of both routers include this command so that each router can be the standby router for the other router. The 1 indicates that this command applies to Hot Standby group 1. If you do not use the standby preempt command in the configuration for a router, that router cannot become the active router.

QUESTION 304
Routers CK1 and CK2 are configured for HSRP as shown below:
Router CK1 :
interface ethernet 0
ip address 20.6.2.1 255.255.255.0
standby 35 ip 20.6.2.21
standby 35 priority 100
interface ethernet 1

ip address 20.6.1.1.2 255.255.255.0
standby 34 ip 20.6.1.21
Router CK2 :
interface ethernet 0
ip address 20.6.2.2 255.255.255.0
standby 35 ip 20.6.2.21
interface ethernet 1
ip address 20.6.1.1.1 255.255.255.0
standby 34 ip 20.6.1.21
standby 34 priority 100
You have configured the routers CK1 & CK2 with HSRP. While debugging router CK2 you notice very frequent HSRP group state transitions. What is the most likely cause of this?

A. physical layer issues
B. no spanning tree loops
C. use of non-default HSRP timers
D. failure to set the command standby 35 preempt

Answer: A

Explanation: CK2 is not able to from the standby state to reach the active state. This could be caused by missing HSRP hello messages. There are several possible causes for HSRP packets to get lost between the peers. The most common problems are Physical Layer Problems or excessive network traffic caused by Spanning-Tree Issues.
Note:
Hot Standby Routing Protocol (HSRP) is a Cisco proprietary protocol used for allowing redundant connections. It can keep core connectivity if the primary routing process fails. HSRP defines six states in which an HSRP router may run: initial, learn, listen, speak, standby, and active.
Incorrect Answers:
B: Spanning tree loops does not affect this problem.
C: Not a likely cause. Besides, in the example here the default values were indeed used.
D: If the Preempt option is set, then an election of the Active router will take place. This process is called a coup. However, an election would take place by default.
Reference:
Understanding and Troubleshooting HSRP Problems in Catalyst Switch Networks
http://www.cisco.com/warp/public/473/62.shtml
RFC 2281, Cisco Hot Standby Router Protocol (HSRP)

QUESTION 305
Which three of the following network features are methods used to achieve high availability? (Select all that apply.)

A. Spanning Tree Protocol (STP)

B. Delay reduction
C. Hot Standby Routing Protocol (HSRP)
D. Dynamic routing protocols
E. Quality of Service (QoS)
F. Jitter management

Answer: A, C, D

Explanation:
Because the importance of high availability networks is increasingly being recognized, many organizations are beginning to make reliability/availability features a key selection criteria for network infrastructure products. With this in mind, Cisco Systems engaged ZD Tag to observe and confirm the results of a series of tests demonstrating the high availability features of Cisco Catalyst Layer 2/Layer 3 switches. In order to maximize the relevance of the results, the demonstration was based on a model of a "real world" campus (in one of Cisco's Enterprise Solution Center labs in San Jose, California). This switched internetwork consisted of wiring closet, wiring center, and backbone switches and conformed to Cisco's modular three-tier (Access/Distribution/Core) design philosophy. The testing demonstrated the following high availability and resilience features of Catalyst switches:
1. per-VLAN Spanning Tree (PVST) using Cisco's InterSwitch Link (ISL) and 802.1Q VLAN Trunking
2. Cisco Spanning Tree Enhancements, including UplinkFast and PortFast
3. Cisco Hot Standby Router Protocol (HSRP) and HSRP Track
4. Cisco IOS per-destination load balancing over equal cost OSPF paths
5. Cisco IOS fast convergence for OSPF
Reference:
http://www.cisco.com/warp/public/779/largeent/learn/technologies/campuslan.pdf

---

QUESTION 306
LDAP is being used throughout the Certkiller wireless network. Which statement about the Lightweight Access Point Protocol (LWAPP) protocol is true?

A. The processing of 802.11 data and management protocols and access point capabilities is distributed between a lightweight access point and a centralized WLAN controller.
B. LWAPP advertises its WDS capability and participates in electing the best WDS device for the wireless LAN.
C. LWAPP aggregates radio management forward information and sends it to a wireless LAN solution engine.
D. LWAPP authenticates all access points in the subnet and establishes a secure communication channel with each of them.
E. None of the above

Answer: A

Explanation:

The control traffic between the access point and the controller is encapsulated with the LWAPP. The control traffic is encrypted via the Advanced Encryption Standard (AES). The data traffic between the access point and controller is also encapsulated with LWAPP. The data traffic is not encrypted. It is switched at the WLAN controller, where VLAN tagging and quality of service (QoS) are also applied.
The lightweight architecture splits the processing of the 802.11 protocol between two devices, the access point and a centralized Cisco WLC. The processing of the 802.11 data and management protocols and the access point functionality is also divided between the two devices. This approach is called split MAC.

QUESTION 307
Study the exhibit shown below carefully. A Cisco Aironet Wireless LAN Client Adapter has been installed and configured through the ADU on the PC. The Aironet System Tray Utility (ASTU) has been enabled during the installation and the icon appears in the system tray area in the lower right of the desktop as shown:

What is the significance of this icon?

A. It indicates that the client adapter is not associated to an access point or another client.
B. It indicates that the radio of the client adapter is disabled.
C. It indicates that the client adapter is associated to an access point or another client, but the user is not EAP authenticated.
D. It indicates that the client adapter is associated to an access point or another client, that the user is authenticated if the client adapter is configured for EAP authentication, and that the signal strength is poor.
E. It indicates that the client adapter is associated to an access point or another client, that the user is authenticated if the client adapter is configured for EAP authentication, and that the signal strength is fair.
F. It indicates that the client adapter is associated to an access point or another client, that the user is authenticated if the client adapter is configured for EAP authentication, and that the signal strength is excellent or good.

Answer: D

Explanation:
heappearance of the ASTU icon indicates the connection status of your client adapter. ASTU reads the client adapter status and updates the icon every 1 to 5 seconds,

depending on the value entered for the Refresh Interval on the Display Settings window.

Interpreting the ASTU Icon	
Icon	**Description**
	A white icon indicates that the client adapter's radio is disabled.
	A dark gray icon indicates that the client adapter is not associated to an access point (in infrastructure mode) or another client (in ad hoc mode).
	A light gray icon indicates that the client adapter is associated to an access point (in infrastructure mode) or another client (in ad hoc mode) but the user is not EAP authenticated.
	A green icon indicates that the client adapter is associated to an access point (in infrastructure mode) or another client (in ad hoc mode), the user is authenticated if the client adapter is configured for EAP authentication, and the signal strength is excellent or good.
	A yellow icon indicates that the client adapter is associated to an access point (in infrastructure mode) or another client (in ad hoc mode), the user is authenticated if the client adapter is configured for EAP authentication, and the signal strength is fair.
	A red icon indicates that the client adapter is associated to an access point (in infrastructure mode) or another client (in ad hoc mode), the user is authenticated if the client adapter is configured for EAP authentication, and the signal strength is poor.

QUESTION 308
LDAP is being utilized throughout the Certkiller Wireless network. Which issue or set of issues does the Lightweight Access Point Protocol (LWAPP) address?

A. Reduction of processing in wireless controllers
B. Distributed approach to authentication, encryption, and policy enforcement
C. Access point discovery, information exchange, and configuration
D. Provides security by blocking communication between access points and wireless clients
E. None of the above.

Answer: C

Explanation:
The control traffic between the access point and the controller is encapsulated with the LWAPP. The control traffic is encrypted via the Advanced Encryption Standard (AES). The data traffic between the access point and controller is also encapsulated with

LWAPP. The data traffic is not encrypted. It is switched at the WLAN controller, where VLAN tagging and quality of service (QoS) are also applied.
Lightweight access points first search for a WLAN controller using LWAPP in Layer 2 mode. Then the access point searches for a WLAN in Layer 3 mode.
The access point requests an IP address via DHCP. The access point then sends a LWAPP discovery request to the management IP address of the WLAN controller via a broadcast.
The WLAN controller responds with a discovery response from the manager IP address. This response includes the number of access points that are currently associated to that access point manager interface and the access point manager IP address.
The access point chooses the access point manager with the least number of associated access points and sends the join request.
All subsequent LWAPP communication is done to the access point manager IP address of the WLAN controller.
Real-timeframe exchange and certain real-time portions of MAC management are accomplished within the access point.
Authentication, securitymanagement, and mobility are handled by WLAN controllers.
Data and control messages are exchanged between the accesspoint and the WLAN controller using LWAPP.
Control messages are encrypted.
All client data trafic is sent via the WLAN controller.

QUESTION 309
A Cisco Aironet Wireless LAN Adapter CB21AG is inserted into a Certkiller user's PC cardbus slot. Both the green status LED and the amber activity LED are blinking slowly. What is the condition of the adapter?

A. The adapter is scanning for the wireless network for which it is configured.
B. The adapter is in power save mode.
C. The adapter is transmitting or receiving data while associated to an access point or another client.
D. The adapter is not receiving power.
E. The adapter is associated to an access point or another client.
F. None of the above

Answer: E

Explanation:
The client adapter shows messages through its two LEDs.

Table, LED Operating Messages		
Status LED (green)	Activity LED (amber)	Condition
Off	Off	Client adapter is not receiving power.
Blinking slowly	Off	Client adapter is in power save mode.
On	Off	Client adapter has awakened from power save mode.
Alternating blink:		Client adapter is scanning for the wireless network for which it is configured.
On	Off	
Off	On	
Blinking slowly	Blinking slowly	Client adapter is associated to an access point (in infrastructure mode) or another client (in ad hoc mode).
Blinking quickly	Blinking quickly	Client adapter is transmitting or receiving data while associated to an access point (in infrastructure mode) or another client (in ad hoc mode).

Reference:
http://www.cisco.com/en/US/products/hw/wireless/ps4555/products_installation_and_configuration_guide_chap

---

QUESTION 310
Study the exhibit below carefully:

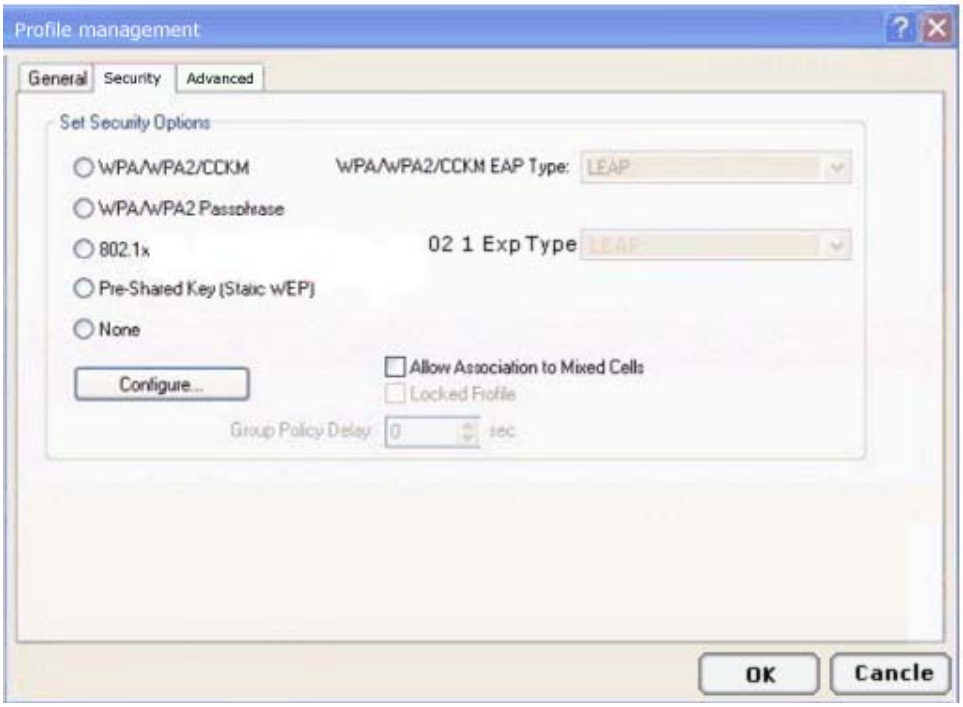

When a profile is configured for a user on the Certkiller network in the Aironet Desktop Utility, which security option permits the configuration of host-based Extensible Authentication Protocol (EAP)?

A. Pre-Shared Key (Static WEP)
B. WPA/WPA2/CCKM
C. 802.1x
D. WPA/WPA2 Passphrase
E. None of the above

Answer: C

Explanation:
The IEEE 802.1x standard defines a port-based access control and authentication protocol that restricts unauthorized workstations from connecting to a LAN through publicly accessible switch ports. The authentication server authenticates each workstation that is connected to a switch port before making available any services offered by the switch or the LAN. Until the workstation is authenticated, 802.1x access control allows only Extensible Authentication Protocol over LAN (EAPOL) traffic through the port to which the workstation is connected. After authentication succeeds, normal traffic can pass through the port. With 802.1x port-based authentication, the devices in the network have specific roles, as follows:
Client: The device (workstation) that requests access to the LAN and switch services, and responds to requests from the switch. The workstation must be running 802.1x-compliant client software, such as what is offered in the Microsoft Windows XP operating system. (The port that the client is attached to is the supplicant [client] in the IEEE 802.1x specification.)
Authentication server: Performs the actual authentication of the client. The authentication server validates the identity of the client and notifies the switch whether or not the client is authorized to access the LAN and switch services. Because the switch acts as the proxy, the authentication service is transparent to the client. The RADIUS security system with Extensible Authentication Protocol (EAP) extensions is the only supported authentication server.
Switch (also called the authenticator): Controls physical access to the network based on the authentication status of the client. The switch acts as an intermediary (proxy) between the client (supplicant) and the authentication server, requesting identifying information from the client, verifying that information with the authentication server, and relaying a response to the client. The switch uses a RADIUS software agent, which is responsible for encapsulating and decapsulating the EAP frames and interacting with the authentication server.

QUESTION 311
The following output was shown on a Certkiller switch:

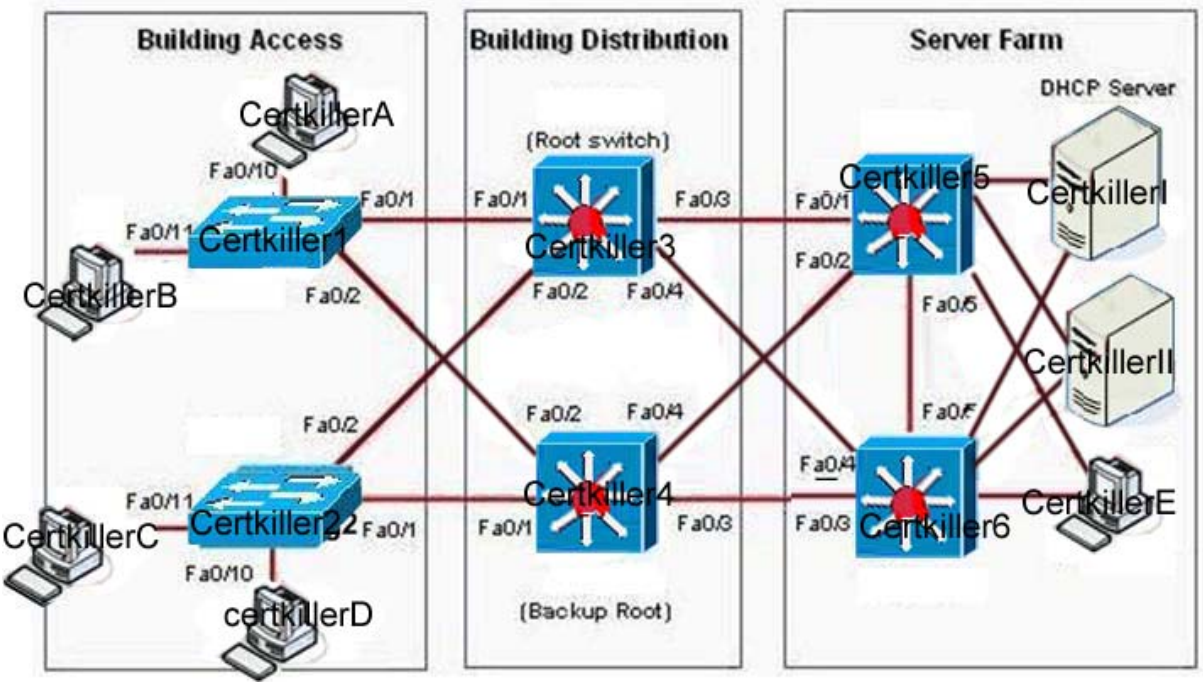

Study the exhibit carefully. An attacker is connected to interface Fa0/11 on switch Certkiller 4 and attempts to establish a DHCP server for a man-in-middle attack. Which recommendation, if followed, would mitigate this type of attack?

A. All switch ports connecting to servers in the Server Farm block should be configured as DHCP untrusted ports.
B. All switch ports connecting to hosts in the Building Access block should be configured as DHCP trusted ports.
C. All switch ports in the Server Farm block should be configured as DHCP untrusted ports.
D. All switch ports connecting to hosts in the Building Access block should be configured as DHCP untrusted ports.
E. All switch ports in the Building Access block should be configured as DHCP untrusted ports.
F. All switch ports in the Building Access block should be configured as DHCP trusted ports.
G. None of the above

Answer: D

Explanation:
One of the ways that an attacker can gain access to network traffic is to spoof responses that would be sent by a valid DHCP server. The DHCP spoofing device replies to client DHCP requests. The legitimate server may reply also, but if the spoofing device is on the same segment as the client, its reply to the client may arrive first.
The intruder's DHCP reply offers an IP address and supporting information that designates the intruder as the default gateway or Domain Name System (DNS) server. In the case of a gateway, the clients will then forward packets to the attacking device, which

will in turn send them to the desired destination. This is referred to as a "man-in-the-middle" attack, and it may go entirely undetected as the intruder intercepts the data flow through the network.
Untrusted ports are those that are not explicitly configured as trusted. A DHCP binding table is built for untrusted ports. Each entry contains the client MAC address, IP address, lease time, binding type, VLAN number, and port ID recorded as clients make DHCP requests. The table is then used to filter subsequent DHCP traffic. From a DHCP snooping perspective, untrusted access ports should not send any DHCP server responses, such as DHCPOFFER, DHCPACK, DHCPNAK.

QUESTION 312
You suspect that a hacker may be performing a MAC address flooding attack somewhere within the Certkiller network. Which description correctly describes a MAC address flooding attack?

A. The attacking device spoofs a source MAC address of a valid host currently in the CAM table. The switch then forwards frames destined for the valid host to the attacking device.
B. Frames with unique, invalid destination MAC addresses flood the switch and exhaust CAM table space. The result is that new entries cannot be inserted because of the exhausted CAM table space, and traffic is subsequently flooded out all ports.
C. The attacking device crafts ARP replies intended for valid hosts. The MAC address of the attacking device then becomes the destination address found in the Layer 2 frames sent by the valid network device.
D. The attacking device crafts ARP replies intended for valid hosts. The MAC address of the attacking device then becomes the source address found in the Layer 2 frames sent by the valid network device.
E. The attacking device spoofs a destination MAC address of a valid host currently in the CAM table. The switch then forwards frames destined for the valid host to the attacking device.
F. Frames with unique, invalid source MAC addresses flood the switch and exhaust CAM table space. The result is that new entries cannot be inserted because of the exhausted CAM table space, and traffic is subsequently flooded out all ports.
G. None of the above

Answer: F

Explanation:
A common Layer 2 or switch attack as of this writing is MAC flooding, resulting in a switch's CAM table overflow, which causes flooding of regular data frames out all switch ports. This attack can be launched for the malicious purpose of collecting a broad sample of traffic or as a denial of service (DoS) attack.
A switch's CAM tables are limited in size and therefore can contain only a limited number of entries at any one time. A network intruder can maliciously flood a switch with a large number of frames from a range of invalid source MAC addresses. If enough new entries are made before old ones expire, new valid entries will not be accepted.

Then, when traffic arrives at the switch for a legitimate device that is located on one of the switch ports that was not able to create a CAM table entry, the switch must flood frames to that address out all ports. This has two adverse effects:
* The switch traffic forwarding is inefficient and voluminous.
* An intruding device can be connected to any switch port and capture traffic that is not normally seen on that port.
If the attack is launched before the beginning of the day, the CAM table would be full when the majority of devices are powered on. Then frames from those legitimate devices are unable to create CAM table entries as they power on. If this represents a large number of network devices, the number of MAC addresses for which traffic will be flooded will be high, and any switch port will carry flooded frames from a large number of devices.

QUESTION 313
The Certkiller switched LAN is shown below:

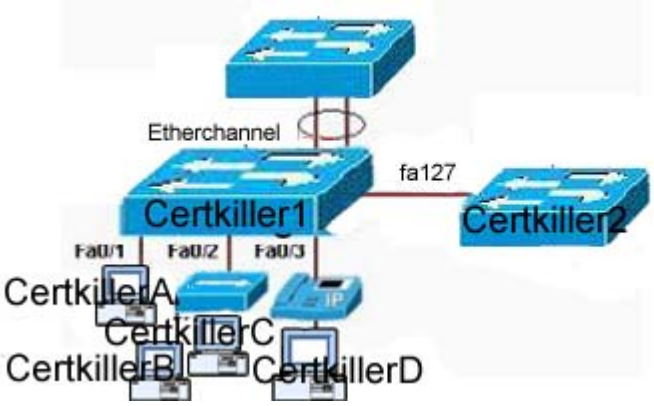

Based on the information shown above, which interface or interfaces on switch Certkiller 1 can have the port security feature enabled?

A. Ports 0/1, 0/2, 0/3 and the trunk port 0/22
B. Ports 0/1 and 0/2
C. Ports 0/1, 0/2, 0/3, the trunk port 0/22 and the EtherChannel ports
D. Ports 0/1, 0/2 and 0/3
E. Port 0/1
F. The trunk port 0/22 and the EtherChannel ports
G. None of the above

Answer: D

Explanation:
Port security is a feature supported on Cisco Catalyst switches that restricts a switch port to a specific set or number of MAC addresses. Those addresses can be learned dynamically or configured statically. The port will then provide access to frames from only those addresses. If, however, the number of addresses is limited to four but no specific MAC addresses are configured, the port will allow any four MAC addresses to be learned dynamically, and port access will be limited to those four dynamically learned addresses. A port security feature called "sticky learning," available on some switch platforms, combines the features of dynamically learned and statically configured

addresses. When this feature is configured on an interface, the interface converts dynamically learned addresses to "sticky secure" addresses. This adds them to the running configuration as if they were configured using the switchport port-security mac-address command.

QUESTION 314
A Certkiller Switch was configured as shown below:

```
Certkiller Switch    #configure terminal
Certkiller Switch (config)#aaa new-model
Certkiller Switch (config)#aaa authentication dot1x default group radius
Certkiller Switch (config)#interface fastethernet 0/1
Certkiller Switch (config-if)#dot1x port-control force-authorized
Certkiller Switch (config-if)#end
```

Based on this information, how will interface FastEthernnet0/1 respond when an 802.1x-enabled client connects to the port?

A. The switch port will enable 802.1x port-based authentication and begin relaying authentication messages between the client and the authentication server.
B. The switch will cause the port to remain in the unauthorized state, ignoring all attempts by the client to authenticate.
C. The switch will uniquely authorize the client by using the client MAC address.
D. The switch port will disable 802.1x port-based authentication and cause the port to transition to the authorized state without any further authentication exchange.
E. None of the above.

Answer: D

Explanation:
The IEEE 802.1x standard defines a port-based access control and authentication protocol that restricts unauthorized workstations from connecting to a LAN through publicly accessible switch ports. The authentication server authenticates each workstation that is connected to a switch port before making available any services offered by the switch or the LAN. Until the workstation is authenticated, 802.1x access control allows only Extensible Authentication Protocol over LAN (EAPOL) traffic through the port to which the workstation is connected. After authentication succeeds, normal traffic can pass through the port.
You control the port authorization state by using the dot1x port-control interface configuration command and these keywords:
force-authorized: Disables 802.1x port-based authentication and causes the port to transition to the authorized state without any authentication exchange required. The port transmits and receives normal traffic without 802.1x-based authentication of the client. This is the default setting.
force-unauthorized: Causes the port to remain in the unauthorized state, ignoring all attempts by the client to authenticate. The switch cannot provide authentication services to the client through the interface.
auto: Enables 802.1x port-based authentication and causes the port to begin in the

unauthorized state, allowing only EAPOL frames to be sent and received through the port. The authentication process begins when the link state of the port transitions from down to up (authenticator initiation) or when an EAPOL-start frame is received (supplicant initiation). The switch requests the identity of the client and begins relaying authentication messages between the client and the authentication server. The switch uniquely identifies each client attempting to access the network by using the client MAC address.
Example:

1	Switch(config)#**aaa new-model**
2.	Create an 802.1x port-based authentication method list. Switch(config)#**aaa authentication dot1x** {**default**} *method1* [*method2...*]
3.	Globally enable 802.1x port-based authentication. Switch(config)#**dot1x system-auth-control**
4.	Enter interface configuration mode and specify the interface to be enabled for 802.1x port-based authentication. Switch(config)#**interface** *type slot/port* 5. Enable 802.1x port-based authentication on the interface. Switch(config-if)#**dot1x port-control auto**
6.	Return to privileged EXEC mode. Switch(config)#**end**

QUESTION 315
A Certkiller switch was configured as shown below:
switchport mode access
switchport port-security
switchport port-security maximum 2
switchport port-security mac-address 0002.0002.0002
switchport port-security violation shutdown
Given the configuration output shown above, what happens when a host with the MAC address of 0003.0003.0003 is directly connected to the switch port?

A. The host will be allowed to connect.
B. The port will shut down.
C. The host can only connect through a hub/switch where 0002.0002.0002 is already connected.
D. The host will be refused access.
E. None of the above

Answer: A

Explanation:
Steps of Implementing Port Security:

Step	Description
1.	Enables port security. Switch(config-if)#**switchport port-security**
2.	Sets a maximum number of MAC addresses that will be allowed on this port. Default is one. Switch(config-if)#**switchport port-security maximum** *value*
3.	Specifies which MAC addresses will be allowed on this port (optional). Switch(config-if)#**switchport port-security mac-** **address** *mac-address* Switch(cor.fig-if)#**switchport port-security mac-address** *mac-address*
4.	Defines what action an interface will take if a nonallowed MAC address attempts access. Switch(config-if)#**switchport port-security violation** {shutdown | restrict | protect}

In Exhibit two MAC addresses are allowed so that host will be allowed to connect.

QUESTION 316
In order to enhance security on the Certkiller network, users must be authenticated using 802.1X. When authentication is required, where must 802.1X be configured in order to connect a PC to a switch?

A. Switch port and local router port
B. Switch port and client PC
C. Client PC only
D. Switch port only
E. None of the above

Answer: B

Explanation:
The IEEE 802.1x standard defines a port-based access control and authentication protocol that restricts unauthorized workstations from connecting to a LAN through publicly accessible switch ports. The authentication server authenticates each workstation that is connected to a switch port before making available any services offered by the switch or the LAN. Until the workstation is authenticated, 802.1x access control allows only Extensible Authentication Protocol over LAN (EAPOL) traffic through the port to which the workstation is connected. After authentication succeeds, normal traffic can pass through the port.

With 802.1x port-based authentication, the devices in the network have specific roles, as follows:
Client: The device (workstation) that requests access to the LAN and switch services, and responds to requests from the switch. The workstation must be running 802.1x-compliant client software, such as what is offered in the Microsoft Windows XP operating system. (The port that the client is attached to is the supplicant [client] in the IEEE 802.1x specification.)
Authentication server: Performs the actual authentication of the client. The authentication server validates the identity of the client and notifies the switch whether or not the client is authorized to access the LAN and switch services. Because the switch acts as the proxy, the authentication service is transparent to the client. The RADIUS security system with Extensible Authentication Protocol (EAP) extensions is the only supported authentication server.
Switch(also called the authenticator): Controls physical access to the network based on the authentication status of the client. The switch acts as an intermediary (proxy) between the client (supplicant) and the authentication server, requesting identifying information from the client, verifying that information with the authentication server, and relaying a response to the client. The switch uses a RADIUS software agent, which is responsible for encapsulating and decapsulating the EAP frames and interacting with the authentication server.

QUESTION 317
Switch Certkiller 1 was configured as shown below:

```
Certkiller (config) #vlan access-map pass 10
Certkiller1 (config access-map) # match ip address Certkillerv
Certkiller1(config access-map) # action forward
Certkiller1 (config) # vlan filter pass vlan-list 5-10
```

Based on the output shown above, which statement is true?

A. IP traffic matching access list Certkiller is forwarded through VLANs 5-10.
B. All VLAN traffic in VLANs 5-10 that match access list Certkiller will be forwarded, and all else will be dropped.
C. IP traffic matching VLAN list 5-10 will be forwarded, and all other traffic will be dropped.
D. All VLAN traffic matching VLAN list 5-10 will be forwarded, and all traffic matching access list Certkiller is dropped.
E. None of the above

Answer: B

Explanation:
VLAN maps, also known as VLAN ACLs or VACLs, can filter all traffic traversing a switch. VLAN maps can be configured on the switch to filter all packets that are routed into or out of a VLAN, or are bridged within a VLAN. VLAN maps are used strictly for security packet filtering. Unlike router ACLs, VLAN maps are not defined by direction (input or output).

To create a VLAN map and apply it to one or more VLANs, perform these steps:
1. Create the standard or extended IP ACLs or named MAC extended ACLs to be applied to the VLAN. This access-list will select the traffic that will be either forwarded or dropped by the access-map. Only traffic matching the 'permit' condition in an access-list will be passed to the access-map for further processing.
2. Enter the vlan access-map access-map-name [sequence] global configuration command to create a VLAN ACL map entry. Each access-map can have multiple entries. The order of these entries is determined by the sequence. If no sequence number is entered, access-map entries are added with sequence numbers in increments of 10.
3. In access map configuration mode, optionally enter an action forward or action drop. The default is to forward traffic. Also enter the match command to specify an IP packet or a non-IP packet (with only a known MAC address), and to match the packet against one or more ACLs (standard or extended).
4. Use the vlan filter access-map-name vlan-list vlan-list global configuration command to apply a VLAN map to one or more VLANs. A single access-map can be used on multiple VLANs.

QUESTION 318
Certkiller is implementing 802.1X in order to increase network security. In the use of 802.1X access control, which three protocols are allowed through the switch port before authentication takes place? (Select three)

A. EAP-over-LAN
B. EAP MD5
C. STP
D. protocols not filtered by an ACL
E. CDP
F. TACACS+

Answer: A, C, E

Explanation:
The IEEE 802.1x standard defines a port-based access control and authentication protocol that restricts unauthorized workstations from connecting to a LAN through publicly accessible switch ports. The authentication server authenticates each workstation that is connected to a switch port before making available any services offered by the switch or the LAN. Until the workstation is authenticated, 802.1x access control allows only Extensible Authentication Protocol over LAN (EAPOL) traffic through the port to which the workstation is connected. After authentication succeeds, normal traffic can pass through the port.
The Authentication server performs the actual authentication of the client. The authentication server validates the identity of the client and notifies the switch whether or not the client is authorized to access the LAN and switch services. Because the switch acts as the proxy, the authentication service is transparent to the client. In this release, the Remote Authentication Dial-In User Service (RADIUS) security system with Extensible Authentication Protocol (EAP) extensions is the only supported authentication server; it

is available in Cisco Secure Access Control Server version 3.0. RADIUS operates in a client/server model in which secure authentication information is exchanged between the RADIUS server and one or more RADIUS clients.
Spanning-Tree Protocol (STP) is a Layer 2 protocol that utilizes a special-purpose algorithm to discover physical loops in a network and effect a logical loop-free topology. STP creates a loop-free tree structure consisting of leaves and branches that span the entire Layer 2 network. The actual mechanics of how bridges communicate and how the STP algorithm works will be discussed at length in the following topics. Note that the terms bridge and switch are used interchangeably when discussing STP. In addition, unless otherwise indicated, connections between switches are assumed to be trunks.
CDP is a Cisco proprietary protocol that operates at the Data Link layer. One unique feature about operating at Layer 2 is that CDP functions regardless of what Physical layer media you are using (UTP, fiber, and so on) and what Network layer routed protocols you are running (IP, IPX, AppleTalk, and so on). CDP is enabled on all Cisco devices by default, and is multicast every 60 seconds out of all functioning interfaces, enabling neighbor Cisco devices to collect information about each other. Although this is a multicast message, Cisco switches do not flood that out to all their neighbors as they do a normal multicast or broadcast.
For STP, CDP and EAP-over-LAN are allowed before Authentication.

QUESTION 319
Part of the Certkiller network is shown in the following diagram:

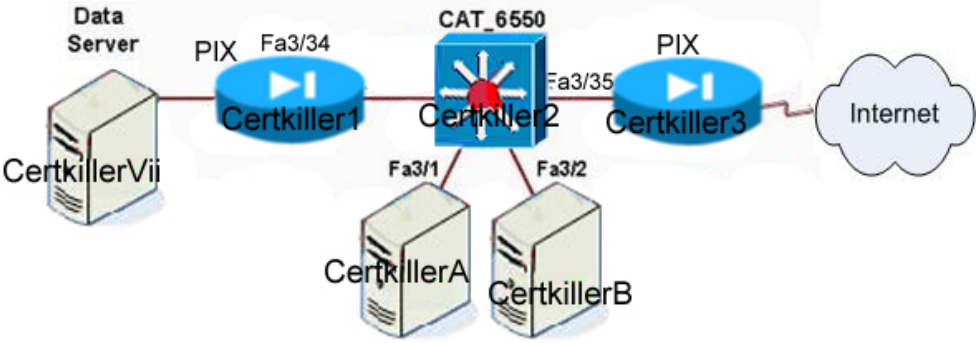

Study the exhibit carefully. The web servers Certkiller A and Certkiller B need to be accessed by external and internal users. For security reasons, the servers should not communicate with each other, although they are located on the same subnet. The servers do need, however, to communicate with a database server located in the inside network. What configuration will isolate the servers from each other?

A. The switch ports 3/1 and 3/2 will be defined as secondary VLAN community ports. The ports connecting to the two firewalls will be defined as primary VLAN promiscuous ports.
B. The switch ports 3/1 and 3/2 will be defined as secondary VLAN isolated ports. The ports connecting to the two firewalls will be defined as primary VLAN promiscuous ports.
C. The switch ports 3/1 and 3/2 and the ports connecting to the two firewalls will be defined as primary VLAN promiscuous ports.
D. The switch ports 3/1 and 3/2 and the ports connecting to the two firewalls will be

defined as primary VLAN community ports.
E. None of the above.

Answer: B

Explanation:
Service providers often have devices from multiple clients, in addition to their own servers, on a single Demilitarized Zone (DMZ) segment or VLAN. As security issues proliferate, it becomes necessary to provide traffic isolation between devices, even though they may exist on the same Layer 3 segment and VLAN. Catalyst 6500/4500 switches implement PVLANs to keep some switch ports shared and some switch ports isolated, although all ports exist on the same VLAN. The 2950 and 3550 support "protected ports," which are functionality similar to PVLANs on a per-switch basis.
A port in a PVLAN can be one of three types:
Isolated: An isolated port has complete Layer 2 separation from other ports within the same PVLAN, except for the promiscuous port. PVLANs block all traffic to isolated ports, except the traffic from promiscuous ports. Traffic received from an isolated port is forwarded to only promiscuous ports.
Promiscuous: A promiscuous port can communicate with all ports within the PVLAN, including the community and isolated ports. The default gateway for the segment would likely be hosted on a promiscuous port, given that all devices in the PVLAN will need to communicate with that port.
Community: Community ports communicate among themselves and with their promiscuous ports. These interfaces are isolated at Layer 2 from all other interfaces in other communities, or in isolated ports within their PVLAN.

QUESTION 320
Certkiller has implemented 802.1X authentication as a security enhancement. Which statement is true about 802.1x port-based authentication?

A. TACACS+ is the only supported authentication server type.
B. If a host initiates the authentication process and does not receive a response, it assumes it is not authorized.
C. RADIUS is the only supported authentication server type.
D. Before transmitting data, an 802.1x host must determine the authorization state of the switch.
E. Hosts are required to have a 802.1x authentication client or utilize PPPoE.
F. None of the above.

Answer: C

Explanation:
The IEEE 802.1x standard defines a port-based access control and authentication protocol that restricts unauthorized workstations from connecting to a LAN through publicly accessible switch ports. The authentication server authenticates each workstation that is connected to a switch port before making available any services offered by the switch or the LAN. Until the workstation is

authenticated, 802.1x access control allows only Extensible Authentication Protocol over LAN (EAPOL) traffic through the port to which the workstation is connected. After authentication succeeds, normal traffic can pass through the port.
Authentication server: Performs the actual authentication of the client. The authentication server validates the identity of the client and notifies the switch whether or not the client is authorized to access the LAN and switch services. Because the switch acts as the proxy, the authentication service is transparent to the client. The RADIUS security system with Extensible Authentication Protocol (EAP) extensions is the only supported authentication server.

QUESTION 321
The DAI feature has been implemented in the Certkiller switched LAN. Which three statements are true about the dynamic ARP inspection (DAI) feature? (Select three)

A. DAI can be performed on ingress ports only.
B. DAI can be performed on both ingress and egress ports.
C. DAI is supported on access ports, trunk ports, EtherChannel ports, and private VLAN ports.
D. DAI should be enabled on the root switch for particular VLANs only in order to secure the ARP caches of hosts in the domain.
E. DAI should be configured on all access switch ports as untrusted and on all switch ports connected to other switches as trusted.
F. DAI is supported on access and trunk ports only.

Answer: A, C, E

Explanation:
To prevent ARP spoofing or "poisoning," a switch must ensure that only valid ARP requests and responses are relayed. DAI prevents these attacks by intercepting and validating all ARP requests and responses. Each intercepted ARP reply is verified for valid MAC-address-to-IP-address bindings before it is forwarded to a PC to update the ARP cache. ARP replies coming from invalid devices are dropped.
DAI determines the validity of an ARP packet based on a valid MAC-address-to-IP-address bindings database built by DHCP snooping. In addition, to handle hosts that use statically configured IP addresses, DAI can also validate ARP packets against user-configured ARP ACLs.
To ensure that only valid ARP requests and responses are relayed, DAI takes these actions:
* Forwards ARP packets received on a trusted interface without any checks
* Intercepts all ARP packets on untrusted ports
* Verifies that each intercepted packet has a valid IP-to-MAC address binding before forwarding packets that can update the local ARP cache
* Drops, logs, or drops and logs ARP packets with invalid IP-to-MAC address bindings

---

QUESTION 322
When authentication is required, where must 802.1x be configured in order to

connect a PC to a switch?

A. Client PC only
B. Switch port only
C. Switch port and client PC
D. Switch port and RADIUS server
E. None of the above.

Answer: D

Explanation:
The IEEE 802.1x standard defines a port-based access control and authentication protocol that restricts unauthorized workstations from connecting to a LAN through publicly accessible switch ports. The authentication server authenticates each workstation that is connected to a switch port before making available any services offered by the switch or the LAN. Until the workstation is authenticated, 802.1x access control allows only Extensible Authentication Protocol over LAN (EAPOL) traffic through the port to which the workstation is connected. After authentication succeeds, normal traffic can pass through the port.
With 802.1x port-based authentication, the devices in the network have specific roles, as follows:
Client: The device (workstation) that requests access to the LAN and switch services, and responds to requests from the switch. The workstation must be running 802.1x-compliant client software, such as what is offered in the Microsoft Windows XP operating system. (The port that the client is attached to is the supplicant [client] in the IEEE 802.1x specification.)
Authentication server: Performs the actual authentication of the client. The authentication server validates the identity of the client and notifies the switch whether or not the client is authorized to access the LAN and switch services. Because the switch acts as the proxy, the authentication service is transparent to the client. The RADIUS security system with Extensible Authentication Protocol (EAP) extensions is the only supported authentication server.
Switch(also called the authenticator): Controls physical access to the network based on the authentication status of the client. The switch acts as an intermediary (proxy) between the client (supplicant) and the authentication server, requesting identifying information from the client, verifying that information with the authentication server, and relaying a response to the client. The switch uses a RADIUS software agent, which is responsible for encapsulating and decapsulating the EAP frames and interacting with the authentication server.

QUESTION 323
VLAN maps have been configured on switch CK1 . Which of the following actions are taken in a VLAN map that does not contain a match clause?

A. Implicit deny feature at end of list.
B. Implicit deny feature at start of list.
C. Implicit forward feature at end of list

D. Implicit forward feature at start of list.

Answer: A

Explanation:
Each VLAN access map can consist of one or more map sequences, each sequence with a match clause and an action clause. The match clause specifies IP, IPX, or MAC ACLs for traffic filtering and the action clause specifies the action to be taken when a match occurs. When a flow matches a permit ACL entry the associated action is taken and the flow is not checked against the remaining sequences. When a flow matches a deny ACL entry, it will be checked against the next ACL in the same sequence or the next sequence. If a flow does not match any ACL entry and at least one ACL is configured for that packet type, the packet is denied.
Reference:
http://www.cisco.com/en/US/products/hw/switches/ps700/products_configuration_guide_chapter09186a008007f

---

QUESTION 324
What is true about access control on bridged and routed VLAN traffic? (Select three)

A. Router ACLs can be applied to the input and output directions of a VLAN interface.
B. Bridged ACLs can be applied to the input and output directions of a VLAN interface.
C. Only router ACLs can be applied to a VLAN interface.
D. VLAN maps and router ACLs can be used in combination.
E. VLAN maps can be applied to a VLAN interface

Answer: A, B, D
Router ACLs are applied on interfaces as either inbound or outbound.
To filter both bridged and routed traffic, VLAN maps can be used by themselves or in conjunction with router ACLs.
VLAN ACLs, also called VLAN maps, which filter both bridged and routed packets.
VLAN maps can be used to filter packets exchanged between devices in the same VLAN.

QUESTION 325
In the use of 802.1X access control, which three products are allowed through the switch port before authentication takes place? Select three.

A. STP
B. CDP
C. EAP MD5
D. TACACS+
E. EAP-over-LAN
F. protocols not filtered by an ACL

Answer: A, B, E

Explanation:
The IEEE 802.1x standard defines a port-based access control and authentication protocol that restricts unauthorized workstations from connecting to a LAN through publicly accessible switch ports. The authentication server authenticates each workstation that is connected to a switch port before making available any services offered by the switch or the LAN. Until the workstation is authenticated, 802.1x access control allows only Extensible Authentication Protocol over LAN (EAPOL) traffic through the port to which the workstation is connected. After authentication succeeds, normal traffic can pass through the port.
The Authentication server performs the actual authentication of the client. The authentication server validates the identity of the client and notifies the switch whether or not the client is authorized to access the LAN and switch services. Because the switch acts as the proxy, the authentication service is transparent to the client. In this release, the Remote Authentication Dial-In User Service (RADIUS) security system with Extensible Authentication Protocol (EAP) extensions is the only supported authentication server; it is available in Cisco Secure Access Control Server version 3.0. RADIUS operates in a client/server model in which secure authentication information is exchanged between the RADIUS server and one or more RADIUS clients.
Spanning-Tree Protocol (STP) is a Layer 2 protocol that utilizes a special-purpose algorithm to discover physical loops in a network and effect a logical loop-free topology. STP creates a loop-free tree structure consisting of leaves and branches that span the entire Layer 2 network. The actual mechanics of how bridges communicate and how the STP algorithm works will be discussed at length in the following topics. Note that the terms bridge and switch are used interchangeably when discussing STP. In addition, unless otherwise indicated, connections between switches are assumed to be trunks.
CDP is a Cisco proprietary protocol that operates at the Data Link layer. One unique feature about operating at Layer 2 is that CDP functions regardless of what Physical layer media you are using (UTP, fiber, and so on) and what Network layer routed protocols you are running (IP, IPX, AppleTalk, and so on). CDP is enabled on all Cisco devices by default, and is multicast every 60 seconds out of all functioning interfaces, enabling neighbor Cisco devices to collect information about each other. Although this is a multicast message, Cisco switches do not flood that out to all their neighbors as they do a normal multicast or broadcast.
For STP, CDP and EAP-over-LAN are allowed before Authentication.

QUESTION 326
A PC host is connected to a switch in the Certkiller network shown below:

Configuration exhibit:

```
Certkiller Switch# show port-security interface fa0/1
Port Security              : Enabled
Port Status                : Secure-down
Violation Mode             : Protect
Aging Time                 : 0 mins
Aging Type                 : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 10
Total MAC Addresses        : 2
Configured MAC Addresses   : 2
Sticky MAC Addresses       : 0
Last Source Address        : 0000.0000.0000:0
Security Violation Count   : 0
```

Study the exhibits carefully. The "show port-security interface fa0/1" command was issued on switch Certkiller 1. Given the output that was generated, which two security statements are true? (Select two)

A. When the number of secure IP addresses reaches 10, the interface will immediately shut down.
B. Interface FastEthernet 0/1 was configured with the switchport port-security aging command.
C. Interface FastEthernet 0/1 was configured with the switchport port-security violation restrict command.
D. When the number of secure MAC addresses reaches 10, the interface will immediately shut down and an SNMP trap notification will be sent.
E. Interface FastEthernet 0/1 was configured with the switchport port-security protect command.
F. None of the above.

Answer: B, E

Explanation:
Port security is a feature supported on Cisco Catalyst switches that restricts a switch port to a specific set or number of MAC addresses. Those addresses can be learned dynamically or configured statically. The port will then provide access to frames from only those addresses. If, however, the number of addresses is limited to four but no specific MAC addresses are configured, the port will allow any four MAC addresses to be learned dynamically, and port access will be limited to those four dynamically learned addresses.
Port Security Implementation:

Step	Description
1.	Enables port security. Switch(config-if)#**switchport port-security**
2.	Sets a maximum number of MAC addresses that will be allowed on this port. Default is one. Switch(config-if)#**switchport port-security maximum** value
3.	Specifies which MAC addresses will be allowed on this port (optional). Switch(config-if)#**switchport port-security mac-** **address** mac-address Switch(config-if)#**switchport port-security mac-address** mac-address
4.	Defines what action an interface will take if a nonallowed MAC address attempts access. Switch(config-if)#**switchport port-security violation** {shutdown | restrict | protect}

When Switch port security rules violate different action can be applied:
1. Protect: Frames from the nonallowed address are dropped, but there is no log of the violation.
2. Restrict: Frames from the nonallowed address are dropped, a log message is created, and a Simple Network Management Protocol (SNMP) trap is sent.
3. Shutdown: If any frames are seen from a nonallowed address, the interface is errdisabled, a log entry is made, an SNMP trap is sent, and manual intervention or errdisable recovery must be used to make the interface usable.
Not D: The port will not be shutdown, because it is in protect mode -- not shutdown.

QUESTION 327
The following show command was issued on switch Certkiller 1:

```
Certkiller1 # show port-security interface fastethernet 5/1

  Port security Enabledport
   Status.Secure up
   Violation mode:  Shutdown
   Maximum MAC Addresses:  11
   Total MAC Addresses:  11
   Configured MAC Addresses:  3
    Aging time: 20 mins
   Aging typer Inactivity
   SecureStatic address aging:  Enabled
   Security Violation count:  0
```

Based on the output shown, what will happen when one additional user is connected to interface FastEthernet 5/1?

A. The interface will be placed into the error-disabled state immediately, and an SNMP trap notification will be sent.

B. The packets with the new source addresses will be dropped until a sufficient number of secure MAC addresses are removed from the secure address list.
C. All secure addresses will age out and be removed from the secure address list. This will cause the security violation counter to increment.
D. The first address learned on the port will be removed from the secure address list and be replaced with the new address.
E. None of the above

Answer: A

Explanation:
Port security is a feature supported on Cisco Catalyst switches that restricts a switch port to a specific set or number of MAC addresses. Those addresses can be learned dynamically or configured statically. The port will then provide access to frames from only those addresses. If, however, the number of addresses is limited to four but no specific MAC addresses are configured, the port will allow any four MAC addresses to be learned dynamically, and port access will be limited to those four dynamically learned addresses.
Port Security Implementation:

Step	Description
1.	Enables port security. Switch(config-if)#**switchport port-security**
2.	Sets a maximum number of MAC addresses that will be allowed on this port. Default is one. Switch(config-if)#**switchport port-security maximum** value
3.	Specifies which MAC addresses will be allowed on this port (optional). Switch(config-if)#**switchport port-security mac-** **address** mac-address Switch(config-if)#**switchport port-security mac-address** mac-address
4.	Defines what action an interface will take if a nonallowed MAC address attempts access. Switch(config-if)#**switchport port-security violation** {shutdown | restrict | protect}

When Switch port security rules violate different action can be applied:
1. Protect: Frames from the nonallowed address are dropped, but there is no log of the violation.
2. Restrict: Frames from the nonallowed address are dropped, a log message is created, and a Simple Network Management Protocol (SNMP) trap is sent.
3. Shutdown: If any frames are seen from a nonallowed address, the interface is

errdisabled, a log entry is made, an SNMP trap is sent, and manual intervention or errdisable recovery must be used to make the interface usable.

QUESTION 328
The following "show" command was issued on Certkiller 1:

```
Certkiller1 # show ip access-lists net_10
  Extended IP access list net_10
      10 permit ip 10.0.0.0.0.255.255.255 and

Certkiller1 # conf t
Certkiller1 (config)# vlan access-map thor 10
Certkiller1 (config-access-map)# match ip address net_10
Certkiller1 (config-access-map)# action forward
Certkiller1 (config-access-map)  exit
Certkiller1 (config)# vlan filter thor vlan-list 12-16
```

Study the exhibit carefully. What will happen to traffic within VLAN 14 with a source address of 172.16.10.5?

A. The traffic will be dropped.
B. The traffic will be forwarded to the router processor for further processing.
C. The traffic will be forwarded without further processing.
D. The traffic will be forwarded to the TCAM for further processing.
E. None of the above

Answer: A

Explanation:
VLAN maps, also known as VLAN ACLs or VACLs, can filter all traffic traversing a switch. VLAN maps can be configured on the switch to filter all packets that are routed into or out of a VLAN, or are bridged within a VLAN. VLAN maps are used strictly for security packet filtering. Unlike router ACLs, VLAN maps are not defined by direction (input or output).
To create a VLAN map and apply it to one or more VLANs, perform these steps:
1. Create the standard or extended IP ACLs or named MAC extended ACLs to be applied to the VLAN. This access-list will select the traffic that will be either forwarded or dropped by the access-map. Only traffic matching the 'permit' condition in an access-list will be passed to the access-map for further processing.
2. Enter the vlan access-map access-map-name [sequence] global configuration command to create a VLAN ACL map entry. Each access-map can have multiple entries. The order of these entries is determined by the sequence. If no sequence number is entered, access-map entries are added with sequence numbers in increments of 10.
3. In access map configuration mode, optionally enter an action forward or action drop. The default is to forward traffic. Also enter the match command to specify an IP packet or a non-IP packet (with only a known MAC address), and to match the packet against one or more ACLs (standard or extended).
4. Use the vlan filter access-map-name vlan-list vlan-list global configuration command

to apply a VLAN map to one or more VLANs. A single access-map can be used on multiple VLANs.

QUESTION 329
In order to enhance worker productivity, a Cisco wireless network has been implemented at all Certkiller locations. Which three statements regarding WLAN are true? (Select three)

A. A WLAN client that is operating in half-duplex mode will delay all clients in that WLAN.
B. The Aironet 1230 access point is an example of an access point that operates solely as a lightweight access point.
C. Ad hoc mode allows mobile clients to connect directly without an intermediate AP.
D. WLANs are designed to share the medium and can easily handle an increased demand of channel contention.
E. A lightweight AP receives control and configuration from a WLAN controller to which it is associated.
F. Another term for infrastructure mode is independent service set (IBSS).

Answer: A, C, E

Explanation:
Ad hoc mode: This mode is called Independent Basic Service Set (IBSS). Mobile clients connect directly without an intermediate access point.
The control traffic between the access point and the controller is encapsulated with the LWAPP. The control traffic is encrypted via the Advanced Encryption Standard (AES).
The data traffic between the access point and controller is also encapsulated with LWAPP. The data traffic is not encrypted. It is switched at the WLAN controller, where VLAN tagging and quality of service (QoS) are also applied.
Lightweight access points first search for a WLAN controller using LWAPP in Layer 2 mode. Then the access point searches for a WLAN in Layer 3 mode.

QUESTION 330
To save on intra-office call costs, the Certkiller network has implemented VOIP at all their locations. Which statement is true about IP telephony calls?

A. The voice carrier stream uses H.323 to set up, maintain, and tear down call endpoints.
B. A Voice over IP (VoIP) packet consists of the voice payload, IP header, TCP header, RTP header, and Layer 2 link header.
C. Call control signaling uses Real-Time Transport Protocol (RTP) packets that contain actual voice samples.
D. The sum of bandwidth necessary for each major application, including voice, video, and data, should not exceed 75 percent of the total available bandwidth for each link.
E. None of the above

Answer: D

Explanation:
Voice over Internet Protocol (VOIP) is probably the most feasible among today's technologies for data, voice and video integration. VOIP is the technology that uses Internet Protocol to transmit voice conversations over a data network such as an intranet or the Internet.
The multisite IP WAN with distributed call processing has the following design characteristics:
* Cisco CallManager or Cisco CallManager cluster at each location (10,000 users maximum per site)
* Cisco CallManager clusters are confined to a single campus and may not span the WAN
* IP WAN as the primary voice path between sites, with the PSTN as the secondary voice path
* Transparent use of the PSTN if the IP WAN is unavailable
* Cisco IOS gatekeeper for E.164 address resolution
* Cisco IOS gatekeeper for admission control to the IP WAN
* Maximum of 100 sites interconnected across the IP WAN using hub and spoke topologies
* Compressed voice calls supported across the IP WAN
* Single WAN codec supported
* DSP resources for conferencing and WAN transcoding at each site
* Voice mail and unified messaging components at each site
* Minimum bandwidth requirement for voice and data traffic is 56 kbps. For voice, interactive video, and data, the minimum requirement is 768 kbps. In each case, the bandwidth allocated to voice, video, and data should not exceed 75% of the total capacity
* Remote sites can use Cisco IOS as well as gateways based on the Skinny Gateway Protocol

---

QUESTION 331
You need to troubleshoot some problems in the Certkiller VOIP network associated with jitter. What is the cause of jitter?

A. Packet drops
B. Transmitting too many small packets
C. Variable queue delays
D. Compression
E. None of the above

Answer: C

Explanation:
Delay variation or jitter is the difference in the delay times of consecutive packets. A jitter buffer is often used to smooth out arrival times, but there are instantaneous and total limits on buffering ability. Any type of buffering used to reduce jitter directly increases

total network delay. In general, traffic requiring low latency also requires a minimum variation in latency.
As a design rule, voice networks cannot cope with more than 30 ms of jitter. Jitter in excess of 30 ms will result in degraded audio performance. Excessive jitter in a streaming video environment will result in jerky motion, loss of video quality or loss of video.

QUESTION 332
An IP phone connects a Certkiller user to a switch as shown below:

Based on the diagram shown above, which statement is true about the voice traffic coming to the switch access port that is connected to the IP phone?

A. A PC connected to a switch port via an IP phone is unaware of the presence of the phone.
B. The traffic on the voice VLAN must be tagged with 802.1p encapsulation in order to coexist on the same LAN segment with a PC.
C. To improve the quality of the voice traffic, no other devices should be attached to the IP phone.
D. The voice VLAN must be configured as a native VLAN on the switch.
E. A PC connected to a switch port via an IP phone must support a trunking encapsulation.

Answer: A

Explanation:
The new voice VLAN is called an auxiliary VLAN in the Catalyst software command-line interface (CLI). In the traditional switched world, data devices reside in a data VLAN. The new auxiliary VLAN is used to represent other types of devices collectively. Today those devices are IP phones (hence the notion of a voice VLAN), but, in the future, other types of non-data devices will also be part of the auxiliary VLAN. Just as data devices come up and reside in the native VLAN (default VLAN), IP phones come up and reside in the auxiliary VLAN, if one has been configured on the switch. When the IP phone powers up, it communicates with the switch using CDP. The switch then provides the phone with its configured VLAN ID (voice subnet), also known as the voice VLAN ID or VVID. Meanwhile, data devices continue to reside in the native VLAN (or default VLAN) of the switch. A data device VLAN (data subnet) is referred to as a port VLAN ID or PVID.

QUESTION 333
VOIP has been implemented at the main office of the Certkiller network. Which two statements are true about voice packets in a LAN? (Select two)

A. Voice carrier stream utilizes Real-Time Transport Protocol (RTP) to carry the audio/media portion of the VoIP communication.
B. Voice traffic data flow involves large volumes of large packets.
C. Because a packet loss involves a small amount of data, voice traffic is less affected by packet losses than traditional data traffic is.
D. Voice packets are encapsulated in TCP segments to allow for proper sequencing during delivery.
E. Voice packets are very sensitive to delay and jitter.

Answer: A, E

Explanation:
Two major factors affect voice quality: lost packets and delayed packets. Packet loss causes voice clipping and skips. Packet delay can cause either voice quality degradation, due to the end-to-end voice latency, or packet loss, if the delay is variable. If the end-to-end voice latency becomes too long (250 Msec, for example), the conversation begins to sound like two parties talking on a CB radio. If the delay is variable, there is a risk of jitter buffer overruns at the receiving end.

QUESTION 334
Which QoS mechanisms can you use on a converged network to improve VoIP quality? (Select three)

A. The use of a queuing method that will give VoIP traffic strict priority over other traffic.
B. The use of RTP header compression for the VoIP traffic.
C. The proper classification and marking of the traffic as close to the source as possible.
D. The use of 802.1QinQ trunking for VoIP traffic.
E. The use of WRED.

Answer: A, C, E

Explanation:
In order to optimize the quality of VOIP calls, QoS should be implemented to ensure that VOIP traffic is prioritized over other traffic types.
By providing a strict queue for VOIP traffic, you will ensure that voice calls take precedence over the other traffic types.
In order to properly provide for QoS across the network, the voice traffic should be marked to give priority as close to the source as possible. This will ensure that the traffic is prioritized end to end.
Finally, WRED (Weighted Random Early Detection) could be configured to prevent congestion. WRED can be used to selectively drop less important traffic types, instead of dropping the voice packets when links become busy.
Incorrect Answers:
B: Compression can be used to lower the bandwidth required to transmit VOIP calls, but

it will not help with improving the voice quality. In general, compression of any kind lowers the quality of VOIP.
D. The trunking method used will have no bearing on the VOIP quality.

QUESTION 335
The Certkiller is rolling out Cisco's Architecture for Voice, Video and Integrated Data (AVVID). Which of the following choices represent the fundamental intelligent network services in Cisco's AVVID? (Select all that apply.)

A. Quality of Service (QoS)
B. Intelligent platforms
C. Mobility and scalability
D. Security
E. High availability

Answer: A, C, D, E

Explanation:
By creating a robust foundation of basic connectivity and protocol implementation, Cisco AVVID Network Infrastructure addresses five primary concerns of network deployment:
1. High availability
2. Quality of service (QoS)
3. Security
4. Mobility and
5. Scalability
Reference:
http://www.cisco.com/en/US/netsol/netwarch/ns19/ns24/networking_solutions_audience_business_benefit09186

QUESTION 336
Which of the characteristics below is associated with the (QoS) Integrated Services Model?

A. QoS classified at layer 3 using IP precedence or DSCP.
B. Guaranteed rate service.
C. Implemented using FIFO queues.
D. All traffic has an equal chance of being dropped.

Answer: B

Explanation:
Cisco IOS QoS includes the following features that provide controlled load service, which is a kind of integrated service:
Resource Reservation Protocol (RSVP) can be used by applications to signal their QoS requirements to the router.
Intelligent queuing mechanisms can be used with RSVP to provide the following kinds

of services:
Ø Guaranteed Rate Service, which allows applications to reserve bandwidth to meet their requirements. For example, a Voice over IP (VoIP) application can reserve 32 Mbps end to end using this kind of service. Cisco IOS QoS uses weighted fair queuing (WFQ) with RSVP to provide this kind of service.
Ø Controlled Load Service, which allows applications to have low delay and high throughput even during times of congestion. For example, adaptive real-time applications such as playback of a recorded conference can use this kind of service. Cisco IOS QoS uses RSVP with Weighted Random Early Detection (WRED) to provide this kind of service.
Reference:
http://www.cisco.com/en/US/products/sw/iosswrel/ps1831/products_configuration_guide_chapter09186a008007

QUESTION 337
Which three QoS mechanisms can be configured to improve VoIP quality on a converged network? Select three.

A. The use of a queuing method that will give VoIP traffic strict priority over other traffic
B. The use of RTP header compression for the VoIP traffic.
C. The proper classification and marking of the traffic as close to the source as possible
D. The use of 802.1QinQ trunking for VoIP traffic
E. The use of WRED for the VoIP traffic

Answer: A, C, E

Explanation:
In order to optimize the quality of VOIP calls, QoS should be implemented to ensure that VOIP traffic is prioritized over other traffic types.
By providing a strict queue for VOIP traffic, you will ensure that voice calls take precedence over the other traffic types.
In order to properly provide for QoS across the network, the voice traffic should be marked to give priority as close to the source as possible. This will ensure that the traffic is prioritized end to end.
Finally, WRED (Weighted Random Early Detection) could be configured to prevent congestion. WRED can be used to selectively drop less important traffic types, instead of dropping the voice packets when links become busy.
Incorrect Answers:
B: Compression can be used to lower the bandwidth required to transmit VOIP calls, but it will not help with improving the voice quality. In general, compression of any kind lowers the quality of VOIP.
D. The trunking method used will have no bearing on the VOIP quality.

QUESTION 338
A Cisco IP phone connects to a host PC and a Certkiller switch as shown below:

Study the exhibit carefully. Which statement is true about a voice VLAN?

A. Physically the voice network and the data network are separate.
B. The voice traffic will normally be on a different IP subnet than will the data traffic.
C. End user intervention is necessary to place the phone into the proper VLAN.
D. The same security policy should be implemented for both voice and data traffic.
E. The data VLAN must be configured as the native VLAN.

Answer: B

Explanation:
The new voice VLAN is called an auxiliary VLAN in the Catalyst software command-line interface (CLI). In the traditional switched world, data devices reside in a data VLAN. The new auxiliary VLAN is used to represent other types of devices collectively. Today those devices are IP phones (hence the notion of a voice VLAN), but, in the future, other types of non-data devices will also be part of the auxiliary VLAN. Just as data devices come up and reside in the native VLAN (default VLAN), IP phones come up and reside in the auxiliary VLAN, if one has been configured on the switch. When the IP phone powers up, it communicates with the switch using CDP. The switch then provides the phone with its configured VLAN ID (voice subnet), also known as the voice VLAN ID or VVID. Meanwhile, data devices continue to reside in the native VLAN (or default VLAN) of the switch. A data device VLAN (data subnet) is referred to as a port VLAN ID or PVID.

QUESTION 339
An IP phone connects a Certkiller user to the network as shown below:

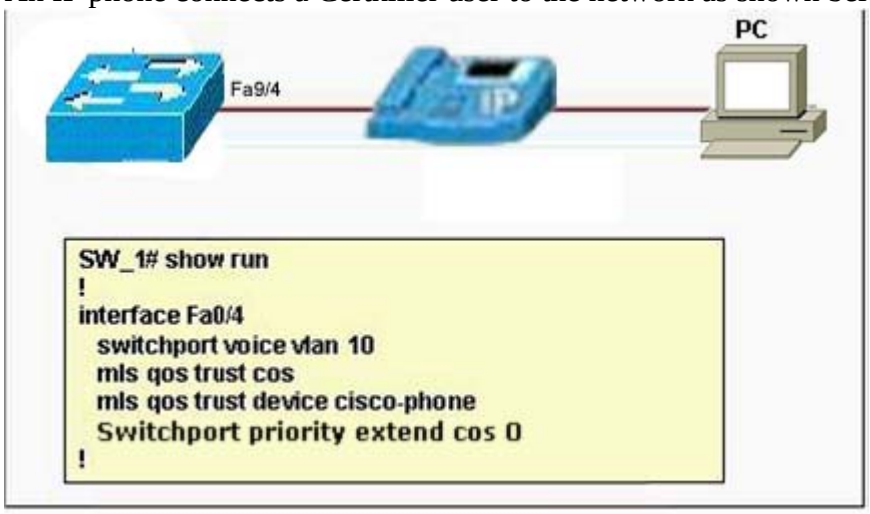

Refer to the exhibit. A workstation PC is connected to the Cisco IP phone access

port. Based on the configuration in the exhibit above, how will the traffic be managed?

A. The IP phone access port will override the priority of the frames received from the PC.
B. The IP phone access port will trust the priority of the frames received from the PC.
C. The switch port Fa0/4 will override the priority of the frames received from the PC.
D. The switch port Fa0/4 will trust the priority for the frames received from the PC.
E. None of the above

Answer: A

Explanation:
The PC connected to the phone, however, should normally be untrusted and have all inbound CoS
values set to 0. This is mentioned here to show how trust boundaries also exist at any connected IP
Phones.
Example:
interface fastethernet 0/1
switchport voice vlan 200
switchport priority extend cos 0
A switch instructs an attached IP Phone through CDP messages as to how it should extend QoS trust
to its own user data switch port. To configure the trust extension, use the following interface
configuration command:
Switch(config-if)# switchport priority extend {cos value | trust}
Normally, the QoS information from a PC connected to an IP Phone should not be trusted. This is
because the PC's applications might try to spoof CoS or Differentiated Services Code Point (DSCP)
settings to gain premium network service. In this case, use the cos keyword so that the CoS bits are
overwritten to value by the IP Phone as packets are forwarded to the switch. If CoS values from the
PC cannot be trusted, they should be overwritten to a value of 0.

QUESTION 340
You are a network administrator of a large investor relations company that uses a switched network to carry both data and IP telephony services. Why should you carry voice traffic on a separate VLAN?

A. IP phones require inline power and must be in separate VLAN to receive inline power.
B. IP telephony applications require prioritization over other traffic as they are more delay sensitive.
C. IP phones can only receive IP addresses through DHCP if they are in separate VLAN.

D. The CDP frames from the IP phone can only be recognized by the switch if the phone is in an auxiliary vlan.

Answer: B

Explanation:
Voice conversations don't take up a lot of bandwidth, but the bandwidth they do is very delicate. If anything happens with the connection or the integrity of the data transfer in either direction the conversation won't seam natural. To ensure the highest degree of integrity you should put voice traffic on its own separate VLAN and give that VLAN the highest priority.

QUESTION 341
A portion of the Certkiller VOIP LAN is shown below:

Configuration exhibit:

```
CertkillerA#show run
!
Interface Fa0/4
 switchport voice vlan 10
 mls qos trust cos
 mls qos trust device cisco-phone
 switchport priority extend cos 0
!
```

Study the exhibits carefully. A workstation PC Certkiller B is connected to the Cisco IP phone access port. Based on the configuration shown in the exhibit above, how will the traffic be managed?

A. The switch port Fa0/4 will trust the priority for the frames received from host Certkiller B.
B. The IP phone access port will override the priority of the frames received from host Certkiller B.
C. The IP phone access port will trust the priority of the frames received from host Certkiller B.
D. The switch port Fa0/4 will override the priority of the frames received from host Certkiller B.
E. None of the above.

Answer: B

Explanation:
When a Cisco IP Phone is connected to a switch port, think of the phone as another switch (which it is). If you install the phone as a part of your network, you can probably trust the QoS information relayed by the phone.

However, remember that the phone also has two sources of data:
1. The VoIP packets native to the phone-The phone can control precisely what QoS information is included in the voice packets because it produces those packets.
* The user PC data switch port-Packets from the PC data port are generated elsewhere, so the QoS information can not necessarily be trusted to be correct or fair.
A switch instructs an attached IP Phone through CDP messages as to how it should extend QoS trust to its own user data switch port. To configure the trust extension, use the following interface configuration command:
Switch(config-if)# switchport priority extend {cos value | trust}
Normally, the QoS information from a PC connected to an IP Phone should not be trusted. This is because the PC's applications might try to spoof CoS or Differentiated Services Code Point (DSCP) settings to gain premium network service. In this case, use the cos keyword so that the CoS bits are overwritten to value by the IP Phone as packets are forwarded to the switch. If CoS values from the PC cannot be trusted, they should be overwritten to a value of 0.
In some cases, the PC might be running trusted applications that are allowed to request specific QoS or levels of service. Here, the IP Phone can extend complete QoS trust to the PC, allowing the CoS bits to be forwarded through the phone unmodified. This is done with the trust keyword.

QUESTION 342
You need to configure a new Cisco router to be installed in the Certkiller VOIP network. Which three interface commands will configure the switch port to support a connected Cisco phone and to trust the CoS values received on the port if CDP discovers that a Cisco phone is attached? (Select three)

A. switchport voice vlan vlan-id
B. mls qos trust device cisco-phone
C. switchport priority extend cos_value
D. mls qos trust cos
E. mls qos trust override cos

Answer: A, B, D

Explanation:
1. To configure the IP Phone uplink, just configure the switch port where it connects. The switch instructs the phone to follow the mode that is selected. In addition, the switch port does not need any special trunking configuration commands if a trunk is wanted. If an 802.1Q trunk is needed, a special-case trunk is negotiated by Dynamic Trunking Protocol (DTP) and CDP. Use the following interface configuration command to select the voice VLAN mode that will be used:
Switch(config-if)# switchport voice vlan { vlan-id | dot1p | untagged | none}
2. mlsqos trust [cos] : Configure the port trust state.
By default, the port is not trusted. All traffic is sent through one egress queue. Use the cos keyword to classify ingress packets with the packet CoS values. The egress queue assigned to the packet is based on the packet CoS value

3. mls qos trust device cisco-phone : Configure the Cisco IP Phone as a trusted device on the interface.

QUESTION 343
Two IP phones connect to switch Certkiller A as shown below:

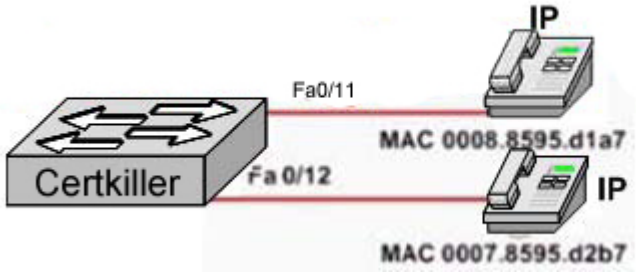

Configuration exhibit:

```
CertkillerA# show running-config
<output omitted>
mis qos map cos dscp 0.8 16 45 45 45 48 56
mls qos
<output omitted>
interface FastEthernet0/11
description IP-Phone, workstation
switchport access vlan 40
switchport mode access
switchport voice vlan dot1p
switchport priority extend cos 4
switchport port-security
switchport port-security maximum 2
switchport port-security mac-address 0008.8595.d1a7
switchport port-security mac-address 0800.20ac.24b8
mls qos trust cos
Spanning-tree ports
!
interface fastEthernet/12
description IP-Phone, work station
switchport access vlan 40
switchport mode access
switchport voice vlan dot1p
switchport priority extend cos 4
switchport port-security
switchport port-security maximum 2
switchport port-security mac-address 0007.8595.d2b7
switchport port-security mac-address 0080.20ac.24b9
mls qos trust cos
spanning-tree portfast
```

Based on the information shown above, what statement is true about the configuration on switch Certkiller A?

A. Untagged Port VLAN ID (PVID) frames will carry voice traffic on VLAN 40.
B. The configuration overrides the Quality of Service value in packets entering ports Fa0/11 and Fa0/12 with a value of 45.
C. Two IP phones with the MAC addresses of 0008.8595.d1a7 and 0007.8595.d2b7 are connected to Certkiller A ports Fa0/11 and Fa0/12, respectively.
D. The configuration overrides 802.1p priorities on packets entering ports Fa0/11 and Fa0/12 with a value of 48.

E. The configuration establishes policed DSCP on ports Fa0/11 and Fa0/12 with values ranging from 8 to 56.
F. Security violation shutdown mode has been activated for ports Fa0/11 and Fa0/12.
G. None of the above.

Answer: C

Explanation:
Port security is a feature supported on Cisco Catalyst switches that restricts a switch port to a specific set or number of MAC addresses. Those addresses can be learned dynamically or configured statically. The port will then provide access to frames from only those addresses. If, however, the number of addresses is limited to four but no specific MAC addresses are configured, the port will allow any four MAC addresses to be learned dynamically, and port access will be limited to those four dynamically learned addresses.
Port Security Implementation:

Step	Description
1.	Enables port security. Switch(config-if)#**switchport port-security**
2.	Sets a maximum number of MAC addresses that will be allowed on this port. Default is one. Switch(config-if)#**switchport port-security maximum** value
3.	Specifies which MAC addresses will be allowed on this port (optional). Switch(config-if)#**switchport port-security mac-** **address** mac-address Switch(config-if)#**switchport port-security mac-address** mac-address
4.	Defines what action an interface will take if a nonallowed MAC address attempts access. Switch(config-if)#**switchport port-security violation** {shutdown | restrict | protect}

QUESTION 344
A Certkiller switch is configured as shown below:

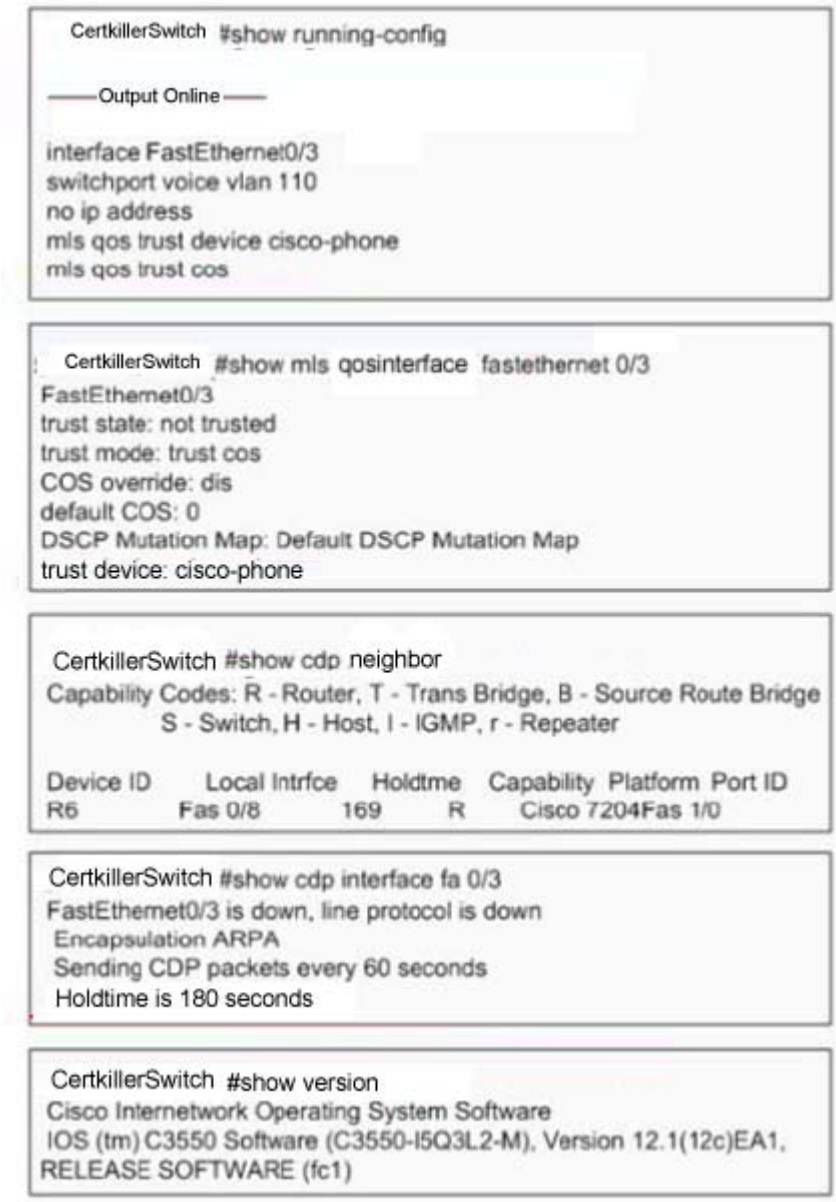

Study the exhibits shown above carefully. Based on the information provided, why does the trust state of interface FastEthernet 0/3 show "not trusted"?

A. DSCP map needs to be configured for VOIP.
B. ToS has been misconfigured.
C. The command mls qos needs to be turned on in global configuration mode.
D. ToS has not been configured.
E. There is not a Cisco Phone attached to the interface.
F. None of the above.

Answer: E

Explanation:
CDP is a device discovery protocol that runs over Layer 2 (the data link layer) on all

Cisco-manufactured devices (routers, bridges, access servers, and switches) and allows network management applications to discover Cisco devices that are neighbors of already known devices. With CDP, network management applications can learn the device type and the Simple Network Management Protocol (SNMP) agent address of neighboring devices running lower-layer, transparent protocols. This feature enables applications to send SNMP queries to neighboring devices.
CDP runs on all media that support Subnetwork Access Protocol (SNAP). Because CDP runs over the data-link layer only, two systems that support different network-layer protocols can learn about each other.
Communication between Switch and IP Phone is performed by CDP protocol. There is no CDP neighbor and trusted state also no trusted.

QUESTION 345
An IP phone connects a Certkiller user to a POE switch as shown below:

Study the exhibit carefully. Which statement is true when voice traffic is forwarded on the same VLAN used by the data traffic?

A. The voice traffic cannot be forwarded to the distribution layer.
B. Quality of service cannot be applied for the voice traffic.
C. The voice traffic cannot use 802.1P priority tagging.
D. Port security cannot be enabled on the switch that is attached to the IP phone.
E. None of the above

Answer: C

Explanation:
You can configure a port connected to the Cisco IP Phone to send CDP packets to the phone to configure the way in which the phone sends voice traffic. The phone can carry voice traffic in 802.1Q frames for a specified voice VLAN with a Layer 2 CoS value. It can use 802.1P priority tagging to give voice traffic a higher priority and forward all voice traffic through the native (access) VLAN. The IP phone can also send untagged voice traffic or use its own configuration to send voice traffic in the access VLAN. In all configurations, the voice traffic carries a Layer 3 IP precedence value (the default is 5).
Beginning in privileged EXEC mode, follow these steps to configure voice traffic on a port:

Command	Purpose
configure terminal	Enter global configuration mode.
Interface *interface-id*	Enter interface configuration mode, and specify the interface connected to the IP phone.
mls qos trust cos	Configure the interface to classify ingress traffic packets by using the packet CoS value. For untagged packets, the port default CoS value is used. **Note** Before configuring the port trust state, you must first globally enable QoS by using the **mls qos** global configuration command.
Switchport voice vlan {*vlan-id* | **dot1p** | **none** | **untagged**)	Configure how the Cisco IP Phone carries voice traffic: • *vlan-id*—Configure the Cisco IP Phone to forward all voice traffic through the specified VLAN. By default, the Cisco IP Phone forwards the voice traffic with an 802.1Q priority of 5. Valid VLAN IDs are from 1 to 4094. • **dot1p**—Configure the Cisco IP Phone to use 802.1P priority tagging for voice traffic and to use the default native VLAN (VLAN 0) to carry all traffic. By default, the Cisco IP Phone forwards the voice traffic with an 802.1P priority of 5. • **none**—Allow the IP phone to use its own configuration to send untagged voice traffic. • **untagged**—Configure the phone to send untagged voice traffic.
End	Return to privileged EXEC mode.

show interfaces *interface-id* **switchport** or **show running-config interface** *interface-id*	Verify your voice VLAN entries. Verify your QoS and voice VLAN entries.
copy running-config startup-config	(Optional) Save your entries in the configuration file.

Reference:
http://www.cisco.com/en/US/products/hw/switches/ps5206/products_configuration_guide_chapter09186a00801a

QUESTION 346
Refer to the following network topology exhibit:

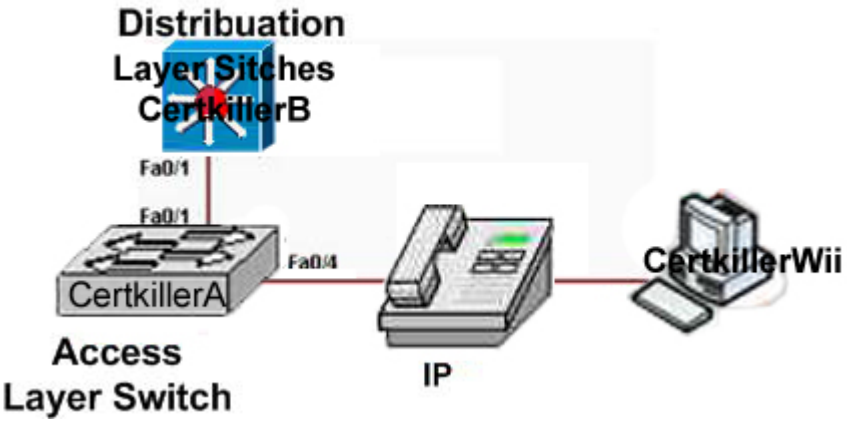

Certkiller B configuration exhibit:

```
CertkillerB# show run

interface Fa0/1
  switchport trunk encapsulation dot1q
  switchport mode trunk
  switchport voice vlan 10
  mls qos
  mls qos trust cos
```

Study the exhibits above carefully. A trunk link is connected between switch Certkiller A and switch Certkiller B. Based on the output shown above, how would the traffic coming from the switch Certkiller A be managed?

A. The trunk port Fa0/1 on switch Certkiller B will trust all CoS values on the frames coming from port Fa0/1 on Certkiller A.
B. The trunk port Fa0/1 on switch Certkiller A will trust all CoS values on the frames coming from the IP phone.
C. The trunk port Fa0/1 on switch Certkiller A will trust all CoS values on the frames received on the IP phone.
D. The trunk port Fa0/1 on switch Certkiller B will trust all CoS values on the frames received on the Certkiller A switch port Fa0/4.
E. The trunk port Fa0/1 on switch Certkiller B will trust all CoS values on the frames received on the IP phone port.
F. None of the above.

Answer: A

QUESTION 347
Jitter is causing problems with the VOIP application in the Certkiller network.
What causes network jitter?

A. Variable queue delays
B. Packet drops

C. Transmitting too many small packets
D. Compression

Answer: A
Delay variation or jitter is the difference in the delay times of consecutive packets. A jitter buffer is often used to smooth out arrival times, but there are instantaneous and total limits on buffering ability. Any type of buffering used to reduce jitter directly increases total network delay. In general, traffic requiring low latency also requires a minimum variation in latency.
Note: Jitter in Packet Voice Networks:
Jitter is defined as a variation in the delay of received packets. At the sending side, packets are sent in a continuous stream with the packets being spaced evenly apart. Due to network congestion, improper queuing, or configuration errors, this steady stream can become lumpy, or the delay between each packet can vary instead of remaining constant.

QUESTION 348
VOIP is being implemented on the Certkiller network. In a properly designed network, what is the maximum amount of time a voice package should spend crossing a network?

A. 90 milliseconds
B. 120 milliseconds
C. 150 milliseconds
D. 240 milliseconds

Answer: C

Explanation:
Delay is the time it takes for VoIP packets to travel between two endpoints and you should design networks to minimize this delay. However, because of the speed of network links and the processing power of intermediate devices, some delay is expected. The human ear normally accepts up to about 150 milliseconds (ms) of delay without noticing problems (the ITU standard recommends no more than 150 ms of one-way delay).
Reference:
http://www.cisco.com/en/US/products/sw/iosswrel/ps5014/products_feature_guide09186a00800880e7.html

---

QUESTION 349
VOIP is being implemented in the Certkiller network and you need to assess the need for QoS. Which of the following network problems would indicate a need to implement QoS features? (Select three)

A. Mis-routed packets
B. Excess jitter
C. Delay of critical traffic
D. Packet loss due to congestion

E. Data link layer broadcast storms
F. FTP connections unsuccessful

Answer: B, C, D

Explanation:
Loss, jitter, and delay are the three reasons for implementing QoS features on modern networks. Loss is when a packet disappears on a network. Jitter is a timing mismatch between two way traffic, and delay is when a packet takes too long to get somewhere.
Incorrect Answers:
A: This would indicate a routing problem, or packets being "black-holed." QoS would not help in this situation.
E: Broadcast storms indicate a problem on a LAN segment, such as a babbling host, too many hosts, a segment that is too large, a bad application, etc. QoS would not help in this situation.
F. If only FTP sessions were having issues, then the FTP application or FTP server should be corrected. Normally, FTP sessions are not delay sensitive due to the re-transmission nature of TCP and do not require QoS.

QUESTION 350
The Certkiller network is displayed in the following topology exhibit:

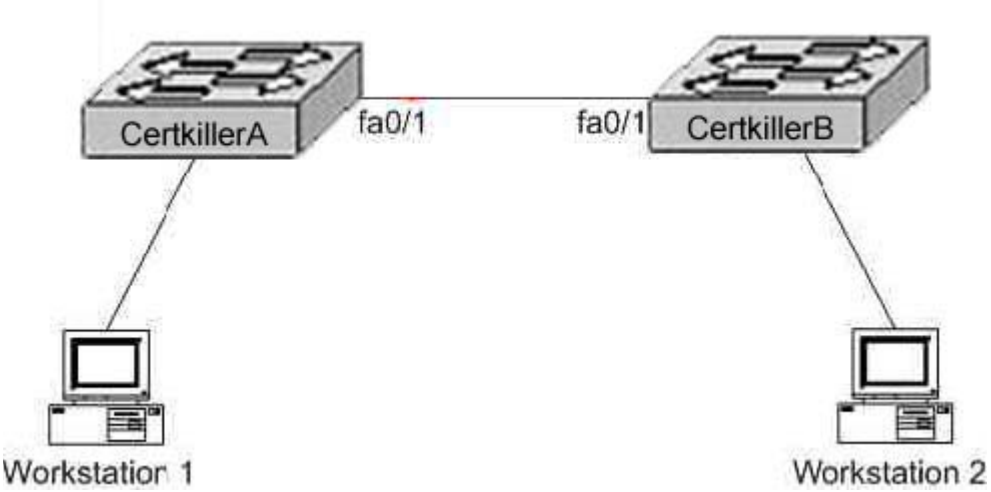

Workstation 1 traffic is set for cos 5 and the switch Certkiller A sends workstation 1 traffic to the switch Certkiller B. However, not all of the traffic from Switch Certkiller A is from workstation 1.
Switch Certkiller A configurations Switch Certkiller B Configuration:
mlsqos mls qos
interfacefa0/1 interface fa0/1
switchportmode trunk switchport trunk mode
switchporttrunk encapsulation dot 1q switchport trunk encapsulation dot 1q
switchporttrunk native vlan 1 switchport trunk native vlan 1
Frames from Workstation 1 are given the rightful priority through Switch Certkiller A, but Switch Certkiller B doesn't reciprocate, and treats Workstation 1 frames as if they have no precedence. Which of the following actions will prioritize

traffic from Workstation 1?

A. Configure qos all command under Switch Certkiller B fa0/1 interface.
B. Configure mls qos trust cos command under Switch Certkiller B fa0/1 interface.
C. Configure mls qos trust cos 5 command under Switch Certkiller B fa0/1 interface.
D. Configure qos cos 5 command under Switch Certkiller B fa0/1 interface.
E. Configure mls qos trust cos command under Switch Certkiller A fa0/1 interface.
F. Configure qos cos 5 command under Switch Certkiller A fa0/1 interface.

Answer: B

Explanation:
The default action is for a switch with QoS features activated not to trust edge devices and any frames that enter the switch have their CoS re-written to the lowest priority of zero. If the edge device can be trusted, this default behaviour must be overridden and the access switch must be configured to switch the frame, leaving the CoS bits untouched. The trust is configured on the switch port using the command:
switch(config-if)#mls qos trust cos

QUESTION 351
Exhibit:

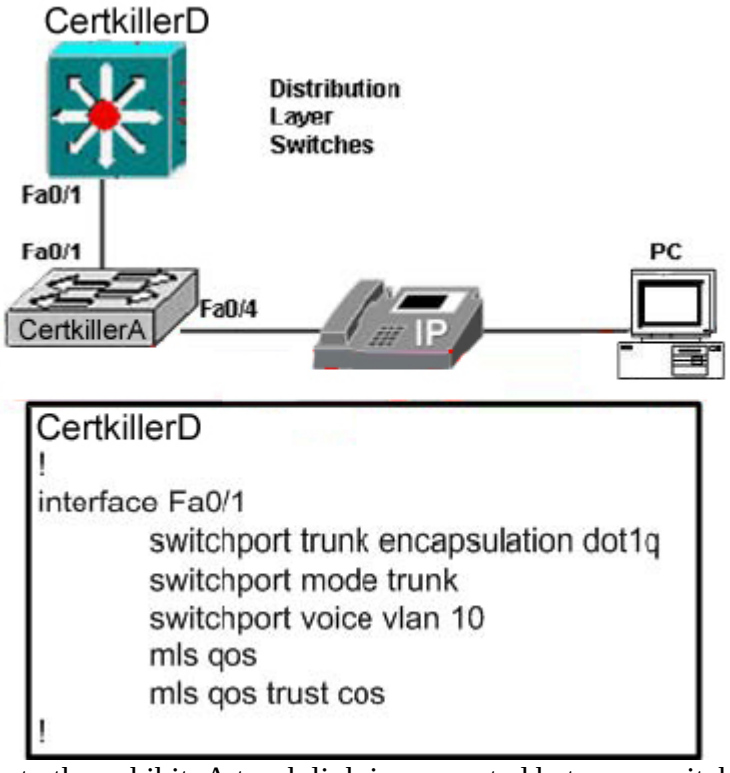

Refer to the exhibit. A trunk link is connected between switch Certkiller A and switch Certkiller D. Based on the configuration shown in the exhibit, how would the traffic coming from the switch Certkiller A be managed?

A. The trunk port Fa0/1 on switch Certkiller A will trust all CoS values on the frames

coming from the IP phone.
B. The trunk port Fa0/1 on switch Certkiller A will trust all CoS values on the frames received on the IP phone.
C. The trunk port Fa0/1 on switch Certkiller D will trust all CoS values on the frames coming from port Fa0/1 on Certkiller A.
D. The trunk port Fa0/1 on switch Certkiller D will trust all CoS values on the frames received on the Certkiller A switch port Fa0/4.
E. The trunk port Fa0/1 on switch Certkiller D will trust all CoS values on the frames received on the IP phone port.

Answer: C

Explanation:
To enable to QOS, you should enter the msl qos command in global configuration mode.
When inbound packets are accepted into a switch, the switch can be selective about which (if any)
of each packet's QoS information will be trusted. If the packets originate from a trusted source, the
QoS information can be safely trusted, too. Usually, it is a best practice to configure switches at
the edge of a trusted QoS domain to verify or overwrite any QoS information that comes into the
domain. This way, any other switch or router within the domain can blindly trust QoS information
that is seen.
You can configure QoS trust in two ways:
1. Per-interface
2. As part of a QoS policy on specific types of traffic
The per-interface trust is described in the next section. Policy trust is described as part of the section,
"Defining a QoS Policy."
Trust QoS on an Interface
On each interface where consistent QoS trust is to be defined, use the following interface configuration command:
Switch(config-if)# mls qos trust {cos | dscp | ip-precedence}
Applying QoS Trust 407
Here, one of the following values can be trusted and used internally as the switch makes forwarding
decisions:
1. The inbound CoS, which is taken from trunking tags
2. DSCP, which is taken from the inbound IP packet headers
3. IP Precedence, which is also taken from the inbound IP packet headers

---

Certkiller .com Scenario

Network topology exhibit:

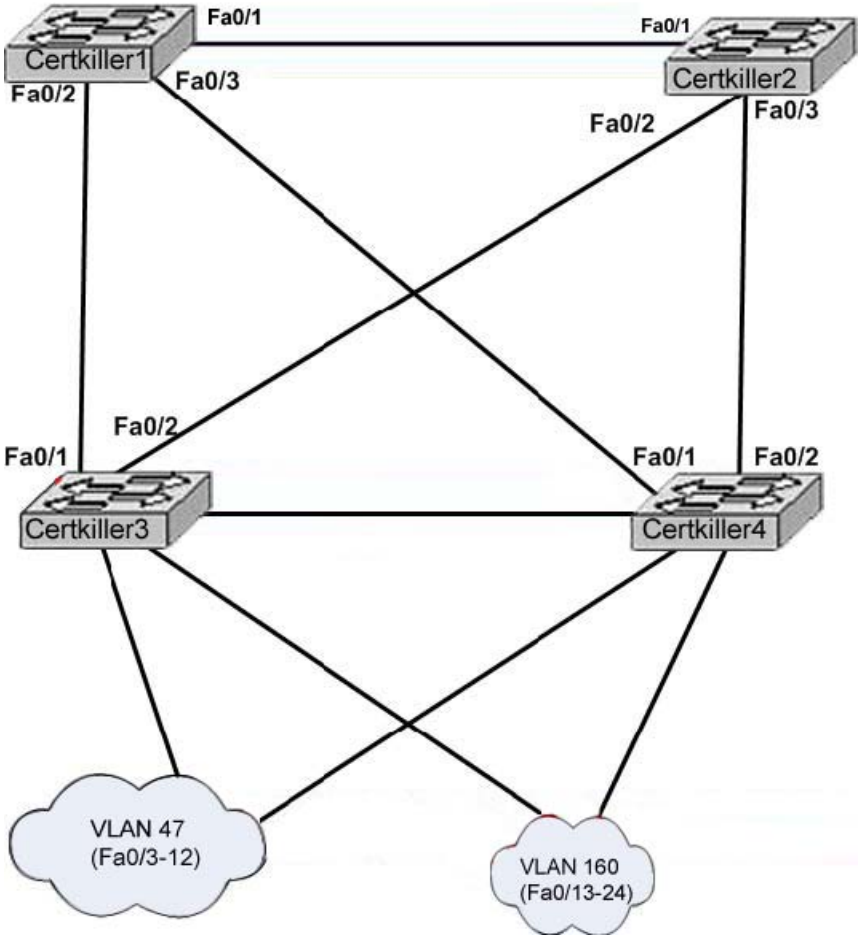

Certkiller 3 configuration exhibit:

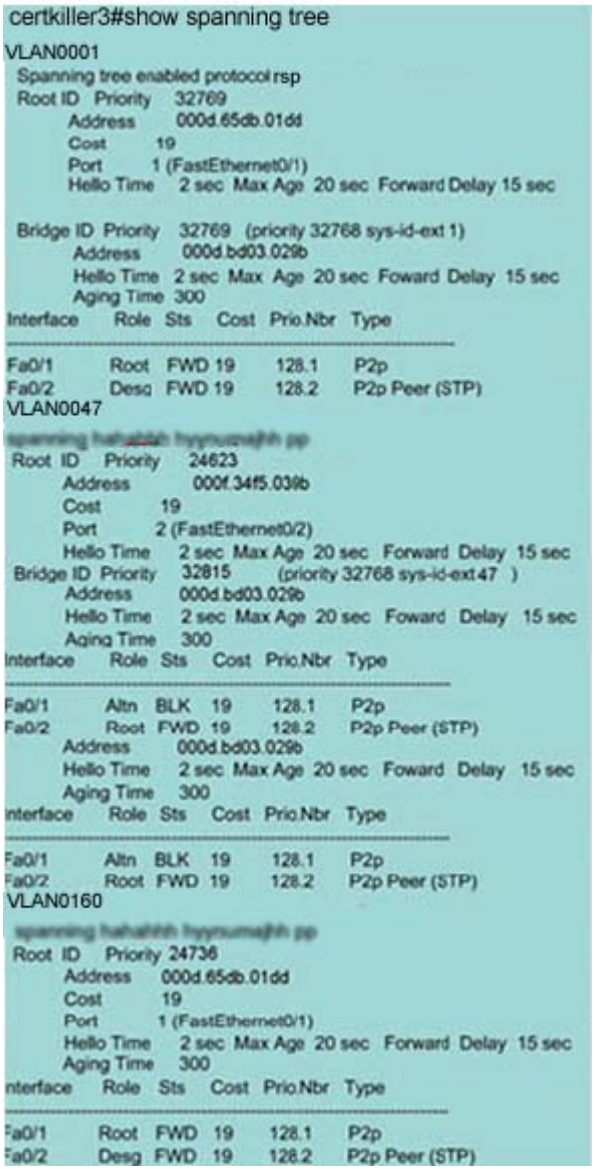

You work as a network technician at Certkiller .com. Certkiller .com is a large international company with offices in the US and in Europe. You work at the New York Office. Recently a new switch block with multiple VLANs configured were added to the Certkiller .com Barcelona LAN.

Your boss at Certkiller .com, Miss Certkiller, tells you that the personal at the Barcelona office failed to document the spanning tree topology during installation. She gives you a plane ticket to Barcelona and orders to provide a baseline of configuration information in regard to the switch block spanning tree topology. Two days later you arrive at the Certkiller .com Barcelona office. You study the network topology carefully, see exhibit. Then you connect to the Certkiller 3 switch; issue the show spanning tree command. Please refer to the exhibit for the output. You are then required to answer the scenario questions using the information that is available.

Certkiller .com (5 Questions)

QUESTION 352
Note: Please refer to the Certkiller .com scenario.
Which spanning Tree Protocol has been implemented on switch Certkiller 2?

A. STP/IEEE 802.1D
B. MSTP/IEEE 802.1s
C. PVST+
D. PVRST
E. None of the above

Answer: C

QUESTION 353
Note: Please refer to the Certkiller .com scenario.
Which bridge ID belongs to switch Certkiller 2?

A. 32928 000d bd33 029b
B. 24623 000f 34f5 039b
C. 32928 000d bd03 029b
D. 32768 000d bd33 029b
E. 32769 000d 65db 01dd
F. 32815 000d bd03 029b
G. None of the above

Answer: B

QUESTION 354
Note: Please refer to the Certkiller .com scenario.
Which port role has interface Fa0/2 of switch Certkiller 1 adopted for VLAN 47?

A. Root port
B. Nondesigned port
C. Designated port
D. Backup port
E. Alternate port
F. None of the above

Answer: C

QUESTION 355
Note: Please refer to the Certkiller .com scenario.
Which port state is interface Fa0/2 of switch Certkiller 2 in for VLANs 1 and 160?

A. listening
B. learning
C. disabled
D. blocking
E. forwarding
F. discarding
G. none

Answer: D

QUESTION 356
Note: Please refer to the Certkiller .com scenario.
Which bridge ID belongs to switch Certkiller 1?

A. 32928 000d bd33 029b
B. 24623 000f 34f5 039b
C. 32928 000d bd03 029b
D. 32768 000d bd33 029b
E. 32769 000d 65db 01dd
F. 32815 000d bd03 029b
G. None of the above

Answer:

Mixed Questions (37 Questions)

QUESTION 357
Which three statements are true about the voice VLAN feature on a Catalyst 2950 switch? Select three.

A. The default CoS value for incoming traffic is set up to 0.
B. The CoS value is trusted for 802.1p or 802.1q tagged traffic.
C. PortFast is automatically disabled when a voice VLAN is configured.
D. The voice VLAN feature is disabled by default.
E. The IP phone accepts the priority of all tagged and untagged traffic and sets the CoS value to 4.
F. When the voice VLAN feature is enabled, all untagged traffic is sent according to the default CoS priority of the port

Answer: A,B,F

QUESTION 358
What does the globacl command udld enable accomplish?

A. enables all fiber-optic LAN ports for Unidirectionla LINK Detection (UDLD)

B. enables all copper media LAN ports Uniderectional Link Detection (UDLD)
C. overrides the default UDLS setting for all ports
D. globally enables all ports on the device for Unidirectional Link Detection (UDLS)

Answer: A

QUESTION 359
What two steps can be taken to help prevent VLAN hopping? Select two.

A. Enable BPD guard
B. Disable CDP on ports where it is not necessary
C. Place unused ports in a common unrouted VLAN
D. Prevent automatic trunk configuration
E. Implement port security

Answer: C,D

QUESTION 360
Exhibit:

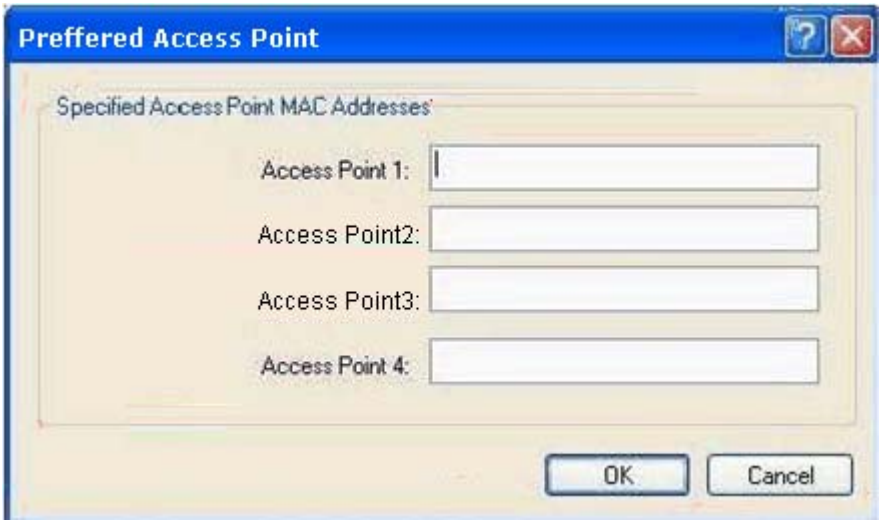

What can be said about the configuration of access Point MAC addresses on the wireless client?

A. The MAC address should consist of 10 hexadecimal characters with every two hex characters separated by hyphens.
B. Each access point MAC address that is specified must have a separate SSID configured on the GENERAL configuration tab.
C. The MAC address should consist of 16 hexadecimal characters with every two four characters separated by hyphens.
D. If the wireless client is out of range of the specified access point or point it will not associate with other access points.
E. Each access point MAC address that is specified must have the same SSID configured on the GENERAL configuration tab.

F. If the wireless client is out of range of the specified access point or point it can associate with other access points.

Answer: F

QUESTION 361
Exhibit:

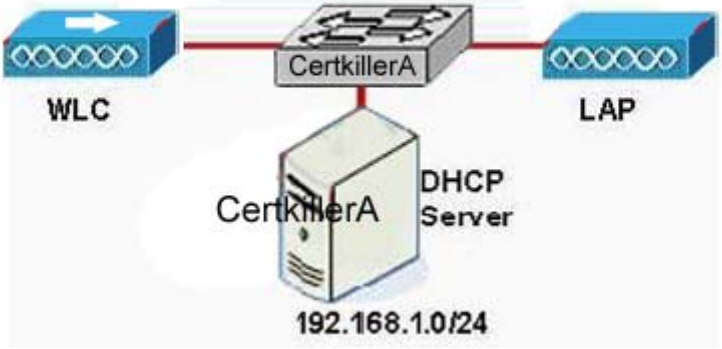

You work as a network technichian at Certkiller .com. Please study the exhibit carefully. The LAP (lightweight access point) attempts to register to a the WLC (Wireless LAN Controller). What kind of message is transmitted?

A. The lightweight access point will send Layer 2 and Layer 3 Lightweight Access Point (LWAPP) mode discovery request messages at the same time.
B. The lightweight access point will send Layer 3 Lightweight Access Point (LWAPP) mode discovery request messages only.
C. The lightweight access point will send Layer 2 Lightweight Access Point (LWAPP) mode discovery request messages. If the attempt fails, the LAP will try Layer 3 LWAPP WLC discovery.
D. The lightweight access point will send Layer 2 Lightweight Access Point (LWAPP) mode discovery request messages only.

Answer: C

QUESTION 362
Exhibit:

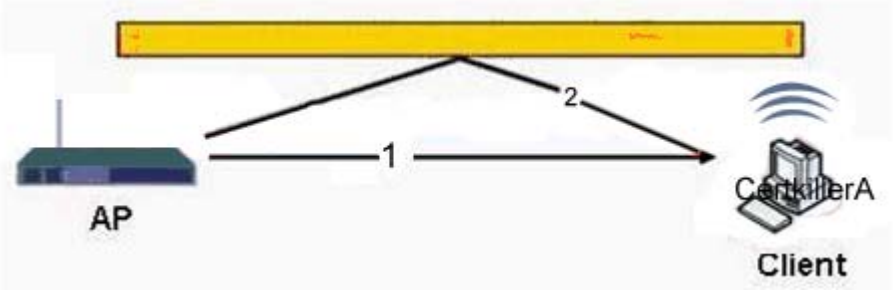

You work as a network technichian at Certkiller .com. Please study the exhibit carefully. The this scenario the signal transmitted from the AP is reflected off a wall resulting in multipath interference at the client end.

Which of the following statements is true?

A. The transmitted signal from the AP arrives at the client at slightly different times resulting in phase shifting.
B. Multipath interference is solved by using dual antennas.
C. If signal 2 is close to 360 degrees out of phase with signal 1, the result is essentially zero signal or a dead spot in the WLAN.
D. Multipath interference is less of an issue when using a DSSS technology because multipath is frequency selective.
E. If signal 1 is in phase with signal 2, the result is essentially zero signal or a dead spot in the WLAN.

Answer: A

QUESTION 363
You work as a network technichian at Certkiller .com. A client is searching for an access point (AP). What is the correct process order that the client and access point goes through to create a connection?

A. association request/response, probe request/response, authentication request/response
B. association request/response, authentication request/response, probe request/response
C. probe request/response, authentication request/response, association request/response
D. probe request/response, association request/response, authentication request/response

Answer: C

QUESTION 364
Exhibit:

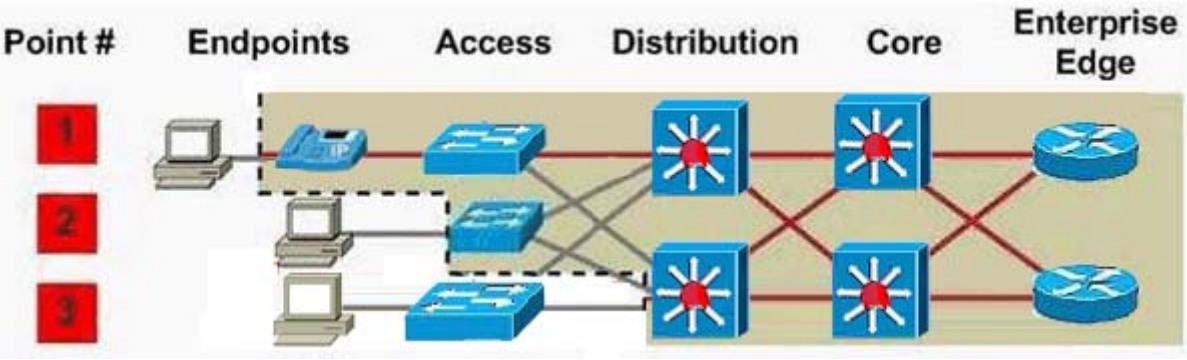

You work as a network technichian at Certkiller .com. Please study the exhibit carefully. Which statement is true about where trust boundaries should be established in a network?

A. Endpoint 2 and 3 are optimal places to establish a trust boundary. Endpoint 1 is an acceptable place to establish a trust boundary
B. Endpoint 2 is the optimal place to establish a trust boundary. Endpoints 1 and 3 are acceptable places to establish a trust boundary.
C. Endpoint 2 is the only acceptable place to establish a trust boundary.

D. Endpoint 1 and 32are optimal places to establish a trust boundary. Endpoint 1 is an acceptable place to establish a trust boundary
E. Endpoint 1 is the only acceptable place to establish a trust boundary.
F. Endpoint 1 is the optimal place to establish a trust boundary. Endpoints 1 and 3 are acceptable places to establish a trust boundary.

Answer: A

QUESTION 365
You work as a network technichian at Certkiller .com. You need to configure DHCP snooping on a switch. Which three steps are required? Select 3.

A. Configure the switch to insert and remove DHCP relay information (option-82 field) in forwarded DHCP request messages.
B. Configure DHCP snooping globally.
C. Configure the switch as a DHCP server.
D. Configure DHCP snooping on an interface.
E. Configure all interfaces as DHCP snooping trusted interfaces.
F. Configure DHCP snooping on a VLAN or range of VLANs.

Answer: B,D,F

QUESTION 366
Exhibit:

You work as a network technichian at Certkiller .com. Please study the exhibit carefully. You are required to automatically configure quality of service (QoS) for voice IP (VoIP) within a QoS domain?

A. mls qos trust
B. switchport priority extend cos 7
C. mls trust qos
D. auto qos voip cisco-phone
E. switchport priority extend trust
F. switchport priority extend

Answer: D

QUESTION 367
You work as a network technichian at Certkiller .com. Your boss, Mrs. Certkiller, is interested in GLBP. In particular she wants to know the multicast address.
What should you tell her?

A. 224.0.0.100
B. 224.0.0.1
C. 224.0.0.102
D. 224.0.0.10
E. 224.0.0.101

Answer: C

QUESTION 368
Exhibit:
iparp inspection vlan 10-12, 15
You work as a network technichian at Certkiller .com. Please study the exhibit carefully.
What is the purpose of the global configuration command in the exhibit?

A. discards ARP packets with invalid IP-to-MAC address bindings on trusted ports
B. validates outgoing ARP requests for interfaces configured on VLAN 10, 11, 12, or 15
C. intercepts, logs, and discards ARP packets with invalid IP-to-MAC address bindings
D. intercepts all ARP requests and responses on trusted ports

Answer: C

QUESTION 369
Network topology exhibit:

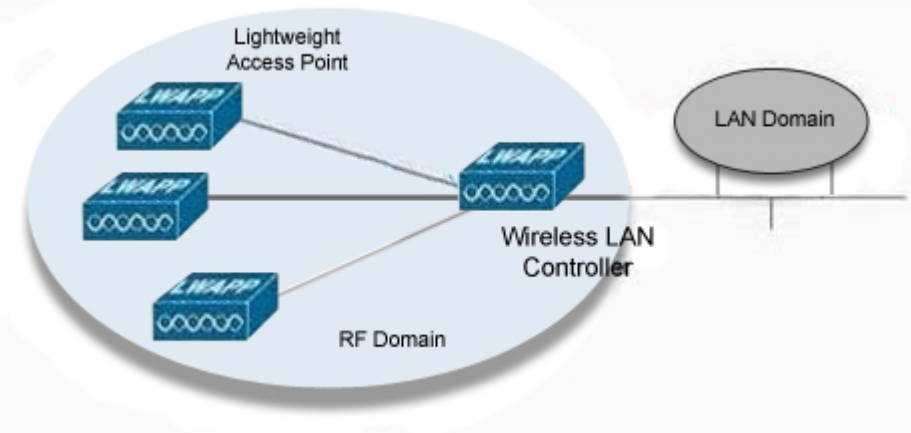

You work as a network technichian at Certkiller .com. Please study the exhibit carefully. Your boss, Mrs. Certkiller, is interested in the Lightweight Access Point technology.
What could you tell her? Select two.

A. Real time events such as authentication, security management, and mobility are handled by the lightweight AP.
B. WLAN controllers provide a single point of management.
C. An AP that has been upgraded from an autonomous AP to lightweight AP will only function in conjunction with a Cisco Wireless controller.
D. lLghtweight APs require local configuration using local management.

E. Autonomous APs receive control and configuration information from a WLAN controller.
F. LWAPP increases the amount of processing within the APs, enabling them to support filtering and policy enforcement features.

Answer: B,C

QUESTION 370
Study the exhibit below carefully:

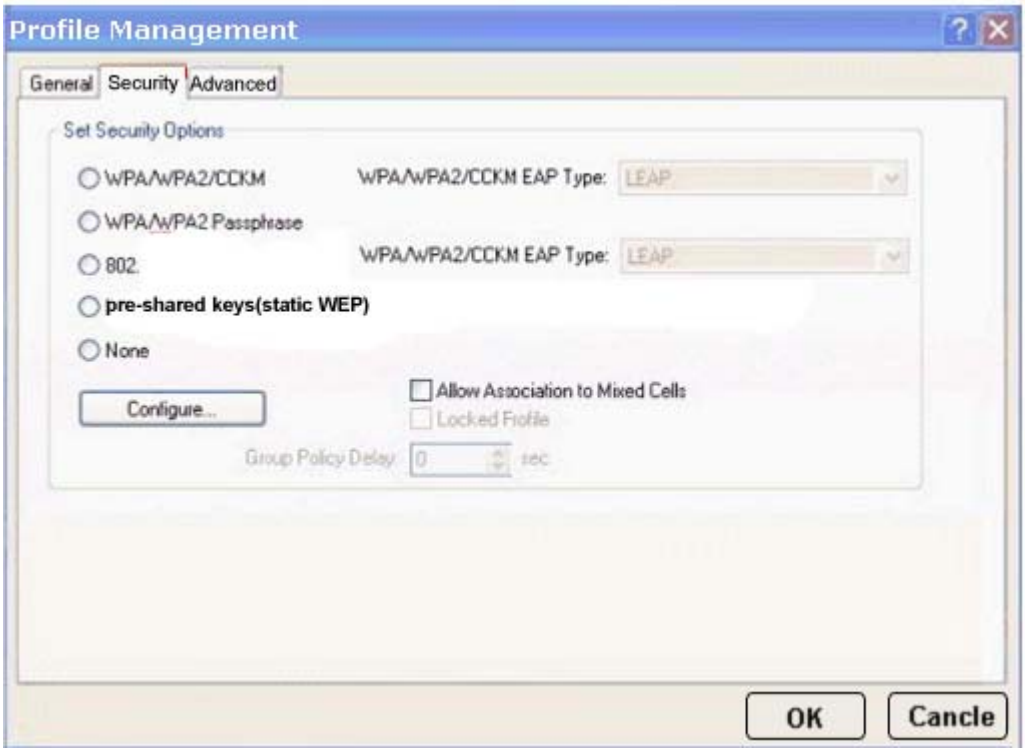

When a profile is configured for a user on the Certkiller network in the Aironet Desktop Utility, which security option permits the configuration of AES (Advaced Encryption Standard) add Extensible Authentication Protocol-Flexible Authentication via Secure Tunneling?

A. Pre-Shared Key (Static WEP)
B. WPA/WPA2/CCKM
C. 802.1x
D. WPA/WPA2 Passphrase
E. None of the above

Answer: B

QUESTION 371
Exhibit:

```
hostname certkiller3

<Output Omitted>
!
interface fastethernet0/5
  switchport mode access
  switchport port-security
  switchport port-security violation protect
  switchport port-security maximum 11
  switchport port-security aging time 0
```

You work as a network technichian at Certkiller .com. Please study the exhibit carefully. Port security has been configured on the switch port Fa0/5.
What would happen if another device is connected to the port after the maximum number of devices has been reached, even if one or more of the original MAC addresses are inactive?

A. Although one or more of the original MAC addresses are inactive, the port will not permit the new MAC address.
B. The port will permit the new MAC address because one or more of the original MAC addresses will age out.
C. Because the new MAC address is not configured on the port, the port will not permit the new MAC address.
D. The port will permit the new MAC address because one or more of the original MAC addresses are inactive.

Answer: A

QUESTION 372
You work as a network technichian at Certkiller .com. Your boss, Mrs. Certkiller, is interested in Spanning Tree Protocol default timers.
What can you tell her? Select three.

A. The hello time is 5 seconds.
B. The forward delay is 10 seconds.
C. The hello time is 2 seconds.
D. The forward delay is 15 seconds
E. The max_age timer is 15 seconds.
F. The hello time is 10 seconds.
G. The max_age timer is 20 seconds.
H. The forward delay is 20 seconds.
I. The max_age timer is 30 seconds.

Answer: C,D,G

QUESTION 373
Exhibit:

```
Certkiller2# show ip access-lists net_10
   Extend IP access list net_10
       10 permit ip 10.0.0.0 0.255.255.255 any

Certkiller2# conf t
Certkiller2( config)# vlan access-map thor 10
Certkiller2 (config-access-map) # match ip address net_10
Certkiller2 (config-access-map) # action forward
Certkiller2 (config-access-map) # exit
Certkiller2 (config)# vlan filter thor vlan-list 12-16
```

You work as a network technichian at Certkiller .com. Please study the exhibit carefully.
Consider traffic with source address of 172.16.10.5 within VLAN 14.
What would happen to this traffic?

A. It will be dropped.
B. I twill be forwarded to the TCAM for further processing.
C. I twill be forwarded to the router processor for further processing.
D. I twill be forwarded to without further processing.

Answer: A

QUESTION 374
Network topology exhibit:

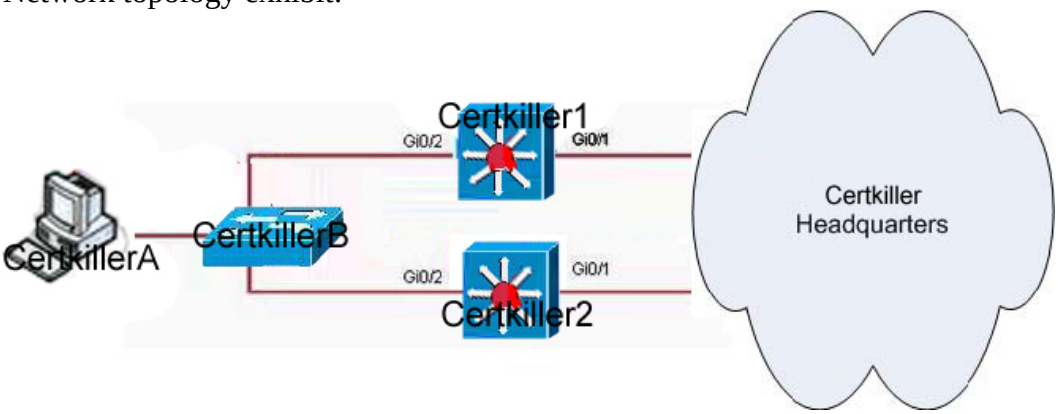

Certkiller 1 configuration exhibit:

```
Certkiller1# show running-config

<Output Omitted>

interface Vlan10
 ip address 10.10.10.2 255.255.255.0
 no ip redirects
 standby 62 priority 105
 standby 62 ip 10.10.10.1
 standby 62 track GigabitEthernet 0/1
```

Certkiller 2 configuration exhibit:

```
certkiller2#show running-config

<Output Omitted>

interface Vlan10
 ip address 10.10.10.3 255.255.255.0
 no ip redirects
 standby 62 priority 150
 standby 62 ip 10.10.10.1
 standby 62 track GigabitEthernet 0/1
```

You work as a network technichian at Certkiller .com. Please study the exhibit carefully.
In this scenario the following is true:
* Certkiller A can ping the Certkiller headquarter office
* HSRP is configured on Certkiller 1
* First Certkiller 1 and then Certkiller 2 are configured and reloaded
What can be said of this network?

A. Certkiller 1 will be the standby router because it has the lower IP address.
B. Certkiller 2 will be the standby router because it has the higher IP address.
C. Certkiller 1 will be the active router because it booted first.
D. Certkiller 2 will be the active router because it booted last.
E. Certkiller 1 will be the active router because it has the lower priority that is configured.
F. Certkiller 2 will be the active router because it has the higher priority that is configured.

Answer: C

QUESTION 375
You work as a network technichian at Certkiller .com. Your boss, Mrs. Certkiller, is interested in the QoS technology in the context of campus network.
What can be said of application of this technology in this type of network? Select three.

A. The access layer is the initial point at which traffic enters the network. Traffic is marked (or remarked) at Layers 2 and 3 by the access switch as it enters the network, or is "trusted" that it is entering the network with the appropriate tag.
B. No traffic marking occurs at the core layer. Layer 2/3 QoS tags are trusted from

distribution layer switches and used to prioritize and queue the traffic as it traverses the core.
C. Traffic inbound from the access layer to the distribution layer can be trusted or reset depending upon the ability of the access layer switches. Priority access into the core is provided based on Layer 3 QoS tags.
D. IP precedence, DSCP, QoS group, IP address, and ingress interface are Layer 2 characteristics that are set by the access layer as it passes traffic to the distribution layer. The distribution layer, once it has made a switching decision to the core layer, strips these off.
E. MAC address, Multiprotocol Label Switching (MPLS); the ATM cell loss priority (CLP) bit, the Frame Relay discard eligible (DE) bit, and ingress interface are established by the voice submodule (distribution layer) as traffic passes to the core layer.
F. The distribution layer inspects a frame to see if it has exceeded a predefined rate of traffic within a certain time frame, which is typically a fixed number internal to the switch. If a frame is determined to be in excess of the predefined rate limit, the CoS value can be marked up in a way that results in the packet being dropped.

Answer: A,B,C

QUESTION 376
Which statement is true about the LWAPP (Lightweight Access Point Protocol)?

A. Real-time frame exchange is accomplished within the access point.
B. The control traffic between the client and the access point is encapsulated with the LWAPP.
C. Authentication, security, and mobility are handled by the access point.
D. Data traffic between the client and the access point is encapsulated with LWAPP.

Answer: A

QUESTION 377
What statement is correct about RSTP port roles?

A. The root port is the switch port on very nonroot bridge that is the chosen path to the root bridge. There can only one root port on very switch. The root port assumes the forwarding state in a stalbe active topology.
B. The designated port is the switch port on every nonroot bridge that is the chosen path to the root bridge. There can be only one designated port on every switch. The designated port assumes the forwarding state in a stalbe active topology. All switches connected to a give segment listen to all BPDUs and determine the switch that will be root switch for a particular segment.
C. The backup port is a switch port that offers an alternate path toward the root bridge. The Backup port assumes a discarding state in a stable, active topology. The backup port will be present on nondesignated switches and will make a transition to a designated port of the current designated path fails.
D. The disabled port is an additional switch port on the designated switch the redundant

link to the segment for which the switch is designated. A disabled port has a higher port ID than the disabled port on the designated switch. The disabled port assumes the discarding state in a stable active topology.

Answer: A

QUESTION 378
You work as a network technichian at Certkiller .com. Your boss, Mrs. Certkiller, is interested in the Cisco Compatible Extensions program.
What at three features of this program?

A. accounting
B. analog and digital voice
C. mobility
D. security
E. routing and switching
F. VLAN and QoS

Answer: C,D,F

QUESTION 379
You work as a network technichian at Certkiller .com. Your boss, Mrs. Certkiller, is interested in the Layer 2 security attacks and mitigation techniques.
What should you tell her?

A. Enable root guard to mitigate ARP address spoofing attacks.
B. Configure DHCP spoofing to mitigate ARP address spoofing attacks.
C. Configure PVLANs to mitigate MAC address flooding attacks.
D. Enable root guard to mitigate DHCP spoofing attacks.
E. Configure dynamic APR inspection (DAI) to mitigate IP address spoofing on DHCP untrusted ports.
F. Configure port security to mitigate MAC address flooding

Answer: F

QUESTION 380
You work as a network technichian at Certkiller .com. Your boss, Mrs. Certkiller, is interested in LWAPP (Lightweight Access Point Protocol).
What should you tell her? Select 2.

A. LWAPP is a proprietary protocol, and because of its very high overhead it is not widely adopted.
B. Control traffic is encapsulated in UDP packets with a source port of 1024 and a destination port of 12223.
C. Layer 3 LWAPP is a UDP/IP Frame that requires a Cisco Aironet AP to obtain an IP address using DHCP.

D. Data traffic is encapsulated in UDP packets with a source port of 1024 and a destination port of 12223.
E. Control traffic is encapsulated in TCP packets with a source port of 1024 and a destination port of 12223.
F. Data traffic is encapsulated in TCP packets with a source port of 1024 and a destination port of 12223.

Answer: B,C

QUESTION 381
Network topology exhibit:

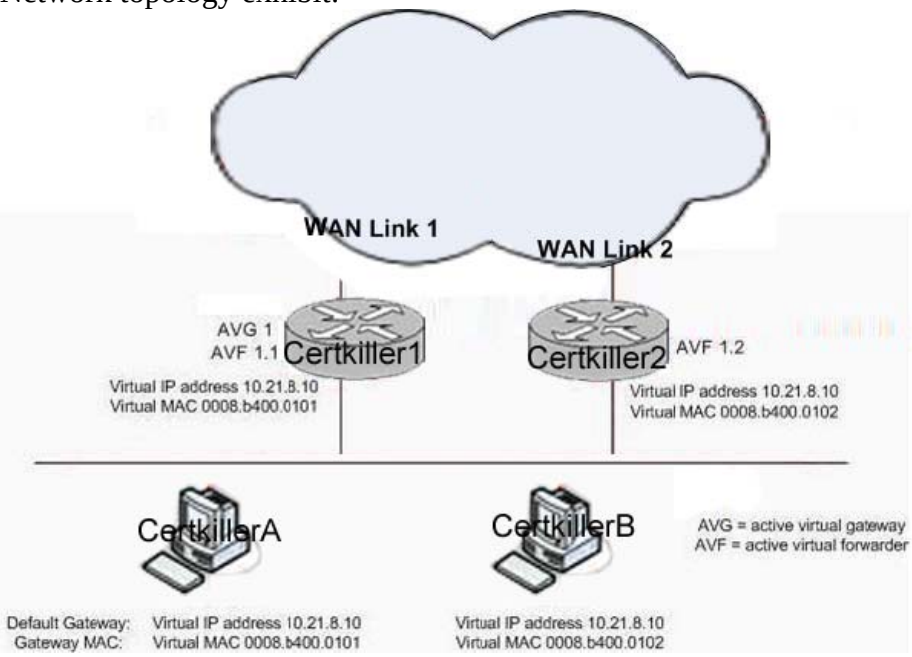

You work as a network technichian at Certkiller .com. Your boss, Mrs. Certkiller, is interested in GLBP.
Study the network toplogy exhibit carefully.
What should you tell her about this GLBP topology? Select 3.

A. If another router were added to this GLGP group, there would be two backup AVGs.
B. Router Certkiller 2 is in GLBP listen state.
C. Router Certkiller 2 will transition from blocking state to forwarding state when it becomes the AVG.
D. Router Certkiller 1 is responsible for answering ARP requests sent to the virtual IP address.
E. If router Certkiller 1 becomes unavailable, Router Certkiller 2 will forward packets sent to the virtual MAC address of Router Certkiller 1.
F. Router Certkiller 1 alternately responds to ARP requests with different virtual MAC addresses.

Answer: D,E,F

QUESTION 382
Network topology exhibit:

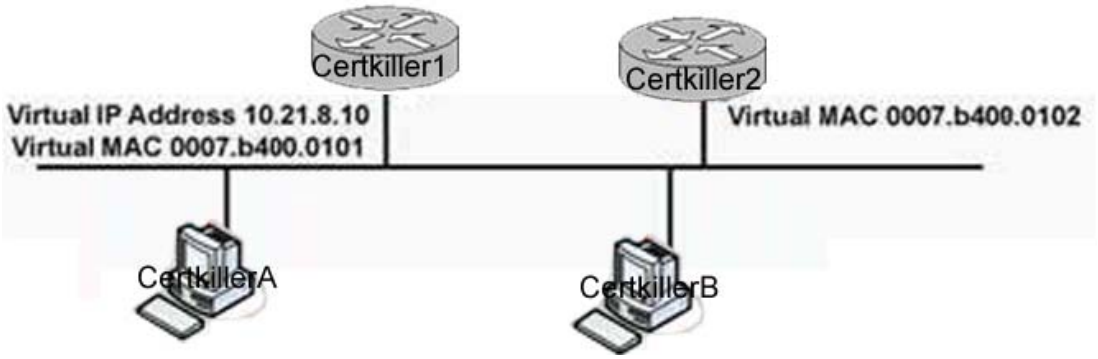

You work as a network technichian at Certkiller .com. Study the network toplogy exhibit carefully.
The two routers on the network are configured for the GLBP (Gateway Load Balancing Protocol).
What can be said?

A. The hosts will have different default gateway IP addresses and different MAC addresses for each rtouter.
B. The default gateway address of each host should be set to the virtual IP address.
C. The hosts will learn the proper default gateway IP address from Router Certkiller 1.
D. The default gateway address of each host should be set to the virtual IP address.

Answer: B

QUESTION 383
You work as a network technichian at Certkiller .com. Your boss, Mrs. Certkiller, is interested in Catalyst 2950 switches, in particular the voice VLAN feature.
What should you tell her? Select 3.

A. PortFast is is automatically disable when a voice VLAN is configured.
B. The CoS value is trusted for 802.1p or 802.1q tagged traffic.
C. The default CoS value for incoming traffic is set to 0.
D. The voice VLAN feature is disable by default.
E. When the voice VLAN features is enabled, all untagged traffic is sent according to the default CoS priority of the port.
F. The IP phone accepts the priority of all tagged and untagged traffic and sets the CoS vale to 4.

Answer: B,C,E

QUESTION 384
Network Topology exhibit:

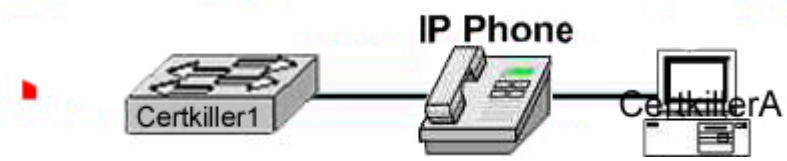

You work as a network technichian at Certkiller .com. Study the network toplogy exhibit carefully.
On the switch port that is connected to the IP phone you issue the command: mls qos trust cos
What is the trust boundary effect?

A. RTP will be used to negotiate a CoS value based upon bandwidth utilization on the link.
B. The host is now establishing the CoS value and has effectively become the trust boundary.
C. The switch will no longer tag incoming voice packets and will trust the distribution layer switch to set the CoS.
D. The switch is rewriting packet it receives from the IP phone and determining the CoS value.
E. Effectively the trust boundary has moved to the IP phone.

Answer: E

QUESTION 385
Exhibit:

You work as a network technichian at Certkiller .com. Study the exhibit carefully. You are required configure the static WEP keys on the wireless client adapter using the Cisco ADU (Aironet Desktop Utility). What should you have in mind?

A. Before the client adapter WEP key is generated, all wireless infrastructure devices (such as access points, servers, etc.) must be properly configured for LEAP authentication.

B. The client adapter WEP key should be generated by the AP and forwarded to the client adapter efore the client adapter can establish communication with the wireless network.
C. In infrastructure mode the client adapter WEP key must match the WEP key used by the access point. In ad hoc mode all client WEP keys within the wireless network must match each other.
D. The client adapter WEP key should be generated by the authentication server and forwarded to the client adapter before the client adapter can establish communication with the wireless network.

Answer: C

QUESTION 386
Network topology exhibit:

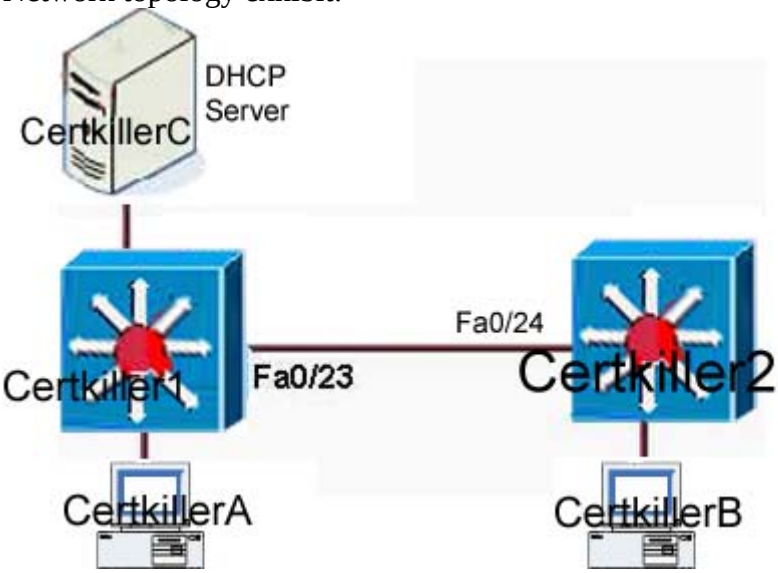

Certkiller 1 configuration exhibit:

```
hostname Certkiller1
!
 <output omitted>
!
ip arp inspection vlan 1
!
interface fastethernet 0/23
```

Certkiller 2 configuration exhibit:

```
hostname Certkiller2
!
<output omitted>
!
interface fastethernet 0/24
  switchport mode trunk
  switchport trunk encapsulation dot1q
```

You work as a network technichian at Certkiller .com. Study the exhibit carefully.
The following is true for this scenario:
* DAI (dynamic ARP inspection) is enabled on switch Certkiller 1, but not on Certkiller 2
* Host Certkiller A receives its IP address from the DHCP server connected to switch Certkiller 1.
* Host Certkiller B receives its IP address from the DHCP server connected to switch Certkiller 1.
* Host Certkiller B initiates an ARP spoof attack towards host Certkiller A.
What would happen with the spoof packets?

A. They will not be inspected at the ingress port of switch Certkiller 1 and will be permitted.
B. They will not be inspected at the ingress port of switch Certkiller 1 and will be dropped.
C. They will be inspected at the ingress port of switch Certkiller 1 and will be permitted.
D. They will be inspected at the ingress port of switch Certkiller 1 and will be dropped.

Answer: A

QUESTION 387
You work as a network technichian at Certkiller .com. Your boss, Mrs. Certkiller, is interested in switch spoofing. She asks you how the attacker how the attacker would collect information with VLAN hoping through switch spoofing.
What should you tell her?
That the attacking station...

A. ...uses VTP to collect VLAN information that is sent out and then tags itself with the domain information in order to capture the data.
B. ...will generate frames with two 802.1Q headers to cause the switch to forward the frames to a VLAN that would be inaccessible to the attacker through legitimate means.
C. ...uses DTP to negotiate trunking with a switch port and captures all traffic that is allowed on the trunk.
D. ...tags itself with all usable VLANs to capture data that is passed throuogh the switch, regardless of the VLAN to which the data belongs.

Answer: C

QUESTION 388
Network topology exhibit:

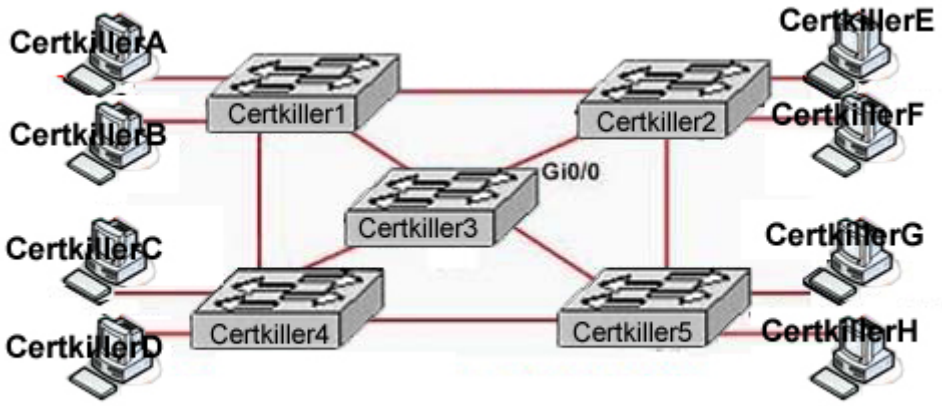

You work as a network technichian at Certkiller .com. Study the exhibit carefully.
The following applies in this scenario:
* Certkiller 3 has the command no spanning-tree portfast bpdufilter default enabled
* At interface G0/0 of Certkiller 3 the following command has been used:
spanning-tree bpdufilter enable
* A hardware failure makes the line between Certkiller 3 and Certkiller 5 go down.
* The spanning tree is recalcuted after the link failure
What effect does this happen between Certkiller 2 and Certkiller 3?
The link between switch Certkiller 2 and Certkiller 3 will...

A. ...participate in Spanning Tree, but will not pass traffic.
B. ...participate in Spanning Tree, and will pass traffic.
C. ...not participate in Spanning Tree, but will not pass traffic.
D. ...not participate in Spanning Tree, and will pass traffic.

Answer: B

QUESTION 389
Network topology exhibit:

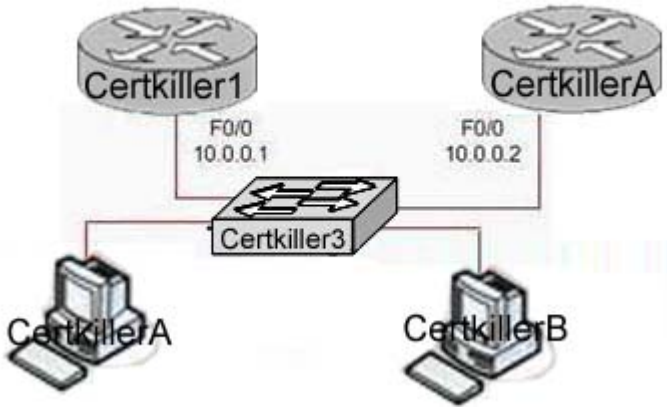

Certkiller 1 configuration exhibit:

```
Certkiller1 (config) # interface f0/0
Certkiller1 (config) #    ip address 10.0.0.2 255.255.255.00
Certkiller1 (config-if)# vrrp 1 priority 100
Certkiller1 (config-if)# vrrp   1 ip 10.0.0.10
```

Certkiller 2 configuration exhibit:

```
Certkiller1 (config) # interface f0/0
Certkiller1 (config) #    ip address 10.0.0.2 255.255.255.0
Certkiller1 (config-if)# vrrp 1 priority 100
Certkiller1 (config-if)# vrrp   1 ip 10.0.0.10
```

You work as a network technichian at Certkiller .com. Study the exhibit carefully. Your boss, Ms. Certkiller, is interested VRRP (Virtual Router Redundancy Protocol). In particular she is curious about the roles of the backup virtual router and the master virtual router in the current scenario.
What can you tell her?

A. Router Certkiller 2 is the master virtual router, and Router Certkiller 1 is the backup virtual router. When Router Certkiller 2 fails, Router Certkiller 1 will become the master virtual router. When Router Certkiller 2 recovers, it will regain the master virtual router role.
B. Router Certkiller 2 is the master virtual router, and Router Certkiller 1 is the backup virtual router. When Router Certkiller 2 fails, Router Certkiller 1 will become the master virtual router. When Router Certkiller 2 recovers, Router Certkiller 1 will maintain the role of master virtual router.
C. Router Certkiller 1 is the master virtual router, and Router Certkiller 2 is the backup virtual router. When Router Certkiller 1 fails, Router Certkiller 2 will become the master virtual router. When Router Certkiller 1 recovers, it will regain the master virtual router role.
D. Router Certkiller 1 is the master virtual router, and Router Certkiller 2 is the backup virtual router. When Router Certkiller 1 fails, Router Certkiller 2 will become the master virtual router. When Router Certkiller 1 recovers, Router Certkiller 2 will maintain the role of master virtual router.

Answer: C

QUESTION 390
You work as a network technichian at Certkiller .com. Your boss, Ms. Certkiller, is interested LWAPP (Lightweight Access Point Protocol). In particular she wants to know which type of activities this protocol defines. What should you tell her? Select two.

A. Layer 3 addressing and distribution
B. SNMP monitoring services
C. Access point certification and software control
D. Packet encapsulation, fragmentation, and formatting
E. Compression and Layer 3 address mapping

Answer: C,D

QUESTION 391
You work as a network technichian at Certkiller .com. Your boss, Ms. Certkiller, is interested in the WLAN client utility. What should you tell her? Select two.

A. The Aironet Desktop Utility (ADU) can be used to enable or disable the adapter radio and to configure LEAP authentication with dynamic WEP.
B. In Windows XP environment, a client adapter can only be configured and managed with the Microsoft Configuration Manager.
C. The Microsoft Wireless Configuration Manager can be configured to display the Aironet System Tray Utility (ASTU) icon in the Windows system tray.
D. The Cisco Aironet Desktop Utility (ADU) and the Microsoft Wirelss Configuration Manager can both be enabled at the same time to setup WLAN client cards.

Answer: A,C

QUESTION 392
You work as a network technichian at Certkiller .com. Your boss, Ms. Certkiller, is interested in the Aironet enterprise solutions. What should you tell her? Select two.

A. The Cisco Aironet AP handles real-time portions of the LWAPP protocol, and the WLAN controller handles those items which are not time sensitive.
B. A Cisco Aironet AP handles the transmission of beacon frames and also handles responses to probe-request frames clients.
C. Virtual MAC architecture allows the splitting of the 802.11 protocol between the Cisco Aironet AP and a LAN switch.
D. A Cisco Aironet solution includes intelligent Cisco Aironet access point (APs) and Cisco Catalyst switches.
E. In the Cisco Aironet solution, each AP is locally configured by the use of either a web interface or the command line interface.

Answer: A,B

QUESTION 393
DRAG DROP
You work as a network technichian at Certkiller .com. Your boss, Ms. Certkiller, is interested in the process of a wireless client associating with a wireless access point. In particular, what is the correct order this takes places?

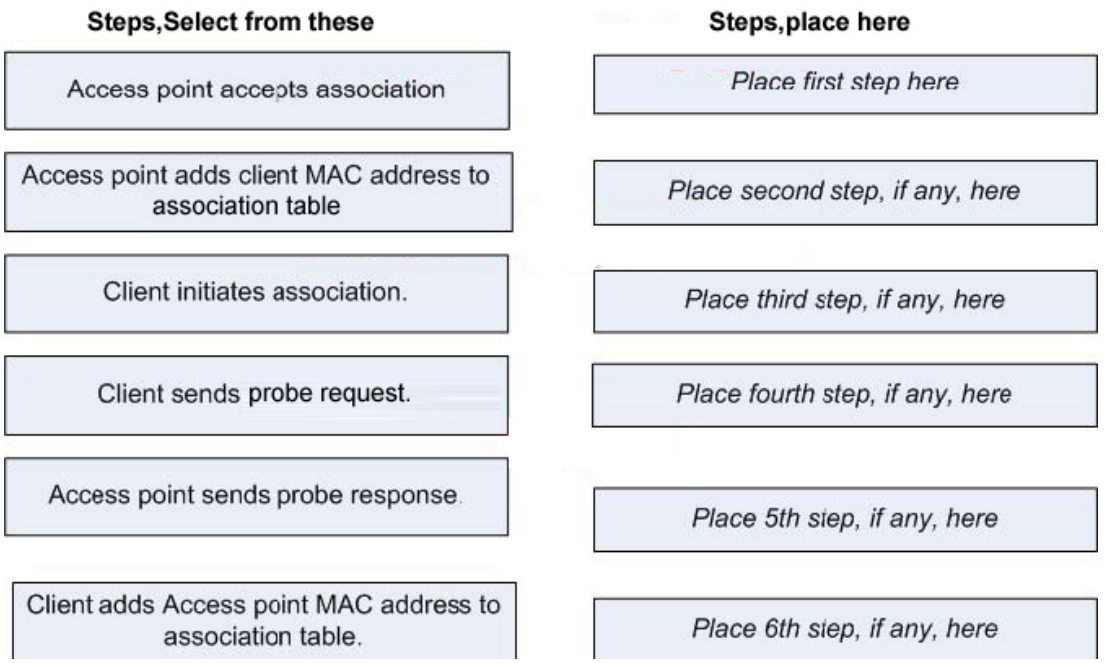
Steps,Select from these
Steps,place here
Access point accepts association
Access point adds client MAC address to association table
Client initiates association.
Client sends probe request.
Access point sends probe response.
Client adds Access point MAC address to association table.
Place first step here
Place second step, if any, here
Place third step, if any, here
Place fourth step, if any, here
Place 5th step, if any, here
Place 6th step, if any, here

Answer:

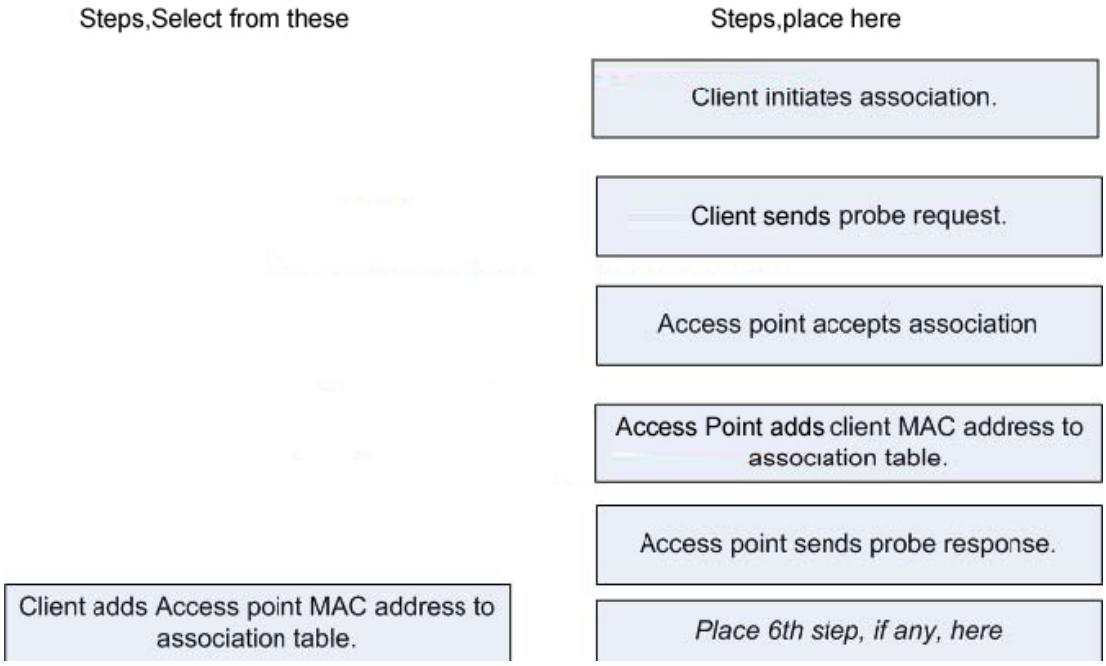
Steps,Select from these
Steps,place here
Client initiates association.
Client sends probe request.
Access point accepts association
Access Point adds client MAC address to association table.
Access point sends probe response.
Client adds Access point MAC address to association table.
Place 6th step, if any, here