

QUESTION 1

Layer 6 of the 7-Layer OSI model is responsible for:

- A. Common Data Compression and Encryption Schemes
- B. Establishing, managing, and terminating communication sessions
- C. Synchronizing communication
- D. Determining resource availability
- E. None of the above

Answer: A

Explanation:

Layer 6 is the Presentation Layer. This layer provides independence from differences in data representation (e.g., encryption and compression) by translating from application to network format, and vice versa. The presentation layer works to transform data into the form that the application layer can accept. This layer formats and encrypts data to be sent across a network, providing freedom from compatibility problems. It is sometimes called the syntax layer.

Incorrect Answers:

- B: This describes layer 5 of the OSI model, which is the Session Layer.
- C, D: These are not responsibilities of the Presentation Layer.

QUESTION 2

Which of the following is a component of the Data Link Layer of the OSI model?

- A. NIC
- B. Repeater
- C. Multiplexer
- D. Hub
- E. Router

Answer: A

Explanation:

The data link layer is layer 2 in the OSI model, and deals with things like MAC addresses, and link level technologies such as Ethernet and Token Ring. Network interface cards (NICs) typically implement a specific data link layer technology, so they are often called "Ethernet cards", "Token Ring cards", and so on. They also include a 48 bit MAC address, also called a burned in address since these addresses are burned into the cards.

Incorrect Answers:

- B, C, D: Repeaters, Hubs, and Multiplexers deal with the physical connections of devices into a network, and they are considered to reside on the physical layer of the OSI

model (layer 1).

E: Routers operate at layer 3 and 4 of the OSI model, since they deal with things like layer 3 IP addresses, and TCP/UDP ports.

QUESTION 3

Which statement is true regarding the use of TFTP?

- A. TFTP lies at the Transport layer and runs over IP.
- B. TFTP lies at the Application layer and runs over FTP.
- C. TFTP lies at the Transport layer and runs over ICMP.
- D. TFTP lies at the Application layer and runs over TCP.
- E. TFTP lies at the Application layer and runs over UDP.

Answer: E

Explanation:

Trivial File Transfer Protocol (TFTP) is a simplified version of FTP that allows files to be transferred from one computer to another over a network, usually without the use of client authentication (for example, username and password). It is an application that uses UDP port 69.

QUESTION 4

In a data communication session between two hosts, the session layer in the OSI model generally communicates with what other layer of the OSI model?

- A. The Physical layer of the peer
- B. The data link layer of the peer
- C. The peer's presentation layer
- D. The peer's application layer
- E. The peer's session layer

Answer: E

Explanation:


Interactions Between the Same Layers on Different Computers

Layer N must interact with Layer N on another computer to successfully implement its functions. For example, the transport layer (Layer 4) can send data, but if another computer does not acknowledge that the data was received, the sender will not know when to perform error recovery. Likewise, the sending computer encodes a destination network layer address (Layer 3) in the network layer header. If the intervening routers do not cooperate by performing their network layer tasks, the packet will not be delivered to the true destination.

To interact with the same layer on another computer, each layer defines a header and, in some cases, a trailer. Headers and trailers are additional data bits, created by the sending computer's software or hardware, that are placed before or after the data given to Layer N by Layer $N+1$. The information needed for this layer to communicate with the same layer process on the other computer is encoded in the header and trailer. The receiving computer's Layer N software or hardware interprets the headers and trailers created by the sending computer's Layer N , learning how Layer N 's processing is being handled, in this case.

Figure 3-3 provides a conceptual perspective on the same-layer interactions. The application layer on Host A communicates with the application layer on Host B. Likewise, the transport, session, and presentation layers on Host A and Host B also communicate. The bottom three layers of the OSI model have to do with delivery of the data; Router 1 is involved in that process. Host A's network, physical, and data link layers communicate with Router 1; likewise, Router 1 communicates with Host B's physical, data link, and network layers. Figure 3-3 provides a visual representation of the same-layer interaction concepts.

Figure 3-3 Same-Layer Interactions on Different Computers



QUESTION 5

Which layers do the OSI model and the TCP/IP models share in common? (Choose all that apply)

- A. Application
- B. Presentation
- C. Session
- D. Transport
- E. Data link
- F. Physical

Answer: A, D

Explanation:

The TCP/IP reference model has the following layers:
Application, Transport, Internet, and Host to Network.

Incorrect Answers:

B, C, E, F. The TCP/IP reference model does not have a presentation layer, a session layer, a physical layer, or a data-link layer.

QUESTION 6

Under the OSPF process of your router's configuration, you type in "redistribute igrp 25 metric 35 subnets" in order to redistribute your OSPF and IGRP routing information. What affect did the "subnets" keyword have in your configuration change?

- A. It resulted in OSPF recognizing non-classful networks.
- B. It had no effect since IGRP will summarize class boundaries by default.
- C. It forced IGRP into supporting VLSM information.
- D. It caused OSPF to accept networks with non-classful masks.

Answer: D

Explanation:

Whenever there is a major net that is subnetted, you need to use the keyword subnet to redistribute protocols into OSPF. Without this keyword, OSPF only redistributes major network boundaries. It is possible to run more than one OSPF process on the same router, but running more than one process of the same protocol is rarely needed, and it consumes the router's memory and CPU.

Incorrect Answers:

- A. OSPF already always recognizes non-classful networks and their VLSM information.
 - B. Although IGRP does indeed summarize by class boundaries, OSPF does not by default. The "subnets" keyword enables OSPF to use VLSM information from the IGRP routes.
 - C. IGRP does not support VLSM routing information.
-

QUESTION 7

Which routing protocols do not need to have their router ID reachable by other routers within any given network in order to maintain proper network connectivity? (Choose all that apply)

- A. EIGRP
- B. OSPF
- C. BGP
- D. LDP
- E. TDP
- F. None of the above

Answer: A, B, C

Explanation:

The router ID of each router does not necessarily need to be reached by other routers in

the network for EIGRP and OSPF. BGP uses TCP as the reliable exchange of information between routers, and BGP routers do not need to even be directly connected.

Incorrect Answers:

D, E. LDP and TDP are not routing protocols.

QUESTION 8

Which of the following does On Demand Routing use to transport ODR information from router to router?

- A. RIP
- B. BGP
- C. CDP
- D. UDP
- E. LSP

Answer: C

Explanation:

ODR uses information from the Cisco Discovery Protocol (CDP).

Incorrect Answers:

A, B, D, E. ODR has nothing to do with RIP, BGP, UDP, or LSP.

QUESTION 9

A router running multiple protocols learns how to reach a destination through numerous different methods. Which of the following information will the router use first to determine the best way to reach the given destination?

- A. The length of the network mask of a route.
- B. The administrative distance of a route.
- C. The metric of a route.
- D. None of the above.

Answer: A

Explanation:

Most specific network match is always used first.

Incorrect Answers:

B, C: The administrative distance and metric is consulted only for routes with the same network mask length.

QUESTION 10

Which of the following routing protocols has a default administrative distance less than the default IS-IS AD?

- A. External EIGRP routes
- B. iBGP routes

- C. Internal EIGRP routes
- D. RIP version 1 routes
- E. eBGP

Answer: C, E

Explanation:

The default IS-IS administrative distance is 115. Internal EIGRP routes are 90, and external BGP is 20.

Incorrect Answers:

- A. External EIGRP routes have an AD of 170.
- B. Interior BGP routes have an AD of 200.
- D. RIP routes have an AD of 120.

QUESTION 11

Which of the following are key differences between RIP version 1 and RIP version 2? (Choose all that apply)

- A. RIP version 1 supports authentication while RIP version 2 does not.
- B. RIP version 2 uses multicasts while RIP version 1 does not.
- C. RIP version 1 uses hop counts as the metric while RIP version 2 uses bandwidth information.
- D. RIP version 1 does not support VLSM while RIP version 2 does.
- E. RIP version 1 is distance vector while RIP version 2 is not.

Answer: B, D

Explanation:

Both Classless Routing and Multicast updates (224.0.0.9) were impossible with RIP v1 and are available with RIP version 2.

Incorrect Answers:

- A. RIPv2 supports neighbor authentication. RIPv1 does not support this.
- C. Both RIP versions use hop counts as the metric.
- E. Both RIP versions are distance vector routing protocols.

QUESTION 12

You are deciding which routing protocol to implement on your network. When weighing the different options, which of the following are valid considerations?

- A. Distance vector protocols have a finite limit of hop counts whereas link state protocols place no limit on the number of hops.
- B. Distance vector protocols converge faster than link state protocols.
- C. RIP is a distance vector protocol. RIP v2 and OSPF are link state protocols.
- D. Distance vector protocols only send updates to neighboring routers. Link state

protocols depend on flooding to update all routers in the within the same routing domain.

Answer: A

Explanation:

Only A is true.

Incorrect Answers:


B. Link state protocols have the benefit of better convergence than distance vector protocols.

C. RIPv2 is a distance vector protocol, just like RIP version 1.

D. Link state protocols do not flood updates to every router within the same domain, just within their area.

QUESTION 13

The Certkiller network contains two routers named Router CK1 and Router CK2 as shown in the following exhibit:



Both Router CK1 and Router CK2 are running RIPv1. Both routers are configured to advertise all of their attached networks via RIP. Which of the networks connected to Router CK2 will be advertised to Router CK1 ?

- A. 10.10.10.0/27 and 134.10.15.0/28
- B. 10.0.0.0/8 and 192.75.0.0/24
- C. 134.10.15.0/28 and 192.75.0.0/22
- D. Only 10.0.0.0/8
- E. Only 134.10.15.0/28
- F. Only 10.10.10.0/27
- G. None of the above

Answer: D

Explanation:

Only one subnet 10.0.0.0/8 will be advertised.

In this scenario we are being tested on the following concepts:

RIP V1 performs auto summarization at network boundaries by default. It treats the subnets to be advertised differently depending upon several attributes of the respective subnets.

Here is the process RIP v1 uses to advertise, assuming that there are no filters (such as

distribute-lists, or route-maps) to block the packet:

Is the route to be advertised part of the major network of the interface?

If it is, then advertise. If it is not, then summarize the network to its classful boundary and send it out.

This is the fate of the 10.10.10.0/27 subnet, which will be summarized as 10.0.0.0/8 and sent out.

Incorrect Answers:

A, C, E. If the route is part of the major network, check to see if the subnet mask matches that of the outgoing interface. If the subnet mask does match then advertise the route out the interface. If the subnet mask of the route does not match the interface's subnet mask, then do not advertise the route out the interface unless the route is a host route (/32). This is the fate of the 134.10.15.0/28 subnet, which will not be sent out (advertised) at all.

B, C. Super net advertisement (advertising any network prefix less than its classful major network) is not allowed in RIP route summarization. This is the fate of the 192.75.1.2/22 subnet, which will not be sent out (advertised) at all.

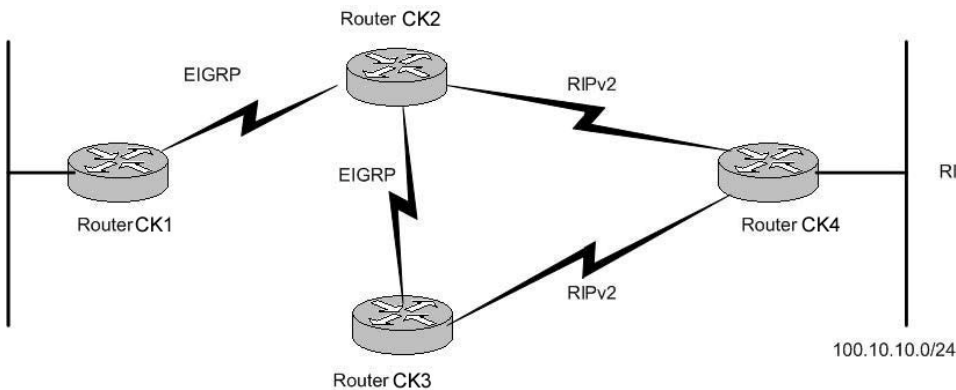
F. The 10.10.10.0/27 network will be summarized and sent as 10.0.0.0/8.

Please note:

If the route is a host route then advertise it out.

QUESTION 14

You are the network administrator at Certkiller . The Routing protocols which run between the different routers in the Certkiller network are shown in the following exhibit:



On Router CK3 RIPv2 is being redistributed into EIGRP. No other redistribution is done to the network.

With regard to this scenario, who owns the route for subnet 100.10.1.0/24 in the routing table of Router CK1 ?

- A. Nobody, because the route is neither in the routing table of Router CK1 , nor EIGRP topology table.
- B. External EIGRP.
- C. The route is only in the EIGRP topology table only and not in the routing table of Router CK1 .
- D. Internal EIGRP.
- E. The route is only but is in the EIGRP topology table as an active route and not in

the routing table of Router CK1 .

Answer: B

Explanation:

External EIGRP will own the route, because the route is from outside the AS. Routes that are redistributed into EIGRP are automatically considered external EIGRP routes.


Incorrect Answers:

- A. Since RIPv2 allows for VLSM information to be carried in the route, there are no concerns about the route not being advertised due to summarization. Since RIPv2 is being redistributed into EIGRP, CK1 will learn about the route via CK2 and CK3 .
- C, E. This route will be in both the EIGRP table, as well as the IP routing table.
- D. Redistributed routes always show up as External routes.

Note: From the perspective of router CK1 , all routes are EIGRP learned, since that is the only protocol running on this router. Although the AD of RIP is lower than external EIGRP routes, RIP is not being configured on CK1 so it will not learn this route via RIP.

QUESTION 15

The router topology for the multi-protocol Certkiller network is shown in the following exhibit:



The current configuration for Router CK1 , Router CK2 , Router CK3 , and Router CK4 are as follows:

Router CK1 :

```
interface loopback0
ip address 1.1.1.1 255.255.255.255
router eigrp 10
network 1.0.0.0
network 10.0.0.0
interface loopback1
ip address 4.4.4.4 255.255.255.255
```

Router CK2

```
router eigrp 10
network 10.0.0.0
network 20.0.0.0
no auto-summary
```

Router CK3

```
router ospf 10
network 30.30.30.0 0.0.0.255 area 0
network 40.40.40.0 0.0.0.255 area 0
```

Router CK4

```
router eigrp 10
redistribute connected metric 1400 230 1 255 1500
network 20.0.0.0
```

no auto-summary

```
router ospf 10
```

```
redistribute connected metric 100 subnets
```

```
network 40.40.40.0 0.0.0.255 area 0
```

```
router bgp 10
```

```
network 100.10.1.0 mask 255.255.255.0
```

```
neighbor 1.1.1.1 remote-as 10
```

```
neighbor update-source loopback
```

```
no auto-summary
```

Your newly appointed Certkiller trainee wants to know who owns the subnet 100.10.1.0/24 in the routing table of Router CK1 .

What would your reply be?

- A. Router CK1 does not have this subnet in its routing table.
- B. EIGRP
- C. OSPF
- D. BGP
- E. RIP
- F. It is there as a static route.

Answer: B

Explanation:

Routers CK1 , CK2 , and CK4 are all EIGRP neighbors with all relevant subnets advertised, so this route will show up as an EIGRP route.

Incorrect Answers:

C, D, E. Router CK1 is only running the EIGRP protocol, so the other routing protocols are completely ruled out.

QUESTION 16

Which of the following are Distance Vector routing protocols? (Choose all that apply)

- A. OSPF
- B. BGP
- C. RIP version 1
- D. ISIS
- E. EIGRP

F. RIP version 2

Answer: C, F

Explanation:

Both RIP version 1 and RIP version 2 are distance vector protocols.

Incorrect Answers:

A, D. OSPF and ISIS are link state routing protocols.

B. BGP is a path vector protocol, which is similar to a distance vector protocol, but with a key difference. A distance vector protocol chooses routes based on hop count, where BGP chooses routes that traverse the least number of Autonomous Systems, among other things.

E. EIGRP is an advanced Cisco-proprietary hybrid routing protocol, which means it uses a metric more sophisticated than distance (hop count) for route selection.

QUESTION 17

As the administrator of the Certkiller network, you are planning to implement a dynamic routing protocol to replace the static routes. When comparing link state and distance vector routing protocols, what set of characteristics best describe Link-State routing protocols?

- A. Fast convergence and lower CPU utilization
- B. High CPU utilization and prone to routing loops
- C. Slower convergence time and average CPU utilization
- D. Fast convergence and greater CPU utilization
- E. None of the above

Answer: D

Explanation:

Link State protocols, such as IS-IS and OSPF, converge more quickly than their distance vector counterparts, through the use of flooding and triggered updates. In link state protocols, changes are flooded immediately and computed in parallel.

Triggered updates improve convergence time by requiring routers to send an update message immediately upon learning of a route change. These updates are triggered by some event, such as a new link becoming available or an existing link failing.

The main drawbacks to Link State protocols are the amount of CPU overhead involved in calculating route changes and memory resources that are required to store neighbor tables, route tables, and a complete topology map.

QUESTION 18

A customer has a router with an interface connected to an OSPF network, and an interface connected to an EIGRP network. Both OSPF and EIGRP have been configured on the router. However, routers in the OSPF network do not have route entries in the route table for all of the routers from the EIGRP network. The default-metric under OSPF is currently set to 16. Based on this information, what is

the most likely cause of this problem?

A. The 'subnets' keyword was not used under the OSPF process when redistributing EIGRP into OSPF.

B. EIGRP is configured as a Stub area, and therefore routes will not be redistributed unless a route-map is used to individually select the routes for redistribution.

C. The 'subnets' keyword was not used the EIGRP process when redistributing between OSPF into EIGRP.

D. The default metric for OSPF is set to 16, and therefore all EIGRP routes that are redistributed are assigned this metric, and are automatically considered unreachable by EIGRP.

E. A metric was not assigned as part of the redistribution command for EIGRP routes redistributing into OSPF, and the default behavior is to assign a metric of 255, which is considered unreachable by OSPF.

Answer: A

Explanation:

When routes are redistributed into OSPF, only routes that are not subnetted are redistributed if the subnets keyword is not specified. It is generally a good idea to include the "subnets" keyword at all times when redistributing routes from other protocols into OSPF.

Incorrect Answers:

B. There is nothing in this question to lead us to believe that stub networks are being used at all. Even if they were, route maps would not be needed to redistribute the EIGRP and OSPF routes.

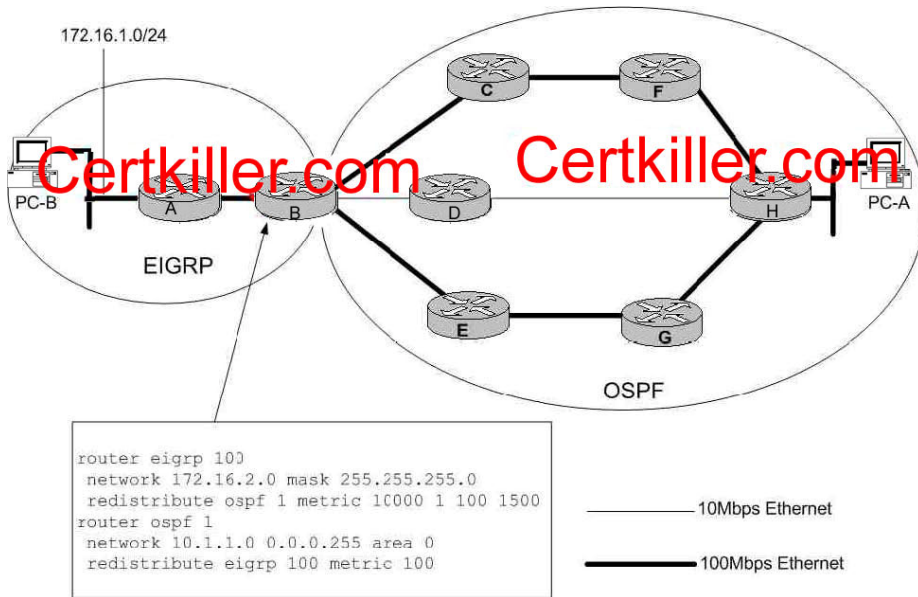
C. The "subnets" keyword needs to be placed under the OSPF process, not the EIGRP process.

D. EIGRP routes with a metric of 16 are acceptable, and not considered unreachable. If the routing protocol used was RIP instead of EIGRP then this would be true.

E. When redistributing from one OSPF process to another OSPF process on the same router, the metric will be carried through from one process to the other if no metric value is specified. When redistributing other processes to an OSPF process, the default metric is 20 when no metric value is specified.

QUESTION 19

The Certkiller WAN consists of an OSPF network portion and an EIGRP routed portion as shown in the display below:



Given the network and OSPF configuration shown in the exhibit, what statement is true regarding traffic flowing from PC-A to PC-B?

- A. Traffic will only flow on the shortest, low-speed path, PC-A-H-D-B-A-PC-B.
- B. Traffic will flow on both of the high speed paths (PC-A-H-F-C-B-A-PC-B and PC-A-H-G-E-B-A-PC-B) but not the slow-speed path.
- C. Traffic will flow on all three of the paths.
- D. Traffic will flow uni-directionally on one of the high-speed paths from PC-A to PC-B, and uni-directionally on the other high-speed path from PC-B to PC-A.
- E. Traffic will flow bi-directionally on only one of the high-speed paths, and the path selected will be based on the OSPF process IDs.


Answer: B

Explanation:

The default metric for OSPF is 100,000,000 divided by the bandwidth. For each 100 Mbps fast Ethernet link, the OSPF cost will be 1. For the slower, 10 Mbps Ethernet link, the OSPF cost will be 10, so the traffic will be routed around the slower link to the high speed links even though more hops are involved, because each high speed link across the entire OSPF cloud will have a total cost of 3 (1+1+1). By default, OSPF will load balance traffic across up to four equal cost paths. Therefore, choice B is correct in that traffic will utilize both high speed links.

QUESTION 20

The Certkiller network is redistributing IS-IS and RIP version 2 routes as shown in the diagram below:



Routers B and G both advertise RIP learned routes into IS-IS. Network is added to Router A via an Ethernet port and Router B is the First router to learn about this new network. After the network has converged, what path will Router G take to reach network 200.155.16.16?

- A. Router G takes the direct path through router A.
- B. Router G takes the path through routers, F, E, D, C, B, A.
- C. Router G will oscillate between the path through router A and the path through router F.
- D. Router G and router B will both think the other router is the best path to network 200.155.16.16, causing a routing loop.
- E. The answer can not be determined unless the default-metric used in the redistribution is known.

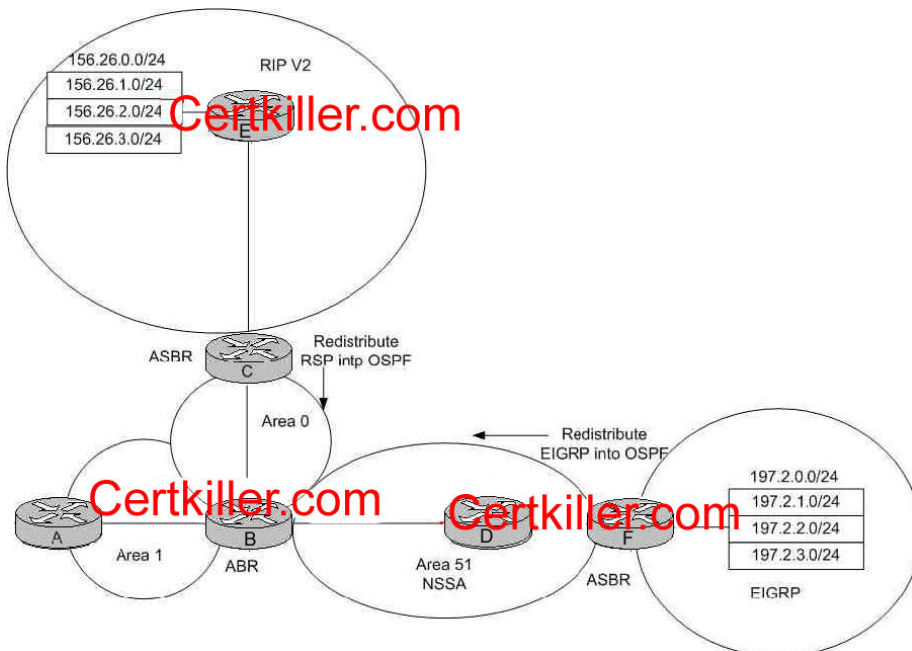
Answer: B

Explanation:

When a router receives identical route and subnet mask information for a given network from two different routing protocols, the route with the lowest administrative distance is chosen. IS-IS has a lower administrative Distance than RIP, so this route is installed in the routing table and used, even though it is obviously not the optimal route in this specific example.

QUESTION 21

The Certkiller network uses multiple IP routing protocols with redistribution, as shown in the diagram below:



Area 51 is configured as a NSSA Totally Stub, using the "area 51 stub no-summary" command. Which routers are in the routing table of Router D?

- A. Redistributed EIGRP and RIP routes, one OSPF default route, OSPF inter and intra-area routes
- B. Redistributed EIGRP routes and OSPF intra-area routes
- C. Redistributed EIGRP routes and OSPF inter and intra-area routes
- D. Redistributed EIGRP routes, an OSPF default route and OSPF intra-area routes
- E. Redistributed EIGRP and RIP routes and an OSPF default route

Answer: D

Explanation:

In the network diagram above, Area 51 is defined as a totally NSSA stub area. EIGRP routes cannot be propagated into the OSPF domain because redistribution is not allowed in the stub area. However, if we define area 51 as NSSA, we can inject EIGRP routes into the OSPF NSSA domain by creating type 7 LSAs. Redistributed RIP routes will not be allowed in area 51 because NSSA is an extension to the stub area. The stub area characteristics still exist, including no type 5 LSAs allowed.

There are two ways to have a default route in an NSSA

A. When you configure an area as NSSA, by default the NSSA ABR does not generate a default summary route. In the case of a stub area or an NSSA totally stub area, the NSSA ABR does generate a default summary route. In addition, all OSPF intra-area routes are allowed in a totally NSSA area.

Incorrect Answers:

A, E. The RIP will become external OSPF routes after the redistribution takes place. Since External OSPF routes from a different area are not injected into NSSA areas, no RIP routes will be seen on router D.

- B. By making the not-so-stubby area a totally not-so-stubby area, a default route is injected, so D is the preferred choice over B.
- C. Inter-area routes are not seen on routers within a totally NSSA.

QUESTION 22

The Certkiller network is displayed below:



Based on the information above, what path would router Certkiller 8 use to reach a network on router Certkiller 1? (Assume that mutual route redistribution takes place at all protocol boundaries)

- A. Router Certkiller 8 takes the path through Certkiller 6.
- B. Router Certkiller 8 takes the path through Certkiller 4.
- C. Router Certkiller 8 takes the path through Certkiller 3.
- D. Router Certkiller 8 takes the path through Certkiller 2.
- E. None of the above.

Answer: A

Explanation:

Assuming that redistribution of routes is taking place on all routers, Certkiller 8 would receive multiple routes to the same network destination. Because of this, Certkiller 8 would choose to install the route with the lowest Administrative Distance into the routing table. The default AD of the routes shown above is:

EIGRP: 90

IGRP: 100

OSPF: 110

RIP: 120

Therefore, router Certkiller 8 will go through the IGRP route via router Certkiller 6.

QUESTION 23

You need to make a new cable that is to be used for connecting a switch directly to another switch using Ethernet ports. What pinouts should be used for this cable?

- A. 1->3, 2->6, 3->1, 6->2
- B. 1->1, 2->2, 3->3, 6->6,

- C. 1->4, 2->5, 4->1, 5->2
- D. 1->5, 2->4, 4->2, 5->1
- E. 1->6, 2->3, 3->2, 6->1

Answer: A


Explanation:

Straight through cables are used when connecting PC hosts and router Ethernet ports to switches. Crossover cables are needed for switch to switch, and router to router connections. More information on crossover cables and their pinouts follows:

Crossover

In the implementation of a crossover, the wires on each end of the cable are crossed. Transmit to Receive and Receive to Transmit on each side, for both tip and ring. Figure 1.19 shows the UTP crossover implementation.

UTP crossover implementation



QUESTION 24

The ITU-T Q.920 and ITU-T Q.921 drafts formally specify which protocol?

- A. HDLC
- B. PPP
- C. LAPD
- D. HSRP
- E. LLC

Answer: C

Explanation:

The LAPD protocol is formally specified in ITU-T Q.920 and ITU-T Q.921.

Incorrect Answers:

A, D. HDLC and HSRP are both Cisco proprietary and are not formally specified in any ITU-T drafts.

QUESTION 25

What is the maximum transmit value for LLC flow control, as defined formally in

the IEEE 802.2 LLC standard?

- A. 15
- B. 127
- C. 256
- D. 1023
- E. 4096

Answer: B

Explanation:

According to the IEEE 802.2 Logical Link Control specification, the maximum transmit value for LLC flow control is 127. The LLC flow control techniques for bridged LANs is described as follows:

Overview:

This annex describes a technique, called dynamic window flow control, to control the offering of frames to the network by an LLC entity when congestion is detected or suspected. It is most effective in a bridged LAN. The technique is one of recovery from congestion and does not prevent congestion in a bridged LAN. It is not a substitute for proper network sizing. The method employs the transmit window already permitted by the standard to regulate the flow between two LLCs using the connection-mode service. Congestion in one direction of a logical link connection is treated independently of congestion in the other direction. The technique does not involve communication with the bridges, but rather relies on a simple algorithm implemented by the LLCs. MAC protocols are unaffected. All actions described in this annex apply to the station transmitting in the direction of the congestion. The receiver does not participate, except through normal LLC procedures, and does not require knowledge of the transmitter's participation. The service interface between the data link layer and the network layer is also unchanged.

Definitions

k: The transmit window size in use at any given time.

kmax: The maximum transmit window size, which is the maximum value that the transmit window k may have. The value of kmax shall not exceed 127.

QUESTION 26

Which is the proper signal for pin 6 of a PHY without an internal crossover MDI Signal according to the IEEE 802.3 CSMA/CD specification?

- A. Receive +
- B. Transmit +
- C. Receive -
- D. Transmit -
- E. Contact 6 is not used.

Answer: C

Explanation:

The four pins that are used are 1, 2, 3, and 6 as shown below:

- 1 Rx+
- 2 Rx-
- 3 Tx+
- 6 Tx-

The table below shows the pin and corresponding signal for the RJ-45 connector pinouts.

RJ-45 Connector Pinout	
Pin	Signal
1	TX+
2	TX-
3	RX+
6	RX-

QUESTION 27

What is the standard transport protocol and port used for SYSLOG messages?

- A. UDP 514
- B. TCP 520
- C. UDP 530
- D. TCP 540
- E. UDP 535

Answer: A

Explanation:

For a complete list of TCP/UDP well known port numbers, see the following link:

<http://www.iana.org/assignments/port-numbers>

UDP 514

This port has been left open for use by the SYSLOG service.

TCP and UDP Ports:

In addition to the standard network ports, Cisco Works uses these TCP and UDP ports:

Port Number	Type	Description
42340	TCP	CiscoWorks2000 Daemon Manager, the tool that manages server processes
42342	UDP	Osagent
42343	TCP	JRun
42344	TCP	ANI HTTP server
7500	UDP	Electronic Switching System (ESS) Service port
7500	TCP	ESS Listening port
7580	TCP	ESS HTTP port
7588	TCP	ESS Routing port
1741	TCP	Port used for the CiscoWorks2000 HTTP server
161	UDP/TCP	Standard port for SNMP Polling
162	UDP/TCP	Standard port for SNMP Traps
514	UDP	Standard port for SYSLOG
69	TCP/UDP	Standard port for TFTP
23	TCP/UDP	Standard port for Telnet

References:

http://www.cisco.com/en/US/products/sw/ps5664/ps5669/products_tech_note09186a0080207605.shtml#udp514

http://www.cisco.com/en/US/products/sw/cscowork/ps4737/products_tech_note09186a00800e2d78.shtml

QUESTION 28

A new Syslog server is being installed in the Certkiller network to accept network management information. What characteristic applies to these Syslog messages? (Select three)

- A. Its transmission is reliable.
- B. Its transmission is secure.
- C. Its transmission is acknowledged.
- D. Its transmission is not reliable.
- E. Its transmission is not acknowledged.
- F. Its transmission is not secure.

Answer: D, E, F

Explanation:

Syslog is a method to collect messages from devices to a server running a syslog daemon. Logging to a central syslog server helps in aggregation of logs and alerts. Cisco devices can send their log messages to a Unix-style SYSLOG service. A SYSLOG service simply accepts messages, and stores them in files or prints them according to a simple configuration file. This form of logging is the best available for Cisco devices because it can provide protected long-term storage for logs. This is useful both in routine troubleshooting and in incident handling. Syslog uses UDP port 514. Since it is UDP based, the transmission is a best effort, and insecure.

Incorrect Answers:

A, C. Syslog uses UDP as the transport layer protocol, not TCP. Since UDP relies on an unreliable method of communication, syslog is not reliable.

B. Syslog has no way of providing a secure transmission by itself. Only by tunneling the syslog data through a secure channel such as IPsec can it be sent securely.

QUESTION 29

A user is having problems reaching hosts on a remote network. No routing protocol is running on the router and it's using only a default to reach all remote networks. An extended ping is used on the local router and a remote file server with IP address 10.5.40.1 is pinged. The results of the ping command produce 5 "U" characters. What does the result of this command indicate about the network?

A. An upstream router in the path to the destination does not have a route to the destination network.

B. The local router does not have a valid route to the destination network.

C. The ICMP packet successfully reached the destination, but the reply from the destination failed.

D. The ping was successful, but congestion was experienced in the path to the destination.

E. The packet lifetime was exceeded on the way to the destination host.

Answer: A

Explanation:

Even though the router is using a default route to get to all networks, at some point the packet is reaching a router that does not know how to reach the destination. The underlying reason for the failure is unknown, but when a ping is used and the response is a series of U replies, then the destination is unreachable by a router. Since the nearest router is using a default route, then the problem must be with an upstream router.

The table below lists the possible output characters from the ping facility:

Character	Description
!	Each exclamation point indicates receipt of a reply.
.	Each period indicates the network server timed out while waiting for a reply.
U	A destination unreachable error PDU was received.
Q	Source quench (destination too busy).
M	Could not fragment.
?	Unknown packet type.
&	Packet lifetime exceeded.

Incorrect Answers:

- B. The local router is using a default route, so all networks are considered to be known and reachable by the local router.
- C. If the Ping packet could reach all the way to the remote host, a "U" response would not be generated.
- D. This type of scenario would most likely result in a source quench response, which would be a Q.
- E. This would mean a "&" response, as shown in the table above.

QUESTION 30

What protocols are considered to be UDP small servers? (Choose all that apply)

- A. Echo
- B. Daytime
- C. Chargen
- D. Discard
- E. DHCP
- F. Finger

Answer: A, C, D

Explanation:

TCP and UDP small servers are servers (daemons, in Unix parlance) that run in the router which are useful for diagnostics.

The UDP small servers are:

- Echo: Echoes the payload of the datagram you send.

- Discard: Silently pitches the datagram you send.
- Chargen: Pitches the datagram you send and responds with a 72 character string of ASCII characters terminated with a CR+LF.

These 3 servers are enabled when the "service UDP-small-servers" command.

Reference:

<http://www.cisco.com/warp/public/66/23.html>

Incorrect Answers:

- B. Daytime: Returns system date and time, if correct. It is correct if you are running Network Time Protocol (NTP) or have set the date and time manually from the exec level. The command to use is telnet x.x.x.x daytime. Daytime is a TCP small server.
- E. Although DHCP uses UDP, it is not considered a UDP small server by Cisco.
- F. The router also offers finger service and async line bootp service, which can be independently turned off with the configuration global commands no service finger and no ip bootp server, respectively. This is in addition to the TCP and UDP small servers.

QUESTION 31

Which protocols are considered to be TCP small servers? (Choose all that apply).

- A. Echo
- B. Time
- C. Daytime
- D. Chargen
- E. Discard
- F. Finger
- G. DHCP

Answer: A, C, D, E

Explanation:

TCP and UDP small servers are servers (daemons, in Unix parlance) that run in the router which are useful for diagnostics.

TCP Small Servers are enabled with the service tcp-small-servers command

The TCP small servers are:

- Echo: Echoes back whatever you type by using the telnet x.x.x.x echo command.
- Chargen: Generates a stream of ASCII data. The command to use is telnet x.x.x.x chargen.
- Discard: Throws away whatever you type. The command to use is telnet x.x.x.x discard.
- Daytime: Returns system date and time, if correct. It is correct if you are running Network Time Protocol (NTP) or have set the date and time manually from the exec level. The command to use is telnet x.x.x.x daytime.

Replace x.x.x.x with the address of your router. Most routers inside Cisco run the small servers.

Incorrect Answers:

- F. DHCP is not considered a UDP small server by Cisco.
- G. The router also offers finger service and async line bootp service, which can be

independently turned off with the configuration global commands no service finger and no ip bootp server, respectively. This is in addition to the TCP and UDP small servers.

QUESTION 32

Which of the following statements are NOT true regarding the TCP sliding window protocol? (Choose all that apply)

- A. It allows the transmission of multiple frames before waiting for an acknowledgement.
- B. The size of the sliding window can only increase or stay the same.
- C. The initial window offer is advertised by the sender.
- D. The receiver must wait for the window to fill before sending an ACK.
- E. The sender need not transmit a full window's worth of data.
- F. The receiver is required to send periodic acknowledgements.

Answer: B, C

Explanation:

The sliding window algorithm allows for the window size to decrease slowing down the transmission of data. TCP uses a window of sequence numbers to implement flow control. The receiver indicates the amount of data to be sent. The receiver sends a window with every ACK that indicates a range of acceptable sequence numbers beyond the last received segment. The window allows the receiver to tell the sender how many bytes to transmit. Therefore, the statements in B and C are not true.

Incorrect Answers:

A, F. In TCP, a sender transmits only a limited amount of data before the receiver must send an acknowledgement. Windows usually include multiple packets, but if the sender doesn't get acknowledgements within a set time, all the packets must be retransmitted.
D, E. Both of these are true statements regarding the TCP sliding window mechanism.

Additional Info:

In Win 2000 and XP, the default TCP window size is 16K bytes - meaning no more than 11 frames can be outstanding without an acknowledgement. For 11 frames at 12 microseconds each, any delay of 132 microseconds or more would cause retransmissions.

QUESTION 33

A new data T1 line is being installed. What choices do you have for provisioning the framing types? (Choose all that apply)

- A. B8ZS
- B. SF
- C. AMI
- D. LLC
- E. ESF
- F. All of the above

Answer: B, E

Explanation:

SuperFraming and Extended SuperFraming are the two T1 framing types.


Incorrect Answers:

A, C. B8ZS and AMI are coding options and are not used for framing. Two typical combinations that T1's are provisioned are B8ZS/ESF and AMI/SF.

D. LLC (Logical Link Control) is not related to T1.

QUESTION 34

The Certkiller network is shown in the following exhibit:



The host sends a 1500 byte TCP packet to the Internet with the DF (Don't Fragment) bit set.

Will router CK1 be able to forward this packet to router CK2 ?

- A. Yes, it will ignore the DF bit and fragment the packet because routers do not recognize the DF bit.
- B. Yes, it will forward the packet without fragmenting it because the DF bit is set.
- C. No, it will drop the packet and wait for the host to dynamically decrease its MTU size.
- D. Yes, it will fragment the packet, and send back ICMP type 3 code 4 (fragmentation needed but DF bit set) messages back to the host.
- E. No, it will drop the packet, and send back ICMP type 3 code 4 (fragmentation needed but DF bit set) message back to the host.

Answer: E

Explanation:

Since the DF bit in the IP packet is set, the router will not be allowed to fragment the packet. Also the MTU size on the routers serial interface is restricted to 576, hence the packet will not be allowed to pass through and it will be dropped.

Incorrect Answers:

- A. Routers do indeed recognize the DF bit and will adhere to it.
- B. With the DF bit set, the packet will not be fragmented, and since 1500 bytes is too large to go through the 576 byte interface, it will be dropped.
- C. In this case, router will always send an ICMP error code back to the source stating what the problem is before dropping it.
- D. With the DF bit set the router is not allowed to fragment the packet.

QUESTION 35

With regard to TCP headers, what control bit tells the receiver to reset the TCP connection?

- A. ACK
- B. SYN
- C. SND
- D. PSH
- E. RST
- F. CLR

Answer: E

Explanation:

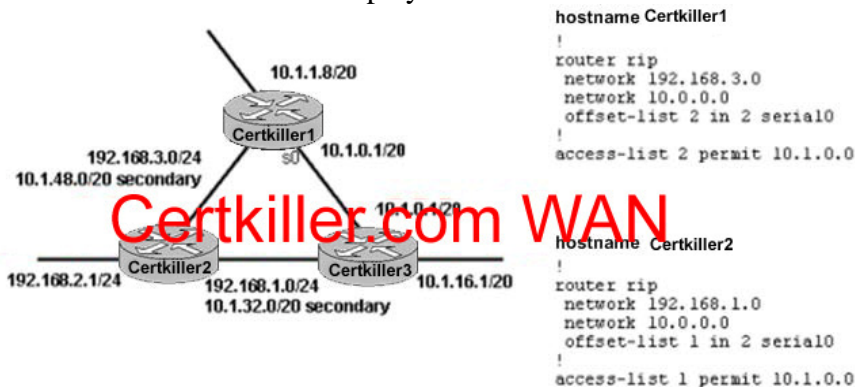
The RST flag resets the TCP connection.

Incorrect Answers:

- A. ACK is used to acknowledge data that has been sent.
- B. SYN is used to synchronize the sequence numbers.
- C. SND is not a TCP control bit.
- D. PSH is used to tell the receiver to pass the information to the application.
- F. CLR is not a valid TCP control bit.

QUESTION 36

The Certkiller RIP network is displayed below:



What statement is correct regarding the configuration in the figure?

- A. The RIP metric between Certkiller 1 and Certkiller 3 remains "1".
- B. The RIP metric between Certkiller 1 and Certkiller 3 is "2" because the offset-list command is changing the metric to "2".
- C. The RIP metric between Certkiller 1 and Certkiller 3 is "3" because the offset-list command is changing the metric to "3" by adding 2 to the existing metric.
- D. The RIP metric cannot be changed and Load sharing will be done between the two paths that exist from Certkiller 1 and Certkiller 3 (and vice-versa) because the offset-list command is specifying that there be a 2:1 load sharing ratio for the two

paths.

Answer: C

Explanation:

The offset value is added to the routing metric. An offset list with an interface type and interface number is considered extended and takes precedence over an offset list that is not extended. Therefore, if an entry passes the extended offset list and the normal offset list, the offset of the extended offset list is added to the metric. In this case, an offset of 2 is added to the routing update, making the total RIP metric 3.

QUESTION 37

Router CK1 is running BGP as well as OSPF. You wish to redistribute all OSPF routes into BGP. What command do you need to change to ensure that ALL available OSPF networks are in the BGP routing table?

- A. redistribute ospf 1 match external
- B. redistribute ospf 1 match external 1
- C. redistribute ospf 1 match external all internal all
- D. redistribute ospf 1 match internal all external 1 external 2
- E. redistribute ospf 1 match internal external 1 external 2
- F. None of the above

Answer: E

Explanation:

In this case, all OSPF routes are redistributed into BGP by using both the internal and external keywords, as shown in this Router configuration:

```
router bgp 100
redistribute ospf 1 match internal external 1 external 2
```

Reference:http://www.cisco.com/en/US/tech/CK365/technologies_tech_note09186a00800943c5.shtml

QUESTION 38

You wish to copy a file from a server into router CK1 . Which of the following IOS "copy" commands is NOT valid?

- A. copy tftp: flash:
- B. copy tftp flash
- C. copy tftp:\\flash\\
- D. copy tftp: //flash:\\

Answer: C

Explanation:

The following are valid copy commands:

/erase Erase destination file system.

bootflash: Copy from bootflash: file system

flash: Copy from flash: file system

ftp: Copy from ftp: file system

null: Copy from null: file system

nvrn: Copy from nvrn: file system

rcp: Copy from rcp: file system

system: Copy from system: file system

tftp: Copy from tftp: file system

The most common use of the copy command is the copy tftp flash command. The full syntax is copy tftp: flash: "flash file name."

QUESTION 39

On router CK1 the IOS command "ospf auto-cost reference-bandwidth 500" command was configured under the Fast Ethernet interface. Based on this, what will be the OSPF metric for this link?

- A. 1
- B. 5
- C. 50
- D. 500
- E. 5000
- F. 50000
- G. None of the above

Answer: B

Explanation:

In Cisco IOS Release 10.3 and later releases, by default OSPF will calculate the OSPF metric for an interface according to the bandwidth of the interface. For example, a 64K link will get a metric of 1562, and a T1 link will have a metric of 64.

The OSPF metric is calculated as the ref-bw value divided by the bandwidth, with ref-bw equal to 108 by default, and bandwidth determined by the bandwidth (interface) command. The calculation gives FDDI a metric of 1.

If you have multiple links with high bandwidth (such as FDDI or ATM), you might want to use a larger number to differentiate the cost on those links.

Note: The value set by the "ip ospf cost" command overrides the cost resulting from the auto-cost command.

Example:

The following example changes the cost of the Fast Ethernet link to 5, while the gigabit Ethernet link remains at a cost of 1. Thus, the link costs are differentiated.

```
router ospf 1
```

```
auto-cost reference-bandwidth 500
```

In this example, the OSPF cost is found by taking the reference bandwidth and dividing it by the bandwidth of the link, which is 100 Mbps for Fast Ethernet ($500/100 = 5$).

QUESTION 40

On your Terminal Server you are seeing spurious signals on line 6 of an asynchronous port due to contention issues. What command will fix this issue?

- A. flowcontrol hardware
- B. transport input none
- C. no exec
- D. exec-timeout 0 0

Answer: C

Explanation:

The "no exec" command is an optional command for reverse telnet configurations. Adding this line lessens the likelihood of contention over the asynchronous port. An executive process exists on all lines and buffer data to each other. At times, it can make it difficult to use a reverse telnet session. The command "no exec" will fix this.

Incorrect Answers:

- A. Console ports do not use flow control. If the terminal server is connecting to Cisco console ports then the "Flowcontrol hardware" would have no bearing.
- B. This will fundamentally cut off all telnet and reverse telnet traffic from the line.
- D. This will disable the timeout value, but will not fix problems relating to spurious signals and contention issues.

QUESTION 41

From the IOS command line interface, you accidentally press the Esc B keys while typing in a configuration line. What is the result of this action?

- A. The cursor will move to the beginning of the entire command
- B. The cursor will move back one character.
- C. The cursor will move back one word
- D. The cursor will remain in the same location.
- E. Noting, this is not a valid shortcut.

Answer: C

Explanation:

The following table describes the different shortcut options and functions that are available from the Cisco Command Line Interface:

Keystroke	Function
Ctrl-A	Jumps to the first character of the command line.

Ctrl-B or the left arrow key	Moves the cursor back one character.
Ctrl-C	Escapes and terminates prompts and tasks.
Ctrl-D	Deletes the character at the cursor.
Ctrl-E	Jumps to the end of the current command line.
Ctrl-F or the right arrow key ₁	Moves the cursor forward one character.
Ctrl-K	Deletes from the cursor to the end of the command line.
Ctrl-L; Ctrl-R	Repeats current command line on a new line.
Ctrl-N or the down arrow key ₁	Enters next command line in the history buffer.
Ctrl-P or the up arrow key ₁	Enters previous command line in the history buffer.
Ctrl-U; Ctrl-X	Deletes from the cursor to the beginning of the command line.
Ctrl-W	Deletes last word typed.
Esc B	Moves the cursor back one word.
Esc D	Deletes from the cursor to the end of the word.
Esc F	Moves the cursor forward one word.
Delete key or Backspace key	Erases mistake when entering a command; re-enter command after using this key.

Incorrect Answers:

- A. This will be the result of the Ctrl-A command.
- B. This will be the result of the Ctrl-B command, not Esc B.

QUESTION 42

Which command will display both the local and all remote SNMP engine Identification information?

- A. Show SNMP ID
- B. Show engine
- C. Show SNMP engineID
- D. Show SNMP engine ID
- E. Show SNMP stats
- F. Show SNMP mib
- G. Show SNMP users

Answer: C

Explanation:

The following is a sample output from a Cisco router:

Certkiller #show snmp ?

```
mib          show mib objects context
engineID   show local and remote SNMP engine
group        IDs show SNMPv3 groups snmp
pending      manager pending requests snmp manager
sessions     sessions show snmp statistics show
stats user   SNMPv3 users
```

| Output modifiers

<cr>

Certkiller #

Reference: CCO login required.

http://www.cisco.com/en/US/products/sw/iosswrel/ps5207/products_command_reference_chapter09186a00801a809e.html#1030651

QUESTION 43

What would occur as a result of the clear ip route * command being issued? (Choose two)

- A. A router would recalculate its entire table and re-establish its neighbor relationships.
- B. A router would recalculate its entire routing table but its neighbor relationship would not be affected.
- C. Only link state routing protocols would be recalculated and only those neighbor relationships re-established.
- D. Only the routing table would be recalculated.
- E. Only its neighbor relationship would be re-established.

Answer: B, D

Explanation:

The use of the * means that all routing table entries will be deleted within the routing table, forcing the router to calculate a new routing table. The underlying neighbor adjacencies are not affected by this command. To force the router to re-establish neighbor relationships, the "clear ip xxx neighbor" command, where xxx is the routing protocol in use. For example, to clear all of the OSPF neighbor relationships, use the "clear ip ospf neighbor" command.

QUESTION 44

Which IOS example will configure a remote user in a group called remotegroup to receive traps at the v3 security model and the authNoPriv security level?

- A. snmp engineid remote 16.20.11.14 000000100a1ac151003
snmp enable traps config
snmp manager
- B. snmp-server group remotegroup v3 noauth
snmp-server user remote remotegroup remote 16.20.11.14 v3
snmp-server host 16.20.11.14 inform version 3 noauth remoteuser config
- C. snmp-server group remotegroup v3 noauth
snmp-server user remoteAuthUser remoteAuthGroup remote 16.20.11.14 v3 auth
md5 password1
- D. snmp-server group remotegroup v3 priv
snmp-server user remote PrivUser remotePrivGroup remote 16.20.11.14 v3 auth
md5 password1 priv des56 password2

Answer: B

Explanation:

snmp-server user:

To configure a new user to a Simple Network Management Protocol group, use the snmp-server user global configuration command. To remove a user from an SNMP group, use the no form of the command.

snmp-server user username [groupname remote ip-address [udp-port port] {v1 | v2c | v3} [encrypted] [auth {md5 | sha} auth-password [priv des56 priv password]] [access access-list]

no snmp-server user

Syntax Description

<i>username</i>	The name of the user on the host that connects to the agent.
<i>groupname</i>	(Optional) The name of the group to which the user is associated.
<i>remote</i>	(Optional) Specifies the remote copy of

	SNMP on the router.
<i>ip-address</i>	(Optional) The IP address of the device that contains the remote copy of SNMP.
<i>udp-port</i>	(Optional) Specifies a UDP port of the host to use.
<i>port</i>	(Optional) A UDP port number that the host uses. The default is 162.
<i>v1</i>	(Optional) The least secure of the possible security models.
<i>v2c</i>	(Optional) The second least secure of the possible security models. It allows for the transmission of informs and counter 64, which allows for integers twice the width of what is normally allowed.
<i>v3</i>	(Optional) The most secure of the possible security models.
<i>encrypted</i>	(Optional) Specifies whether the password appears in encrypted format (a series of digits, masking the true characters of the string).
<i>auth</i>	(Optional) Initiates an authentication level setting session.
<i>md5</i>	(Optional) The HMAC-MD5-96 authentication level.
<i>sha</i>	(Optional) The HMAC-SHA-96 authentication level.
<i>auth-password</i>	(Optional) A string (not to exceed 64 characters) that enables the agent to receive packets from the host.
<i>priv</i>	(Optional) The option that initiates a privacy authentication level setting session.
<i>des56</i>	(Optional) The CBC-DES privacy authentication algorithm.
<i>priv password</i>	(Optional) A string (not to exceed 64 characters) that enables the host to encrypt the contents of the message it sends to the agent.
<i>access</i>	(Optional) The option that enables you to specify an access list.
<i>access-list</i>	(Optional) A string (not to exceed 64 characters) that is the name of the access list.

Incorrect Answers:

- A. The SNMP engineid is an invalid command.
- C, D. In these examples, the MD5 authentication level is used. In this question we want the user to use no authentication.

QUESTION 45

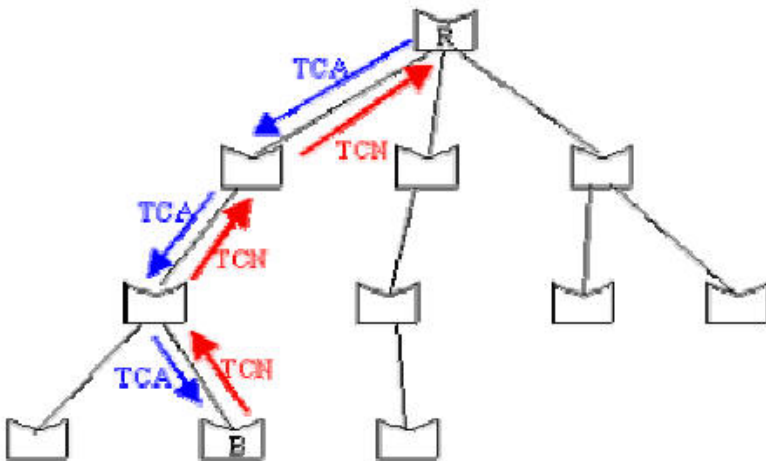
In a bridged LAN, the number of BPDU's with the TCA bit set is incrementing rapidly. What could be the cause of this? (Choose all that apply).

- A. BPDU's with the TCA bit set is part of the normal operation of a bridged LAN.
- B. Improper cabling is being used in the network.
- C. There is no spanning tree portfast configured on the ports connecting 2 workstations.
- D. The root switch is experiencing problems due to high CPU utilization and is not sending any BPDUs.
- E. None of the above.

Answer: B, C

Explanation:

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



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 notifying the topology change will

not stop sending its TCN until the designated bridge has acknowledged it, so the designated bridge answers the TCN even though it does not receive configuration BPDU from its root.

The portfast feature is a Cisco proprietary change in the STP implementation. The command is applied to specific ports and has two effects:

- Ports coming up are put directly in the forwarding STP mode, instead of going through the learning and listening process. Note that STP is still running on ports with portfast.
- The switch never generates a TCN when a port configured for portfast is going up or down.

Reference:

http://www.cisco.com/en/US/tech/CK389/CK621/technologies_tech_note09186a0080094797.shtml#portfastcommand

QUESTION 46

The Certkiller LAN is a bridged network running the 802.1D spanning tree protocol. Which of the following are parameters that a bridge will receive from the root bridge.

- A. Maxage
- B. Root Cost
- C. Forward delay
- D. A,B, and C
- E. None of the above

Answer: D

Explanation:

A, B and C are all located in the BPDU which each switch gets from the root bridge. The BPDUs are in the following format:

2	1	1	1	8	4	8	2	2	2	2	2	Octets
Protocol ID	Version	Message Type	Flags	Root ID	Root Cost	Bridge ID	Port ID	Message Age	Max Age	Hello Time	Forward Delay	

- Protocol ID - indicates that the packet is a BPDU.
- Version - the version of the BPDU being used.
- Message Type - the stage of the negotiation.
- Flags - two bits are used to indicate a change in topology and to indicate acknowledgement of the TCN BPDU.
- Root ID - the root bridge priority (2 bytes) followed by the MAC address (6 bytes).
- Root Path Cost - the total cost to from this particular bridge to the designated root bridge.
- Bridge ID - the bridge priority (2 bytes) followed by the MAC address (6 bytes), lowest value wins! The default bridge priority is 0x8000 (3276810).
- Port ID - the ID of the port from which are transmitted the BPDUs, a root port, this is made up of the configured port priority and the bridge MAC address.
- Message Age - timers for aging messages (only has effect on the network if the

root bridge is configured with this parameter).


- Maximum Age - the maximum message age before information from a BPDU is dropped because it is too old and no more BPDUs have been received. (only has effect on the network if the root bridge is configured with this parameter). The default value for this is 20 seconds.
- Hello Time - the time between BPDU configuration messages sent by the root bridge (only has effect on the network if the root bridge is configured with this parameter). The default value for this is 2 seconds.
- Forward Delay - this temporarily stops a bridge from forwarding data to give a chance for information of a topology change to filter through to all parts of the network. This means that ports that need to be turned off in the new topology have a chance to be switched off before the new ports are turned on (only has effect on the network if the root bridge is configured with this parameter).

Reference:

<http://www.rhyshaden.com/ethernet.htm>

QUESTION 47

A small office LAN contains only one switch, which was put in place without any of the default configurations changed. You have noticed that somebody in the office has looped a cable by connecting one end to port 4/37 and the other to port 4/38 as shown below:



All links are 10/100

Configuration is default

Which of the following statements is true?

- A. Port 4/38 will be blocked.
- B. Both ports will be forwarding.
- C. Port 4/37 will be blocking.
- D. Both ports will be blocked.
- E. Port 4/38 will continuously move between the listening and learning states.
- F. Port 4/37 will be stuck in the learning state.

Answer: A

Explanation:

Port priority is based on lowest priority, and lowest port number. Because of this, then 4/37 would become the root port and 4/38 would be blocking. The default mode of a

Catalyst switch is to enable the STP process for all VLANs.

Incorrect Answers:

B. Even though this switch will effectively become the root switch, and all ports in a root switch should be in the "forwarding state" a loop will occur in this case, and so one of the ports must be blocking. Since the priority of 4/38 is lower by default, it will be blocking.

QUESTION 48

Which of the following statements regarding Transparent Bridge tables are FALSE? (Choose all that apply.)

- A. Decreasing the bridge table aging time would reduce flooding.
- B. Increasing the bridge table aging time would reduce flooding.
- C. Bridge table entries are learned by way of examining the source MAC address of each frame.
- D. Bridge table entries are learned by examining destination MAC addresses of each frame.
- E. The bridge aging time should always be more than the aggregate time for detection and recalculation of the spanning tree.

Answer: A, D

Explanation:

Basic fundamental behind TB is to learn the network topology by means of storing the source MAC address of a packet, and the corresponding interface from which the packet came in on the network. This information is stored in the bridge table. To keep the bridge table small and manageable entries are deleted after a specified period of time, known as bridge table aging time. Once an entry is removed from the bridge table, and a packet arrives for which the information is no longer there in the bridge table, the packet will be flooded out of all interfaces except the interface on which it was received.

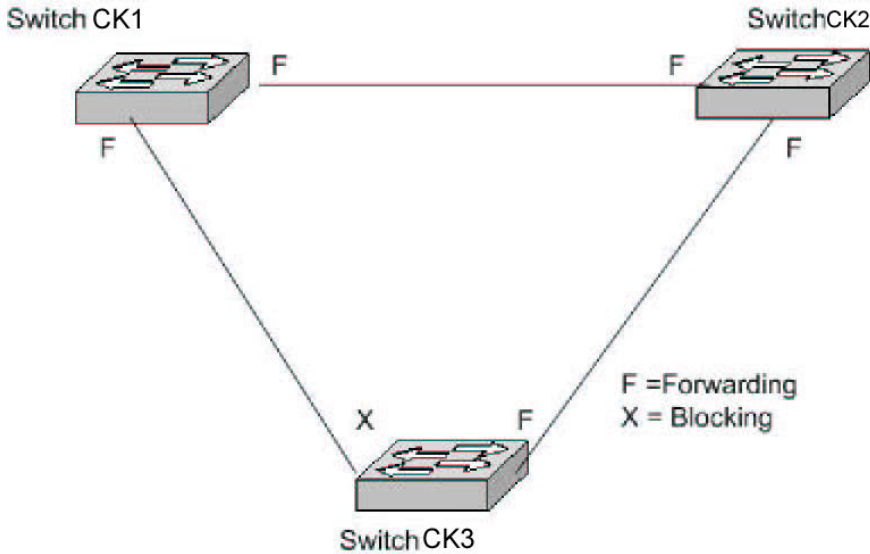
An increase in the bridge table aging time will reduce flooding.

Incorrect Answers:

- B. Increasing the aging table time will indeed reduce the flooding, since the source MAC addresses are cached for a longer period of time.
- C. This statement is also true. Bridge tables are built by looking at the source MAC address to learn which stations are attached to the bridge ports.
- E. The aging time should indeed be longer than the convergence time for the spanning tree algorithm in order to prevent information from timing out and being re-learned, which will just begin the STP process again.

QUESTION 49

The Certkiller network is shown in the following exhibit:



All links are 1 Gb
Configuration is default.

You issue the "set spantree root 1" command on Switch CK1 . What will happen as a result of this change? (Choose all that apply).

- A. No other switch in the network will be able to become root as long as Switch CK1 remains up and running in this topology.
- B. Switch CK1 will change its Spanning Tree priority to become the root for Vlan 1, only.
- C. The port that used to be blocking on Switch CK3 will be changed to forwarding.
- D. The link between Switch CK1 and Switch CK2 will remain forwarding throughout the reconvergence of the Spanning Tree domain.

Answer: B, C, and D

Explanation:

The syntax specified in this question only makes CK1 root for Vlan 1 only.

The set spantree root {vlan_id} command sets the priority of the switch to 8192 for the VLAN or VLANs specified in {vlan_id} .

Note: The default priority for switches is 32768. This command setting means that the Certkiller switch will be selected as the root switch because it has the lowest priority.

```
Certkiller -Switch> (enable) set spantree root 1
```

```
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.

Certkiller -Switch> (enable)

Because Switch 1 will become the root, the link between CK 3 and CK 1 will be forwarding, which means the link between CK2 and CK3 will change from forwarding to blocking.


Incorrect Answers:

A. It is important to remember that the "set spantree root" command merely changes the spanning tree switch priority from 32768 to 8192 so that it is much more likely to become the root for the particular VLAN. If another switch already has a priority lower than 8192, then this command will make the switch the root by lowering it to one below the existing value. For example, if another switch is already configured with a priority of 8192, then issuing the "set spantree root" command will configure the new switch with a priority of 8191. However, another switch could still become the root if it were configured with a lower priority after this command was issued on another switch.

Reference: Cisco LAN Switching, Clark and Hamilton, Cisco Press, Page 197.

QUESTION 50

The Certkiller network consists of only 2 routers as below:



Router CK1

Router CK2

You perform the following router configurations:

Router CK1 :

```
no ip routing
```

```
!
```

```
interface Ethernet 0
```

```
no ip address
```

```
bridge-group 1
```

```
!
```

```
interface Ethernet 1
```

```
no ip address
```

```
bridge-group 1
```

```
!
```

```
bridge 1 protocol ieee
```

Router CK2 :

```
no ip routing
```

```
!
```

```
interface Ethernet 0
```

```
no ip address
```

```
bridge-group 2
```

```
!
```

```
interface Ethernet 1
```

```
no ip address
```

```
bridge-group 2
```

```
!
```

```
bridge 2 protocol ieee
```

```
bridge 2 priority 63500
```

Based on this configuration, which router will become the root, and which ports will be forwarding?

A. Router CK2 will become the root.

One port on Router CK1 will be forwarding, and the other will be blocking.

One port on Router CK2 will be forwarding, and the other will be blocking.

B. Both Router CK1 and Router CK2 will become the root in an independent spanning tree.

All ports on Router CK1 and Router CK2 will be forwarding.

C. Router CK1 will become the root.

Both ports on Router CK1 will be forwarding.

Both ports on Router CK2 will be forwarding.

D. Router CK2 will become the root.

Both ports on Router CK1 will be forwarding.

One port on Router CK2 will be forwarding, and the other will be blocking.

E. Router CK1 will become the root.

Both ports on Router CK1 will be forwarding.

One port on Router CK2 will be forwarding, and the other will be blocking.

Answer: E

Explanation:

Bridge 1's priority is at default 32768, Bridge 2 is at 63500, Bridge 1 (with a lower Bridge ID) will be Root Bridge. All ports on the root bridge are always in forwarding state, hence both the ports on Bridge 1 will be in forwarding state. As per STP any other Bridge can only have one connection to the Root Bridge in the forwarding state, hence only one port on Bridge 2 will be forwarding.

Incorrect Answers:

A, D. CK2 has a bridge priority configured as 63500, while CK1 is left with the default.

Since the default value is 32768 and lower is preferred, CK1 will become the root.

C. Only the single root port will be forwarding.

QUESTION 51

What spanning-tree protocol timer determines how often the root bridge send configuration BDPUs?

A. STP Timer

B. Hold Timer

C. Hello Timer

D. Max Age Timer

E. Forward Delay Timer

Answer: C

Explanation:

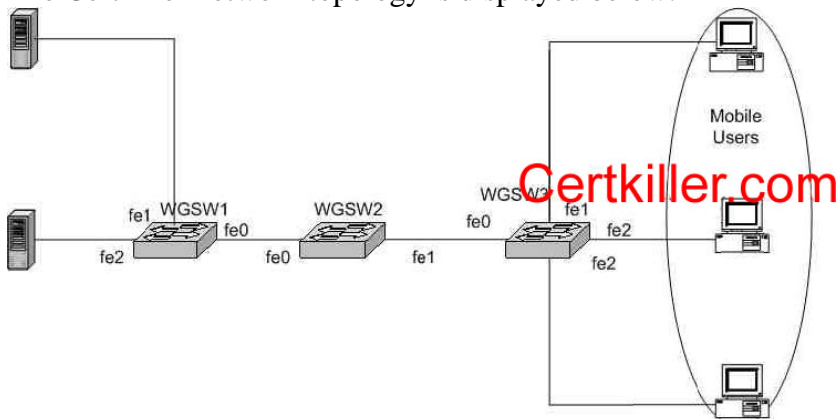
The STP Hello Time is the time between each Bridge Protocol Data Unit (BPDU) that is sent on a port. This is equal to two seconds by default, but can be tuned to be between one and ten seconds.

Incorrect Answers:

- A. The Max Age, Forward Delay, and Hello Timers are all considered to be STP timers.
- B. Hold timers are used in routing protocols to avoid inconsistent information and loops, but they are not an STP timer.
- D. 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.
- E. The Forward Delay Timer 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.

QUESTION 52

The Certkiller network topology is displayed below:



WGSW3 has been set up to provide access to mobile users in a conference room. Portfast has been enabled on all access ports. The following command is entered on WGSW3:

```
Switch(config)#spanning-tree portfast bpduguard
```

What happens if a switch or bridge is connected to one of the access ports?

- A. Any access port that receives a BPDU packet will be disabled.
- B. The access port will reject any BPDU packets that they receive.
- C. Portfast will be disabled on any access port that receives a BPDU packet.
- D. The bridge can join the BPDU topology, but it is blocked from becoming the root bridge.
- E. Only BPDU packets that are NOT superior to the current root bridge will be accepted on the access port.

Answer: A


Explanation:

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:

```
2000 May 12 15:13:32 %SPANTREE-2-RX_PORTFAST:Received BPDU on PortFast enable port.  
Disabling 2/1  
2000 May 12 15:13:32 %PAGP-5-PORTFROMSTP:Port 2/1 left bridge port 2/1
```

QUESTION 53

You are having connectivity problems with the network shown below:



Router CK2 is able to ping the Catalyst switch CK3 , but router CK1 cannot. What is the probable cause of this problem?

- A. There is no VTP domain on the Catalyst switch.
- B. The incorrect VLAN is attached to the command interface of the Catalyst.
- C. There is no default route configured on the switch.
- D. An incorrect IP address on the switch.
- E. ICMP packets are being filtered on the switch CK3

Answer: C

Explanation:

Without a default route on Cat CK3 , CK3 will not know how to get packets back to CK1 . Catalyst CK3 would be able to ping router CK2 without a default route, however, because they share the same IP subnet.

Incorrect Answers:

- A, B. VTP and VLAN information that is configured incorrectly could explain problems associated with local LAN users attached to the CK3 , but this would not explain why CK1 would not be able to reach CK3 .
- D. If CK3 had an incorrect IP address, then CK2 would not be able to ping CK3 .
- E. If all ICMP packets were filtered, then CK2 would also not be able to ping CK3 . This answer could be the problem only if ICMP were being filtered from router CK1 .

QUESTION 54

What is the Cisco recommended best practice PaGP setting for ports Etherchannel trunks?

- A. on - on
- B. auto - auto
- C. desirable - on
- D. desirable - auto
- E. desirable - desirable

Answer: E

Explanation:

Using PAgP to Configure EtherChannel (Recommended)

PAgP facilitates the automatic creation of EtherChannel links by exchanging packets between channel-capable ports. The protocol learns the capabilities of port groups dynamically and informs the neighboring ports.

After PAgP identifies correctly paired channel-capable links, it groups the ports into a channel. The channel is then added to the spanning tree as a single bridge port. A given outbound broadcast or multicast packet is transmitted out one port in the channel only, not out every port in the channel. In addition, outbound broadcast and multicast packets transmitted on one port in a channel are blocked from returning on any other port of the channel.

There are four user-configurable channel modes: on, off, auto, and desirable. PAgP packets are exchanged only between ports in auto and desirable mode. Ports configured in on or off mode do not exchange PAgP packets. For switches to which you want to form an EtherChannel, it is best to have both switches set to desirable mode. This gives the most robust behavior if one side or the other encounters error situations or is reset. The default mode of the channel is auto.

Both the auto and desirable modes allow ports to negotiate with connected ports to determine if they can form a channel. The determination is based on criteria such as port speed, trunking state, and native VLAN.

Ports can form an EtherChannel when they are in different channel modes as long as the modes are compatible. This list provides examples:

- A port in desirable mode can successfully form an EtherChannel with another port that is in desirable or auto mode.
- A port in auto mode can form an EtherChannel with another port in desirable mode.
- A port in auto mode cannot form an EtherChannel with another port that is also in auto mode, since neither port initiates negotiation.
- A port in on mode can form a channel only with a port in on mode because ports in on mode do not exchange PAgP packets.
- A port in off mode cannot form a channel with any port.

Reference:http://www.cisco.com/en/US/tech/CK3_89/CK2_13/technologies_tech_note09186a00800949c2.shtml#pagptoconfig

Additional Information:

The Best practices for Cisco Catalyst switch configurations can be found in this document:

http://www.cisco.com/en/US/products/hw/switches/ps700/products_white_paper09186a00801b49a4.shtml#cg6

From this Best Practices document:

Cisco Configuration Recommendation for L2 Channels

Cisco recommends enabling PAgP and using a setting of desirable-desirable on all EtherChannel links. Refer to the output below for more information:

```
Switch(config)#interface type slot/port
```

```
Switch(config-if)#no ip address
```

!--- Ensures that there is no IP

!--- address assigned to the LAN port.

```
Switch(config-if)#channel-group <number> mode desirable
```

!--- Specify the channel number and the PAgP mode.

Verify the configuration, as shown below.

```
Switch#show run interface port-channel number
```

```
Switch#show running-Config interface type slot/port
```

```
Switch#show interfaces type slot/port etherchannel
```

```
Switch#show etherchannel <number> port-channel
```

QUESTION 55

You wish to implement Ethernet Channels in your switched LAN. Which of the following are valid statements that should be kept in mind before this implementation? (Choose all that apply)

- A. Ports within a Fast Ether Channel need to have identical duplex and speed settings.
- B. Port Aggregation Protocol (PAgP) facilitates the automatic creation of Fast Ether channels links.
- C. Ports within a Fast Ether Channel may be assigned to multiple VLANs.
- D. Fast Ethernet Channels can not be configured as a trunk.
- E. Only Fast Ethernet ports can be channeled.

Answer: A, B

Explanation:

You can not mix and match different types of Ethernet ports, such as 10M, 100M, GIGE, etc into the same channel. All ports in the channel need to have the same speed settings.

Similarly, all ports need to be configured to have identical duplex settings.

The Port aggregation protocol (PAgP) aids in the automatic creation of Fast EtherChannel links. PAgP packets are sent between Fast EtherChannel-capable ports in order to negotiate the forming of a channel.

Incorrect Answers:

C. Ports in the channel can only be assigned to one VLAN.

D. Ethernet channels can indeed be set up as trunks.

E. Ethernet channels can be set up for fast Ethernet, Gigabit Ethernet, and 10 Gigabit Ethernet.

QUESTION 56

A new Catalyst switch is added to the Certkiller switched LAN. Users attached to the

new switch are having connectivity problems. Upon troubleshooting, you realize that the new switch is not dynamically learning any VLAN information via VTP from the other switches. What could be causing this problem?

- A. The other switches are different Catalyst models.
- B. There are no users on one of the existing switches.
- C. The other upstream switches are VTP clients.
- D. The VTP domain name is not properly configured.
- E. The native VLAN on the trunk is VLAN 1.

Answer: D

Explanation:

In order for VTP information to be propagated throughout the network, every LAN switch participating in the VTP domain must have the exact same VTP domain name configured.

Incorrect Answers:

- A. All Catalyst switch models support VTP.
- B. The number of users or types of devices attached to any switch has absolutely no bearing on the functionality of VTP.
- C. VTP clients can pass updates to each other to propagate VLAN info throughout the network. All VTP client switches do not necessarily need to be directly connected to a VTP server.
- E. VLAN 1 is the default VLAN for all Catalyst switches. Although it is not necessarily recommended that all switches use this default VLAN, VTP information would be able to pass throughout the network if they did.

QUESTION 57

The Certkiller network is implementing a new Layer 3 Switching architecture. When an IP packet is Layer 3-switched from a source in one VLAN to a destination in another VLAN, what field in a packet will be rewritten?

- A. Layer 3 destination address
- B. Layer 3 source address
- C. Layer 2 TTL
- D. Layer 3 TTL
- E. Layer 3 Transport Protocol

Answer: D

Explanation:

When a packet is Layer 3 switched, the source and destination MAC address, as well as the IP TTL and IP checksum is rewritten.

	Layer 2 Ethernet		Layer 3 IP Header				Data	FCS
--	-----------------------------	--	------------------------------	--	--	--	-------------	------------

		Header							
	Destination MAC	Source MAC	Destination IP	Source IP	TTL	Checksum			
Received Frame	Router MAC Address	Host-A MAC Address	Host-B	Host- A	n	value1			
Rewritten Frame	Next Hop MAC Address	Router MAC Address	Host-B	Host- A	n-1	value2			

The Table above displays the details of the received frame that are indicated and then the details required for the rewritten frame that is transmitted after routing are shown. Notice that the following fields must be modified for the rewritten frame that is forwarded to the next hop routing device:

- Destination MAC address: The MAC address of the next hop must be written to the rewritten frame.
- Source MAC address: The source MAC address must be written to the MAC address of the router.
- IP TTL: This must be decremented by one, as per the normal rules of IP routing.
- IP Header Checksum: This must be recalculated, as the TTL field changes.

The process of how the data plane operations shown in Table 6-1 are implemented is where the difference between a traditional router and Layer 3 switch lie. A traditional router uses the same general purpose CPU used to perform control plane operations to also implement data plane operations, meaning data plane operations are handled in software. A Layer 3 switch on the other hand uses an ASIC to perform data plane operations because it is very easy to program the very simple operations required for the data plane into an ASIC. In this respect, the data plane is implemented in hardware because a series of hardware operations are programmed into the ASIC that perform the data plane operations required for routing a packet.

Reference: Justin Menga, CCNP Practical Studies: Layer 3 switching.

QUESTION 58

By default, which of the following VLANs are eligible for pruning in a Catalyst 6509 switch? (Choose all that apply)

- A. VLAN 1
- B. VLAN 2
- C. VLAN 999
- D. VLAN 1000
- E. VLAN 1001
- F. VLAN 4094

Answer: B, C, D

Explanation:

By default, VLANs 2-1000 are pruning eligible in a Catalyst switch. For the default VLAN settings in Catalyst switches see the following document:

http://www.cisco.com/en/US/partner/products/hw/switches/ps708/products_configuration_guide_chapter09186a008007d4c3.html

QUESTION 59

You have ISL trunks configured between two Catalyst switches, and you wish to load share traffic between them. Which method of load sharing can you utilize?

- A. Load sharing of traffic over parallel ISL trunks on a per flow basis.
- B. Load sharing of traffic over parallel ISL trunks on a per VLAN basis.
- C. Load sharing of traffic over parallel ISL trunks on a per packet basis.
- D. Automatic round robin load sharing of VLAN traffic.

Answer: B

Explanation:

It is possible to load share over parallel ISL trunks on a per-VLAN basis, using either path costs or port priorities, or a combination of these two methods. However, this will only load share traffic from different VLANs, and not evenly distribute traffic from the same VLAN as the STP process will only allow a single VLAN to use one of the ISL trunks.

Incorrect Answers:

- A, C. It is not possible to load share on a per flow or per packet basis as any given VLAN will only traverse over one of the ISL trunks. The other trunk will be in a blocking state for that particular VLAN.
 - D. Automatic load sharing is not possible over parallel ISL trunks.
-

QUESTION 60

You are trying to bring up an ISL trunk link between two switches. The trunk mode on the local end is set to auto. However, the ISL trunk never comes up. What is the probable cause of this problem? (Choose all that apply.)

- A. The trunk mode on the remote end is set to on.
- B. The trunk mode on the remote end is set to off.
- C. The trunk mode on the remote end is set to auto.
- D. The trunk mode on the remote end is set to desirable.
- E. The trunk mode on the remote end is set to nonegotiate.

Answer: B, C, E

Explanation:

The trunk mode can be: auto, Desirable, On, nonegotiate, and Off. When set to "off" ISL is not allowed on this port regardless of the mode configured on the other end. When set to "auto" the port listens for Dynamic Trunking Protocol (DTP) frames from the remote

device. Auto does not propagate any intent to become a trunk; it is solely dependent on the remote device to make the trunking decision. Thus, if both ends are set to Auto, no trunking will occur. When set to nonegotiate, DTP is not spoken to the neighboring switch. nonegotiate automatically enables ISL trunking on its port, regardless of the state of its neighboring switch. However, according to Cisco when one end is set to auto, and the other end is set to nonegotiate, then the result is a non-trunking port (see the table at the middle of the Cisco link, used as a reference).

Incorrect Answers:

A. When set to "on", DTP is spoken to the neighboring switch. On automatically enables ISL trunking on its port, regardless of the state of its neighboring switch. It remains an ISL trunk unless it receives an ISL packet that explicitly disables the ISL trunk. The Cisco TAC recommends that desirable trunking mode be configured on the ports.

D. In desirable mode, DTP is spoken to the neighboring switch. Desirable communicates to the neighboring switch that it is capable of being an ISL trunk, and would like the neighboring switch to also be an ISL trunk.

Reference:

http://www.cisco.com/warp/public/793/lan_switching/2.html

QUESTION 61

The Certkiller corporate LAN consists of numerous Catalyst switches and a large number of VLANs. You are seeing an excessive amount of broadcasts across your trunk links. In an effort to reduce unnecessary traffic, VLAN Trunk Protocol (VTP) pruning is enabled. Which of the following statements is true regarding this change?

- A. Traffic on VLAN 1 can be pruned.
- B. Pruning eligibility is determined by the amount of ports assigned to a VLAN.
- C. VTP pruning is a way to detect the removal of a VLAN within a VTP domain.
- D. VTP version 2 is backward compatible with VTP version 1.
- E. VTP pruning only affects traffic from VLANs that are pruning eligible.

Answer: E

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 default, VTP pruning is disabled. VTP pruning does not prune traffic from VLANs that are pruning-ineligible.

Incorrect Answers:

A. VLAN 1 is always pruning-ineligible, meaning traffic from VLAN 1 cannot be pruned.

B. Pruning eligibility is based only on the VLANs that need the given broadcast information across the trunks. It has nothing to do with the number of ports assigned to

that VLAN.

C. VTP Pruning simply reduces the broadcast and multicast traffic. It does not change, add, or delete the VLANs in a VTP domain.

D. VTP version 1 and VTP version 2 are not interoperable on network devices in the same VTP domain. Every network device in the VTP domain must use the same VTP version. Do not enable VTP version 2 unless every network device in the VTP domain supports version 2.

Reference:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/rel_4_2/config/vlans.htm#xtocid79802

QUESTION 62

After performing some testing on a Catalyst switch in a lab, it is connected to the production network to another Catalyst switch via the supervisor Gigabit Ethernet port. Soon after this, users complain that they have lost all connectivity to the network.

What could have caused this to happen?

A. You did not issue the set spantree uplinkfast enable 1/1 command before connecting to the corporate switch.

B. You did not make the trunk mode set to on or desirable for the trunk to the supervisor of the other switch.

C. You did not make the VTP mode transparent in the new switch.

D. The dynamic CAM entries were not cleared after the new switch was connected to the network.

E. The new switch had the wrong VTP domain name.

Answer: C

Explanation:

The most likely cause of this happening is that the new switch was configured to participate in the VTP domain, but that it was set to server mode. The default mode is VTP server, which can override the VLAN information and get propagated to other switches in the network. In transparent mode, the switch will not participate in VTP, and it cannot override existing VTP settings.

Understanding the VTP Domain:

A VTP domain (also called a VLAN management domain) is made up of one or more interconnected network devices that share the same VTP domain name. A network device can be configured to be in one and only one VTP domain. You make global VLAN configuration changes for the domain using either the command-line interface (CLI) or Simple Network Management Protocol (SNMP).

By default, the Catalyst 6500 series switch is in VTP server mode and is in the nomanagement domain state until the switch receives an advertisement for a domain over a trunk link or you configure a management domain.

If the switch receives a VTP advertisement over a trunk link, it inherits the management domain name and the VTP configuration revision number. The switch ignores

advertisements with a different management domain name or an earlier configuration revision number.

If you configure the switch as VTP transparent, you can create and modify VLANs but the changes affect only the individual switch.

When you make a change to the VLAN configuration on a VTP server, the change is propagated to all network devices in the VTP domain. VTP advertisements are transmitted out all trunk connections.

VTP maps VLANs dynamically across multiple LAN types with unique names and internal index associations. Mapping eliminates excessive device administration required from network administrators.

Understanding VTP Modes:

You can configure a Catalyst 6500 series 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 network devices in the same VTP domain and synchronize their VLAN configuration with other network devices 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 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.

Incorrect Answers:

A. The set spantree uplinkfast enable command increases the path cost of all ports on the switch, making it unlikely that the switch will become the root switch. This obviously would not cause the problem described in this question.

B. This would affect the trunk coming up between the switches, but would not cause this kind of connectivity issue in this question. In this case, even if the trunk did not come up, end users would not even notice.

D. The CAM entries would have no impact, especially since no end stations were plugged into it in the lab.

E. The wrong VTP domain name would mean that this switch would not be participating in this particular VTP domain. In this specific case, this would have actually fixed the problem.

QUESTION 63

You are trying to set up an Ethernet channel between switch A and switch B. After issuing the command "set port channel 3/1-2 on" on switch B, connectivity to switch B is lost. The following messages appear on switch B as a result of this.

Switch-B> (enable)

```
%SPANTREE-2-CHNMISCFG: STP loop - channel 3/1-2 is disabled in vlan 1.
```

%PAGP-5-PORTFROMSTP:Port 3/1 left bridge port 3/1-2
%PAGP-5-PORTFROMSTP:Port 3/2 left bridge port 3/1-2
You then disable the Ethernet Channel on Switch-B, but you still have no connectivity to Switch-A.
Which command will restore connectivity to switch A?

- A. clear port error 3/1-2
- B. set port enable 3/1-2
- C. set trunk channel 3/1-2 desirable isl
- D. set port channel 3/1-2 enable

Answer: B

Explanation:

The message clearly indicates that the ports 2/1-4 have been disabled. This is a consequence of spantree as shown by the "channel 3/1-2 is disabled in vlan 1" message. This will make the ports affected go into an err-disable state. To fix this, the ports need to be manually re-enabled with the "set port enable" command.

QUESTION 64

A switch is configured for an ISL trunk, with the trunk mode set to on. A new switch is added to the network, but the trunk will not come up.
What is the probable cause of this problem?

- A. The native VLANs are not the same.
- B. The trunks need to be set to "on" or "auto".
- C. The trunks need to be set to "desirable" or "nonegotiate".
- D. The VTP domain names carried in the Dynamic Inter-Switch Link (DISL) messages are not the same.
- E. The Unidirectional Link Detection timers are shorter than the Spanning Tree Protocol (STP) timers.

Answer: D

VTP domain names on an ISL trunk must be the same. DTP packets will not pass between switches that are in different VTP domains.

Incorrect Answers:

- A. The VLANs can be different for each switch and the trunk will still come up if set up correctly.
- B, C. Since one end of the trunk is set to on, the other end can be set to either on, auto, desirable, or nonegotiate for the trunk to come up.
- E. These timers will have no bearing on the trunk formation.

Reference:

http://www.cisco.com/warp/public/793/lan_switching/2.html

QUESTION 65

You are designing a new switched LAN and VLAN information will need to be shared between switches. What VLAN trunking protocol contains the following features?

- 26 byte header and a 4 byte frame check sum
- Supports up to 1024 VLANs
- Supports a single instance of spanning tree per-VLAN

- A. ISL
- B. 802.1d
- C. 802.1q
- D. 802.1v
- E. 802.10

Answer: A

Explanation:

ISL is a Cisco proprietary protocol for interconnecting multiple switches and maintaining VLAN information as traffic goes between switches. ISL uses Per VLAN Spanning Tree (PVST) which runs one instance of Spanning Tree Protocol (STP) per VLAN. PVST allows for optimizing the root switch placement for each VLAN and supports load balancing of VLANs over multiple trunk links.

With ISL, an Ethernet frame is encapsulated with a header that transports VLAN IDs between switches and routers. A 26-byte header that contains a 10-bit VLAN ID is prepended to the Ethernet frame. This 10 byte-VLAN ID provides for up to 1024 VLANs. The FCS field consists of four bytes in an ISL packet. This sequence contains a 32-bit CRC value, which is created by the sending MAC and is recalculated by the receiving MAC to check for damaged frames. The FCS is generated over the DA, SA, Length/Type, and Data fields. When an ISL header is attached, a new FCS is calculated over the entire ISL packet and added to the end of the frame

QUESTION 66

A switch can belong to how many VTP domains?

- A. 1
- B. 2
- C. 1 to 1005
- D. 1 to 4096
- E. It depends upon memory
- F. It depends on the number of available IDB blocks

Answer: A

Explanation

A Catalyst switch can only be configured to belong in only one VTP domain, using the "set VTP domain" command. If you attempt to use additional "set vtp domain"

commands, you will simply overwrite the previous command and the switch will belong to the newly configured domain.

QUESTION 67

What is a key advantage to configuring all switches in an enterprise network to VTP transparent mode?

- A. Ensures consistency between VLAN numbering for all switches in the switched network.
- B. Prevents network administrator's from accidentally deleting VLAN information from all switches.
- C. Allows for more rapid deployment of VLANs throughout the enterprise.
- D. Reduces the size of the spanning tree network, and as a result, improves STP convergence time.
- E. Reduces the total number of VLANs required in the enterprise network.

Answer: B

Explanation:

A major advantage to configuring all switches within a domain to transparent mode is that VLAN configuration settings on other switches will not get overridden. A mistake that many administrators make is installing a new switch into a domain when it is configured as a VTP server. The default mode is VTP server, which can override the VLAN information and get propagated to other switches in the network. This can mean the deletion of all of the other VLANs within the switched network. In transparent mode, the switch will not participate in VTP, and it cannot override existing VTP settings.

Understanding the VTP Domain:

A VTP domain (also called a VLAN management domain) is made up of one or more interconnected network devices that share the same VTP domain name. A network device can be configured to be in one and only one VTP domain. You make global VLAN configuration changes for the domain using either the command-line interface (CLI) or Simple Network Management Protocol (SNMP).

By default, the Catalyst 6500 series switch is in VTP server mode and is in the nomanagement domain state until the switch receives an advertisement for a domain over a trunk link or you configure a management domain.

If the switch receives a VTP advertisement over a trunk link, it inherits the management domain name and the VTP configuration revision number. The switch ignores advertisements with a different management domain name or an earlier configuration revision number.

If you configure the switch as VTP transparent, you can create and modify VLANs but the changes affect only the individual switch.

When you make a change to the VLAN configuration on a VTP server, the change is propagated to all network devices in the VTP domain. VTP advertisements are transmitted out all trunk connections.

VTP maps VLANs dynamically across multiple LAN types with unique names and internal index associations. Mapping eliminates excessive device administration required

from network administrators.

Understanding VTP Modes:

You can configure a Catalyst 6500 series 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 network devices in the same VTP domain and synchronize their VLAN configuration with other network devices 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 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 68

A new Certkiller switch has been configured as a VTP client, and added to the existing VTP domain. Shortly after the ISL link is brought up to the rest of the network, the whole network goes down. What could have caused this to happen? (Choose the most likely option).

- A. The configuration revision of the switch inserted was higher than the configuration revision of the VTP domain.
- B. This is not an issue that could be related to the inserted switch since it was configured as a VTP client.
- C. The inserted switch was incorrectly configured for VTP v2 and caused an unstable condition.
- D. VLAN 1 was incorrectly deleted on the switch before insertion causing an unstable condition.

Answer: A

Explanation:

Even though the Catalyst switch is configured as a VTP client, and not a server, it can erase the information of an existing network. Cisco explains the problem as follows:

How a Recently Inserted Switch Can Cause Network Problems

This problem occurs when you have a large switched domain, which is all in the same VTP domain, and you want to add one switch in the network.

This switch was previously used in the lab, and a good VTP domain name was entered. It was configured as a VTP client, and connected to the rest of the network. Then, the ISL link was brought up to the rest of the network. In just a few seconds, the whole network is down. What could have happened?

The configuration revision of the switch you inserted was higher than the configuration

revision of the VTP domain. Therefore, your recently-introduced switch, with almost no configured VLANs, has erased all VLANs through the VTP domain.

This happens whether the switch is a VTP client or a VTP server. A VTP client can erase VLAN information on a VTP server. You can tell that this has happened when many of the ports in your network go into inactive state, but continue to be assigned to a nonexistent VLAN.

Solution:

Quickly reconfigure all of the VLANs on one of the VTP servers.

What to Remember:

Always make sure that the configuration revision of all switches inserted into the VTP domain is lower than the configuration revision of the switches already in the VTP domain.

Reference: http://www.cisco.com/warp/customer/473/21.html#vtp_ts_cav

QUESTION 69

The Certkiller network is bonding some of the Ethernet connections via PaGP in order to increase the backbone bandwidth. In PagP, what mode combination will allow a channel to be formed?

- A. Auto-auto
- B. Desirable-on
- C. On-auto
- D. Auto-desirable

Answer: D

Explanation:

The Port Aggregation Protocol (PAgP) modes are off, auto, desirable, and on. Only the combinations auto-desirable, desirable-desirable, and on-on will allow a channel to be formed.

The PAgP modes are explained below.

1. off: PAgP will not run. The channel is forced to remain down.
2. auto: PAgP is running passively. The formation of a channel is desired; however, it is not initiated.
3. desirable: PAgP is running actively. The formation of a channel is desired and initiated.
4. On: PAgP will not run. The channel is forced to come up.

Only the combinations of auto-desirable, desirable-desirable, and on-on will allow a channel to be formed. If a device on one side of the channel does not support PAgP, such as a router, the device on the other side must have PAgP set to on.

QUESTION 70

Certkiller is using extended VLANs (VLAN IDs 1006-4094) on their switches. What should he VTP mode be set to before configuring extended-range VLANs?

- A. Client

- B. Server
- C. Transparent
- D. Client or Server
- E. Client or Transparent
- F. Server or Transparent

Answer: 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 from other switches from their trunk interfaces. You can create, modify, and delete VLANs on a switch in VTP transparent mode. The switch must be in VTP transparent mode when you create extended-range VLANs.

QUESTION 71

Both ISL and 802.1Q is being used in the Certkiller network. When comparing the differences in ISL and 802.1Q, which of the following are true? (Select three)


- A. 802.1q allows the encapsulation of multiple trunks within a single trunk.
- B. 802.1q supports fewer VLANs than ISL.
- C. ISL is more efficient than 802.1q due to its smaller header size.
- D. ISL supports the processing of untagged frames.
- E. 802.1q uses a tag protocol ID of 0x8100

Answer: A, D, E

Both 802.1Q and ISL allows for the use of multiple trunks within any single trunk.

ISL supports the use of untagged frames on the trunk. All untagged frames are associated with the native VLAN, which is VLAN 1 by default.

The IEEE 802.1Q specification defines the Ethertype field to be 8100 in the presence of a VLAN ID. The entire packet format is shown below:



Incorrect Answers:

- B. 802.1Q supports up to 4096 VLANs, while ISL supports a maximum of 1024.
- C. ISL encapsulation adds 30 bytes to the entire frame, while the 802.1Q tag is only 4 bytes in length.

QUESTION 72

Your Catalyst switch is configured to support Multi Layer Switching (MLS). The switch contains an access list designed to prevent certain users from using ports 20 and 21 to reach the Internet. Because of this, which flow mask will be needed to

create each MLS shortcut?

- A. Destination flow mask
- B. Full flow mask
- C. Source flow mask
- D. Partial flow mask
- E. Destination-source mask
- F. Session flow mask

Answer: B

Explanation:

The three types of IP MLS modes are destination-ip, destination-source-ip, and full-flowip. Full flow-ip is in effect when an extended access list is applied..

To Block FTP traffic we require an extended access-list, which acts on layer 3 as well as layer 4 information in a packet. Because of this, the full flow mask is needed, which uses layer 3 and layer 4 information to create the shortcuts.

Incorrect Answers:

A. Destination-ip mode is the default mode. It is used when no access list is applied to the router's MLS-enabled interface.

C, D, F. These types of flow masks do not exist.

E. Source-destination-ip mode is in use when a standard access list is applied.

Reference:

<http://www.cisco.com/warp/public/473/13.html#flowchart>

QUESTION 73

While looking through the log files of your Catalyst switch, you notice that the following two messages are displayed somewhat infrequently:

```
%MLS-4-MOVEOVERFLOW:Too many moves, stop MLS for 5  
sec(20000000)
```

```
%MLS-4-RESUMESC:Resume MLS after detecting too many moves
```

What is the most likely cause of this problem?

- A. A transitory Spanning Tree loop
- B. A permanent Spanning Tree loop
- C. A faulty cable
- D. Faulty switch port
- E. A Pinnacle sync failure

Answer: A

Explanation:

If you see these messages infrequently, it is most likely a transitory L2 (spanning-tree) loop, resulting in packet flooding in one or more VLANs. However, if you are seeing an excessive number of these messages (for example, if your syslog server log file or your

switch console are being flooded with these messages), the problem might be due to the following reasons:

- a permanent L2 (spanning-tree) loop
- one or more faulty switch ports
- a bad cable (for example, a unidirectional fiber link)
- other bad hardware (not necessarily on the switch generating the messages)
- misconfigured device (for example, a traffic generator sending traffic to two switch ports using the same MAC address)

Incorrect Answers:

B, C, D. These are all possible causes, but not the most probably cause. The fact that only a few of these error messages are appearing tells us that A is the best choice.

E. This choice is the least likely to be the cause of the errors messages. A pinnacle Sync failure is a hardware error and Cisco does not cite this as a reason for the MLS errors at all.

Reference:

Common CatOS Error Messages on Cisco Catalyst Switches

<http://www.cisco.com/warp/public/473/34.shtml>

QUESTION 74

You have just recently implemented the Multilayer switching feature on your Catalyst Switch. How will this change affect your network?

- A. The MLS Switching Engine will forward the first packet in every flow.
- B. The MLS Route Processor will forward the first packet in every flow.
- C. The MLS Switching Engine will forward all traffic.
- D. The MLS Route Processor will forward all traffic.

Answer: B

Explanation:

Multi-Layer Switching (MLS) has become a highly desired method of accelerating routing performance through the use of dedicated Application Specific Integrated Circuits (ASICs). Traditional routing is done through a central CPU and software. MLS offloads a significant portion of routing (packet rewrite) to hardware, and thus has also been termed switching. MLS and Layer 3 switching are equivalent terms. It works by utilizing the MLS Route Processor, which forwards only the first packet in every sourcedestination flow. The remaining packets in the flow are then switched by the Switching Engine.

Incorrect Answers:

A, C. The Switching Engine is utilized after the first packet is processed by the Route processor. The packets in each flow are then routed once, and then switched.

D. MLS works by only running the first packet in any flow through the relatively resource intensive routing process.

QUESTION 75

A workstation has been connected to the Certkiller LAN using a Category 5e cable.

350-001

The workstation can connect to the rest of the network through the switch (i.e has full connectivity), but is suffering from much slower than expected performance. Looking at the interface statistics on the switch, many "runs" are being detected. Using software to read the counters on the workstation NIC, many FCS and alignment errors are occurring. What is the most likely cause of these errors?

- A. Bad Network Interface Card on the workstation
- B. Bad cable between the workstation and the switch
- C. The port has erroneously been configured as an 802.1q trunk port
- D. Mismatched speed settings between the workstation and the switch
- E. Mismatched duplex setting between the workstation and the switch
- F. None of the above.

Answer: E

Explanation:

In half-duplex environments, it is possible for both the switch and the connected device to sense the wire and transmit at exactly the same time and result in a collision. Collisions can cause runs, FCS, and alignment errors, caused when the frame is not completely copied to the wire, which results in fragmented frames.

When operating at full-duplex, FCS, cyclic redundancy checks (CRC), alignment errors, and runt counters should be minimal. If the link is operating at full-duplex, the collision counter is not active. If the FCS, CRC, alignment, or runt counters are incrementing, check for a duplex mismatch. Duplex mismatch is a situation in which the switch is operating at full-duplex and the connected device is operating at half-duplex, or the other way around. The result of a duplex mismatch is extremely slow performance, intermittent connectivity, and loss of connection.

The following describes the errors and their meanings:

Alignment Errors: Alignment errors are a count of the number of frames received that do not end with an even number of octets and have a bad CRC.

FCS Errors: FCS error count is the number of frames that were transmitted or received with a bad checksum (CRC value) in the Ethernet frame. These frames are dropped and not propagated onto other ports.

Runs: These are frames smaller than 64 bytes with a bad FCS value.

QUESTION 76

How much data can be carried in a standard Ethernet frame?

- A. Up to 4096 bytes
- B. No limit
- C. Up to 1500 bytes
- D. Up to 1518 bytes
- E. Up to 4400 bytes

Answer: C

Explanation:

A standard Ethernet frame MTU is 1500 bytes. The MTU size or packet size refers only to Ethernet payload. Ethernet frame size refers to the whole Ethernet frame, including the header and the trailer. This question asked for the amount of data that can be carried, which is the payload.

Note: Preamble is not calculated in frame size so DA (6 bytes) SA (6 bytes) Type (2 bytes) data + pad (1500 bytes) FCS (4bytes) = a total of 1518

Incorrect Answers:

D. A standard Ethernet frame MTU is 1500 bytes. This does not include the Ethernet header and Cyclic Redundancy Check (CRC) trailer, which is 18 bytes in length, to make the total Ethernet frame size of 1518. So the total size of an Ethernet packet can be as large as 1518 bytes, but the maximum payload is only 1500 bytes.

QUESTION 77

Which of the following will cause a switch port to go into the err-disable state?

(Choose all that apply)

- A. Duplex mismatch.
- B. Unidirectional Link Detection.
- C. AN incorrect VTP domain name is configured on the switch.
- D. Ethernet channeling is configured on the port.
- E. VLANs on the trunk were not matching on both sides.

Answer: A, B

Explanation:

If the interface status is err-disable in the output of the "show interface status" command, refer to the common reasons below. When a port is error-disabled, the LED associated with the port on the front panel will be solid orange.

The reasons for the interface going into "err-disable" state are varied. Some of the possibilities include the following:

- duplex mismatch (A is correct)
- port-channel misconfiguration
- Bridge Protocol Data Unit (BPDU) Guard violation
- UniDirectional Link Detection (UDLD) condition (B is correct)
- late-collision detection
- link-flap detection
- security violation
- Port Aggregation Protocol (PAgP) flap
- Layer 2 Tunneling Protocol (L2TP) Guard
- DHCP snooping rate-limiting
- EtherChannel guard detects a misconfigured EtherChannel

Incorrect Answers:

C. The VTP configuration relates to the switch as a whole and has no impact on individual ports.

D. Although the port could be in the err-disable state if the Ethernet channeling feature is

not set correctly on both ends, simply configuring channeling will not cause the port to go into this state by itself.

E. VLAN mismatches have no bearing on port status.

QUESTION 78


802.1Q trunking uses which Ethertype to identify itself?

- A. 8100
- B. 8021
- C. 802A
- D. 2020
- E. None of the above

Answer: A

Explanation:

The IEEE 802.1Q specification defines the Ethertype field to be 8100 in the presence of a VLAN ID. The entire packet format is shown below:



QUESTION 79

In an 802.3 LAN, PAUSE frames are used for inhibiting data transmissions for a period of time. Which MAC address does this PAUSE mechanism use in order to accomplish this?

- A. 00-00-00-00-00-00
- B. 00-00-0C-00-00-0F
- C. 01-04-0C-07-AC-3C
- D. 01-80-C2-00-00-01
- E. FF-FF-FF-FF-FF-FF

Answer: D

Explanation:

The globally assigned 48-bit multicast address 01-80-C2-00-00-01 has been reserved for use in MAC Control PAUSE frames for inhibiting transmission of data frames from a DTE in a full duplex mode IEEE 802.3 LAN. IEEE 802.1D-conformant bridges will not forward frames sent to this multicast destination address, regardless of the state of the bridge's ports, or whether or not the bridge implements the MAC Control sub-layer.

Reference: <http://www.techfest.com/networking/lan/ethernet3.htm>

QUESTION 80

A single end station failure can be prevented from disrupting the Spanning Tree

algorithm in a LAN according to the 802.1D specification. 802.1D recommends preventing this by:

- A. Clearing the Topology Change flag.
- B. Re-setting the Topology Change flag to one (1).
- C. Configuring the Bridge Forward Delay to less than 1/2 of the Bridge Maxage.
- D. Disabling the 801.1D Change Detection Parameter.
- E. Disabling the Topology Change Notifications.

Answer: D

Explanation:

The intent of the 802.1D standard is that the detectable failure of a MAC should cause the Bridge Port supported by that MAC to enter the Disabled state. A transition to the Disabled Port state causes the Bridge to initiate a topology change notification, unless, for the Port concerned, topology change detection has been explicitly disabled. Disabling this change detection will result in the prevention of the MAC failure to disrupt the Spanning Tree.

QUESTION 81

What trunking protocol uses an internal tagging mechanism that inserts a 4 byte tag field in the original Ethernet frame?

- A. ISL
- B. 802.1P
- C. DTP
- D. 802.1Q
- E. DVP

Answer: D

Explanation:

802.1Q is the IEEE standard for tagging frames on a trunk and supports up to 4096 VLANs. IEEE 802.1q uses an internal tagging mechanism which inserts a 4 byte tag field in the original Ethernet frame itself between the Source Address and Type/Length fields. Since the frame is altered, the trunking device re-computes the frame check sequence (FCS) on the modified frame.

Incorrect Answers:

- A. In ISL, a 26-byte header that contains a 10-bit VLAN ID is inserted at the beginning of the Ethernet frame.
- B, C, E. These are not trunk encapsulation options.

QUESTION 82

Which of the following are true regarding Unidirectional Link Detection? (Choose all that apply.)

- A. UDLD uses auto-negotiation to take care of physical signaling and fault detection.
- B. Both devices on the link need to support Unidirectional Link Detection.
- C. It works by exchanging protocol packets between the neighboring devices.
- D. It performs tasks that autonegotiation cannot perform.
- E. UDLD is a layer one protocol.

Answer: A, B, C, and D

Explanation:

In order to detect the unidirectional links before the forwarding loop is created, Cisco designed and implemented the UDLD protocol.

UDLD is a Layer 2 (L2) protocol that works with the Layer 1 (L1) mechanisms to determine the physical status of a link. At L1, auto-negotiation takes care of physical signaling and fault detection. UDLD performs tasks that auto-negotiation cannot perform, such as detecting the identities of neighbors and shutting down misconnected ports. When you enable both auto-negotiation and UDLD, L1 and L2 detections work together to prevent physical and logical unidirectional connections and the malfunctioning of other protocols.

UDLD works by exchanging protocol packets between the neighboring devices. In order for UDLD to work, both devices on the link must support UDLD and have it enabled on respective ports.

Each switch port configured for UDLD will send UDLD protocol packets containing the port's own device/port ID, and the neighbor's device/port IDs seen by UDLD on that port. Neighboring ports should see their own device/port ID (echo) in the packets received from the other side.

If the port does not see its own device/port ID in the incoming UDLD packets for a specific duration of time, the link is considered unidirectional.

Incorrect Answers:

E. UDLD is a Layer 2 (L2) protocol that works with the Layer 1 (L1) mechanisms to determine the physical status of a link.


Reference:

<http://www.cisco.com/warp/public/473/77.html>

QUESTION 83

The Certkiller network has 2 servers that are to be load balanced. They are connected to a Cisco switch via etherchannels, using 2 ports for each server as shown below:

350-001



With regard to this network, which of the following statements are true?

- A. Both channels should be given the same channel-id.
- B. Load balancing of traffic between two servers will not work.
- C. Spanning Tree needs to be disabled on the VLAN for the channel to come up.
- D. Channeling to a server is not supported.
- E. Channeling to the servers will work only for Fast Ethernet.
- F. Up to 4 links can be aggregated per channel

Answer: B

Explanation:

Traffic to each individual server will be load balanced over the Ethernet links in each channel, but traffic can not be load balanced between the servers. In order to do this, a load balancing device will need to be installed.

Incorrect Answers:

- A. In this instance there are 2 separate channels, so they will need to have different channel ID's. A single channel consists of Ethernet connections that terminate on the same device on each end.
- C. Spanning tree is supported over etherchannel, and should not be disabled.
- D. Channeling works between switches, routers, and servers.
- E. Channeling works over Fast Ethernet, Gigabit Ethernet, and 10 Gigabit Ethernet links.
- F. Up to 8 links can be aggregated per channel.

QUESTION 84

The Certkiller switched LAN network is upgrading many of the switch links to Gigabit Ethernet. Which of the following IEEE standards are used for Gigabit Ethernet? (Choose all that apply)

- A. 802.3z
- B. 802.3ab
- C. 802.3ad
- D. 802.3af
- E. All of the above

Answer: A, B

Explanation:

The Gigabit Ethernet standard is described in the IEEE 802.3z standard, which was defined in 1998. The 802.3ab document specifically describes the 1000BASE-T standard, which was done in 1999. Both describe Gigabit speed implementations, with 802.3z using fiber and 802.3ab using copper.

Incorrect Answers:

C. This standard describes Ethernet Link Aggregation.

D. The 802.3af standard describes a method for providing DTE power via MDI. This is useful for power over Ethernet implementations such as VOIP phones, providing for 15.4 Watts of power per port.

QUESTION 85

Which of the following statements regarding the use of SPAN on a Catalyst 6500 are true?

- A. With SPAN an entire VLAN can be configured to be the source.
- B. If the source port is configured as a trunk port, the traffic on the destination port will also be tagged, irrespective of the configuration on the destination port.
- C. In any active SPAN session, the destination port will not participate in Spanning Tree.
- D. It is possible to configure SPAN to have a Gigabit port as the destination port.
- E. In one SPAN session it is possible to monitor multiple ports that do not belong to the same VLAN.

Answer: A, C, D, E

Explanation:

A destination port (also called a monitor port) is a switch port where SPAN sends packets for analysis. If the trunking mode of a SPAN destination port is "on" or "nonegotiate" during SPAN session configuration, the SPAN packets forwarded by the destination port have the encapsulation as specified by the trunk type; however, the destination port stops trunking, and the show trunk command reflects the trunking status for the port prior to SPAN session configuration.

For a detailed discussion on SPAN and RSPAN refer the link below.

Reference:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_6_3/config_gd/span.htm

QUESTION 86

In order to maximize the speed and duplex setting resulting from auto-negotiation, the Certkiller network administrator has configured all Ethernet ports of a workgroup switch to 100 Mbps, full-duplex. When a workstation NIC configured for auto-negotiation is connected to the switch, the resulting negotiated parameters

are 100 Mbps, half-duplex.

What statement best accounts for this result?

- A. The workstation NIC must not be properly set for auto-negotiation as the highest port speed and duplex should result from this setup.
- B. The port speed is auto-negotiated by the burst of Fast link pulses sent upon port initialization, but duplex negotiation does not use FLPs.
- C. The switch should be configured for portfast since the spanning tree protocol leaves the port in the blocking state as it initializes, causing the auto-negotiation process to fail.
- D. Without auto-negotiation on the switch, FLPs will not be sent to the workstation, and as a result, the workstation will configure itself to half-duplex.
- E. There is problem with the NIC, most likely resulting from order drivers since auto-negotiation will allow the workstation NIC to learn what speed and duplex setting have been configured on the switch.

Answer: D

Explanation:

Speed determination issues may result in no connectivity. However, issues with autonegotiation of duplex generally do not result in link establishment issues. Instead, autonegotiation issues mainly result in performance-related issues. The most common problems when investigating NIC issues deal with speed and duplex configuration. The table below summarizes all possible settings of speed and duplex for FastEthernet NICs and switch ports.

The following table displays all of the various options:

350-001

Configuration NIC (Speed/Duplex)	Configuration Switch (Speed/Duplex)	Resulting NIC Speed/Duplex	Resulting Catalyst Speed/Duplex	Comments
AUTO	AUTO	1000 Mbps, Full-duplex	1000 Mbps, Full-duplex	Assuming maximum capability of Catalyst switch, and NIC is 1000 Mbps, full-duplex.
1000 Mbps, Full-duplex	AUTO	1000 Mbps, Full-duplex	1000 Mbps, Full-duplex	Link is established, but the switch does not see any autonegotiation information from NIC. Since Catalyst switches support only full-duplex operation with 1000 Mbps, they default to full-duplex, and this happens only when operating at 1000 Mbps.
1000 Mbps, Full-duplex	1000 Mbps, Full-duplex	1000 Mbps, Full-duplex	1000 Mbps, Full-duplex	Correct Manual Configuration
100 Mbps, Full-duplex	1000 Mbps, Full-duplex	No Link	No Link	Neither side establishes link, due to speed mismatch
100 Mbps, Full-duplex	AUTO	100 Mbps, Full-duplex	100 Mbps, Half-duplex	Duplex Mismatch ¹
AUTO	100 Mbps, Full-duplex	100 Mbps, Half-duplex	100 Mbps, Full-duplex	Duplex Mismatch ¹
100 Mbps,	100 Mbps,	100 Mbps,	100 Mbps,	Correct Manual

Full-duplex	Full-duplex	Full-duplex	Full-duplex	Configuration ²
100 Mbps, Half-duplex	AUTO	100 Mbps, Half-duplex	100 Mbps, Half-duplex	Link is established, but switch does not see any autonegotiation information from NIC and defaults to half-duplex when operating at 10/100 Mbps.
10 Mbps, Half-duplex	AUTO	10 Mbps, Half-duplex	10 Mbps, Half-duplex	Link is established, but switch does not see Fast Link Pulse (FLP) and defaults to 10 Mbps half-duplex.
10 Mbps, Half-duplex	100 Mbps, Half-duplex	No Link	No Link	Neither side establishes link, due to speed mismatch.
AUTO	100 Mbps, Half-duplex	100 Mbps, Half-duplex	100 Mbps, Half-duplex	Link is established, but NIC does not see any autonegotiation information and defaults to 100 Mbps, half-duplex.
AUTO	10 Mbps, Half-duplex	10 Mbps, Half-duplex	10 Mbps, Half-duplex	Link is established, but NIC does not see FLP and defaults to 10 Mbps, half-duplex.

1 A duplex mismatch may result in performance issues, intermittent connectivity, and loss of communication. When troubleshooting NIC issues, verify that the NIC and switch are using a valid configuration.

2 Some third-party NIC cards may fall back to half-duplex operation mode, even though both the switchport and NIC configuration have been manually configured for 100 Mbps, full-duplex. This behavior is due to the fact that NIC autonegotiation link detection is still operating when the NIC has been manually configured. This causes duplex inconsistency between the switchport and the NIC. Symptoms include poor port performance and frame check sequence (FCS) errors that increment on the switchport. To troubleshoot this issue, try manually configuring the switchport to 100 Mbps, half-duplex. If this action resolves the connectivity problems, you may be running into this NIC issue. Try updating to the latest drivers for your NIC, or contact your NIC card vendor for additional support

Note: Per the IEEE 802.3u specification, it is not possible to manually configure one link partner for 100 Mbps full-duplex and still auto-negotiate to full-duplex with the other link partner. Attempting to configure one link partner for 100 Mbps full-duplex and the other link partner for auto-negotiation will result in a duplex mismatch. This is a result of one link partner auto-negotiating and not seeing any auto-negotiation parameters from the other link partner and defaulting to half-duplex.

QUESTION 87

During routine maintenance, your issue the show interface Fast Ethernet 0

command on Router CK1 . The output from the command is shown in the following exhibit:

```
FastEthernet0 is up, line protocol is up
Hardware is DEC21140, address is 00e0.1ea8.e299 (bia
00e0.1ea8.e299)
Description: Ethernet 100Mbps
Internet address is 10.11.11.1/24
MTU 1500 bytes, BW 100000 Kbit, DLY 100 usec, rely 255/255,
load 3/255
Encapsulation ARPA, loopback not set, keepalive set (10
sec)
Half-duplex, 100Mb/s, 100BaseTX/FX
ARP type: ARPA, ARP Timeout 04:00:00
Last input 00:00:00, output 00:00:00, output hang never
Last clearing of "show interface" counters 6 weeks, 3 days
Queuing strategy: fifo
Output queue 0/40, 0 drops; input queue 0/75, 0 drops
5 minute input rate 1953000 bits/sec, 652 packets/sec
5 minute output rate 1407000 bits/sec, 600 packets/sec
47250970 packets input, 3285704002 bytes, 0 no buffer
Received 257038 broadcast, 1056 runts, 0 giants, 0
throttles
1918 input errors, 462 CRC, 0 frame, 0 overrun, 0
ignored, 0 abort
0 watchdog, 0 multicast
311 input packets with dribble condition detected
46457848 packets output, 3093573182 bytes, 0 underruns
0 output errors 759 collisions, 0 interface resets
0 babbles, 0 late collision, 0 deferred
0 lost carrier, 0 no carrier
0 output buffer failures, 0 output buffers swapped out
Based on the information above, should you be concerned with the operation of this
interface?
```

- A. Yes. There is a physical problem with the connection since there are recorded Runts and CRCs.
- B. No. The interface is normal for a 100mb full duplex environment.
- C. Yes. There are an excessive amount of collisions that could result from a cable that is too long.
- D. No. Collisions, runts and CRC's are normal for a 100mb half-duplex connection.
- E. None of the above.

Answer: D

Explanation:

Many performance issues with NICs may be related to data link errors. Excessive errors

usually indicate a problem. When operating at half-duplex setting, some data link errors such as FCS, alignment, runts, and collisions are normal. Generally, a one percent ratio of errors to total traffic is acceptable for half-duplex connections. If the ratio of errors to input packets is greater than two or three percent, performance degradation may be noticed.

Incorrect Answers:

A, C. The error counts are normal considering the number of packets that have passed through since the last clearing of counters.

B. The output above clearly says that this interface is operating in half duplex mode. However, if it were actually running in full duplex mode, then there would be reason for alarm as collisions are not possible on a full duplex link.

QUESTION 88

You are connecting a new 10/100 NIC to a Catalyst 5000 switch port. You want to achieve the most optimal settings possible. Which settings should you use?

A. NIC: 100 Mbps & Full-duplex

Catalyst: Auto

B. NIC: Auto

Catalyst: 100 Mbps & Full-duplex

C. NIC: 100 Mbps & Half-duplex

Catalyst: Auto

D. NIC: 100 Mbps & Half-duplex

Catalyst: 10 Mbps & Half-duplex

Answer: C

Explanation:

The speed and duplex cannot be Hard-Coded as full duplex on only one link. This will result a duplex mismatch. The default setting, when set to auto, is always (half-duplex) for port switch or NIC card. Note that setting the connections to auto on both devices is acceptable.

Incorrect Answers:

A. The Catalyst will default to half duplex, causing a mismatch.

B. The NIC will default to half duplex, causing a mismatch.

D. In this case both the Catalyst and the NIC have their information hard coded, but the throughput speeds do not match.

Reference:

<http://www.cisco.com/warp/public/473/46.html>

QUESTION 89

You are connecting a new PC with a 10/100 NIC to a Catalyst switch. The switch port is configured for auto negotiation for the speed and duplex settings. Which of the following settings on the PC will cause a duplex mismatch?

- A. 100mb, half duplex
- B. auto-negotiation speed, half duplex
- C. Auto-negotiation
- D. 100mb, full duplex
- E. 10mb, half duplex

Answer: D

Explanation:

Auto set on the switch side with 100mbps full-duplex will result in a duplex-mismatch because auto negotiation always defaults to half-duplex. Per the IEEE 802.3u specification, it is not possible to manually configure one link partner for 100 Mbps full duplex and still auto-negotiate to full-duplex with the other link partner. Attempting to configure one link partner for 100 Mbps full-duplex and the other link partner for autonegotiation will result in a duplex mismatch. This is a result of one link partner autonegotiating and not seeing any auto-negotiation parameters from the other link partner and defaulting to half-duplex.

Incorrect Answers:

A, B, C, E. With auto-negotiation set on the switch, any combination will be acceptable with the exception of full duplex. In the past, there were some issues with having each end set to "auto" but these issues have been resolved and is now a supported configuration from Cisco.

Reference:

<http://www.cisco.com/warp/public/473/46.html>

QUESTION 90

The Certkiller network includes a Full Duplex Gigabit link between a Router and a Switch. Periodically, you notice the collision counter incrementing slowly. What could be the cause of this problem?

- A. The Router is receiving too much traffic and is asserting the Collision signal to be slow down the rate that the switch is sending traffic.
- B. Both the Router and the Switch are transmitting at the same time.
- C. The switch and the router might be running an ISL trunk.
- D. A bug or faulty equipment.
- E. A few collisions are normal.

Answer: D

Explanation:

In full duplex mode collisions are impossible so it could only be a bug or problem with hardware. Full-duplex mode allows stations to transmit and receive data simultaneously. This makes for more efficient use of the available bandwidth by allowing open access to the medium. Conversely, this mode of operation can function only with Ethernet switching hubs or via Ethernet cross-over cables between interfaces capable of full duplex Ethernet. Full-duplex mode expects links to be point-to-point links. There are also

no collisions in full-duplex mode, so CSMA/CD is not needed.

Incorrect Answers:

- A. There is no such slow down mechanism as described here for a LAN.
- B. This would indeed be the cause of a collision in a half duplex LAN, but full duplex allows stations to listen and send at the same time.
- C. Trunking alone does affect the number of collisions on a segment.
- E. While a few collisions are indeed normal operation for a half duplex LAN, this does not apply for a full duplex segment.

QUESTION 91

You are a technician at Certkiller . You are connecting a new PC to a Catalyst 5000 switch port. After a short time, you notice some performance and intermittent connectivity issues with the PC. As a result of troubleshooting the issue, it is determined that the cause is a duplex mismatch between the PC and switch. Which combination below would cause this?

- A. NIC: 100 Mbps & Half-duplex
Catalyst: Auto
- B. NIC: 100 Mbps & Full-duplex
Catalyst: Auto
- C. NIC: Auto
Catalyst: 100 Mbps & Half-duplex
- D. NIC: Auto
Catalyst: Auto

Answer: B

Explanation:

The speed and duplex cannot be Hard-Coded on only one link. This will result a duplex mismatch. When set to auto, the duplex always defaults to half-duplex for both the switch port and for a NIC.

Reference:

<http://www.cisco.com/warp/public/473/46.html>

QUESTION 92

Which of the following is used in Ethernet networks? (Choose all that apply)

- A. Non Canonical format MAC addresses.
- B. CSMA/CD for media access.
- C. Canonical format MAC addresses.
- D. 802.5 encapsulated frames.
- E. 802.3 encapsulated frames

Answer: B, C, and E

Explanation:

B. Carrier Sense Multi Access with Collision Detection (CSMA/CD) is the media access method on Ethernet network.

C. Ethernet uses the Canonical MAC address format, which means the Least Significant Bit is transmitted first and the Most Significant Bit is transmitted last. The canonical transmission is also known as LSB first.

E. Ethernet is 802.3.

Incorrect Answers:

A. Ethernet and Token Ring topologies read MAC addresses differently. For example, a MAC address of 4040.4040.4040 on Ethernet is read as 0202.0202.0202 on Token Ring.

Token Rings use Non-canonical MAC address formats, also known as MSB first.

D. Token Ring uses 802.5

Reference:

http://www.cisco.com/en/US/tech/CK331/CK660/technologies_tech_note09186a008012811e.shtml

QUESTION 93

You are seeing a few errors on the LAN port of your Cisco router and suspect that the problem is with the link between the router and the switch. This link is configured for 100MB full duplex operation. In order to verify the problem, you connect a hub between the router and the switch so that you can connect your PC on this link and capture the packets. With your PC, you see a very large number of CRC errors, alignment errors, and late collisions. You are seeing the number of these errors increment quickly. What could be the cause of this?

- A. Either the Router or the Switch is faulty.
- B. These errors will not cause a performance problem.
- C. The cabling is causing these errors and should be replaced.
- D. Adding the Hub in between might have caused these errors.

Answer: D

Explanation:

The errors cited can all be attributed to increased cable distance, and the fact that the hub most likely does not support Full Duplex.

Incorrect Answers:

A. There were only a few errors on the port before the insertion of the hub into the network. If the router or switch were faulty, we would be seeing these errors at all times.

B. Although some errors, especially collisions, are normal in an Ethernet network, anything more than 2-3% of the packets having errors is excessive and indicates a network problem.

C. Similar to

A. The fact that the excessive errors were seen only after the hub was

placed in between the connection indicates that the errors were actually caused by the troubleshooting.

QUESTION 94

Troubleshooting STP convergence errors reveals that a switched network has multiple bridging loops, which is periodically causing problems. What Cisco IOS switching feature, if used improperly, would most likely cause these errors?

- A. Port Fast
- B. Uplink Fast
- C. Backbone Fast
- D. Dot1q Trunking
- E. Fast EtherChannel

Answer: A

Explanation:

Spanning tree PortFast causes a spanning tree port to enter the 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, rather than waiting for spanning tree to converge. PortFast should be used only when connecting a single end station to a switch port. Otherwise, you might create a network loop.

Incorrect Answers:

B. UplinkFast provides fast convergence after a spanning tree topology change and achieves load balancing between redundant links using uplink groups. An uplink group is a set of ports (per VLAN), only one of which is forwarding at any given time.

Specifically, an uplink group consists of the root port (which is forwarding) and a set of blocked ports, except for self-looping ports. The uplink group provides an alternate path in case the currently forwarding link fails.

C. BackboneFast is initiated when a root port or blocked port on a switch receives inferior BPDUs from its designated bridge. An inferior BPDU identifies one switch as both the root bridge and the designated bridge. When a switch receives an inferior BPDU, it indicates that a link to which the switch is not directly connected (an indirect link) has failed (that is, the designated bridge has lost its connection to the root bridge). Under normal spanning tree rules, the switch ignores inferior BPDUs for the configured maximum aging time.

The switch tries to determine if it has an alternate path to the root bridge. If the inferior BPDU arrives on a blocked port, the root port and other blocked ports on the switch become alternate paths to the root bridge. (Self-looped ports are not considered alternate paths to the root bridge.) If the inferior BPDU arrives on the root port, all blocked ports become alternate paths to the root bridge. If the inferior BPDU arrives on the root port and there are no blocked ports, the switch assumes that it has lost connectivity to the root bridge, causes the maximum aging time on the root to expire, and becomes the root switch according to normal spanning tree rules.

If the switch has alternate paths to the root bridge, it uses these alternate paths to transmit

a new kind of PDU called the Root Link Query PDU. The switch sends the Root Link Query PDU out all alternate paths to the root bridge. If the switch determines that it still has an alternate path to the root, it causes the maximum aging time on the ports on which it received the inferior BPDU to expire. If all the alternate paths to the root bridge indicate that the switch has lost connectivity to the root bridge, the switch causes the maximum aging times on the ports on which it received an inferior BPDU to expire. If one or more alternate paths can still connect to the root bridge, the switch makes all ports on which it received an inferior BPDU its designated ports and moves them out of the blocking state (if they were in blocking state), through the listening and learning states, and into the forwarding state.

D. The 802.1Q trunking method is the industry standard for trunk links, and can be used as an alternative to ISL. The use of either trunking method alone will not cause any bridging loops.

E. Fast Etherchannel simply provides a way to bond multiple Ethernet links into one larger channel. It will not introduce any STP loops into the network.

QUESTION 95

The speed and duplex settings are being configured for each port in a Catalyst switch. When trying to set the duplex mode on Port 1/1, what does the following message mean: "Port 1/1 is in auto-sensing mode"?

- A. Port 1/1 has auto-negotiated the duplex correctly.
- B. An error has occurred - the duplex setting of auto is not valid.
- C. CDP has detected that both sides are set for auto-negotiating.
- D. An error has occurred - the duplex is now mismatched.

Answer: B

Explanation:

When a port is in auto-sensing mode, both its speed and duplex are determined by autosensing. An error message is generated if you attempt to set the transmission type of autosensing ports. On a 10/100 module, if a port speed is set to auto, its transmission type (duplex) will also set to auto automatically, i.e., the duplex of an auto-speed port is not settable. The only two configurable choices for duplex settings are full and half.

QUESTION 96

The Certkiller network is experiencing network connectivity problems soon after an end-user disconnected her PC and connects a switch with an unknown configuration into an access layer switch port, which has spanning-tree portfast configured. What should be configured on the access layer switch to prevent the network connectivity problems?

- A. Certkiller 2950(config-if)# spanning-tree portfast bpduguard enable
- B. Certkiller 2950(config-if)# spanning-tree portfast bpduguard enable
- C. Certkiller 2950(config-if)# no spanning-tree portfast
- D. Certkiller 2950(config-if)# spanning-tree link-type point-to-point

- E. Certkiller 2950(config-if)# spanning-tree link-type shared
- F. Certkiller 2950(config-if)# no spanning-tree backbonefast
- G. Certkiller 2950(config-if)# no spanning-tree uplinkfast

Answer: B

Explanation:

The following explains the portfast Bridge Protocol Data Unit (BPDU) guard feature.

This feature is one of the Spanning-Tree Protocol (STP) enhancements created by Cisco to enhance switch network reliability, manageability, and security.

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 if the port is to be a part of the loop, 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 represent 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. This is done via the use of the "spanning-tree portfast bpduguard enable" command.

Reference:[http://www.cisco.com/en/US/tech/CK3 89/ CK6 21/technologies_tech_note09186a008009482f.shtml](http://www.cisco.com/en/US/tech/CK3_89/CK6_21/technologies_tech_note09186a008009482f.shtml)

QUESTION 97

A Certkiller LAN switch has been configured as shown below:

```
Switch(config)# wrr-queue bandwidth 10 20 70 1
Switch(config)# no wrr-queue cos-map
Switch(config)# wrr-queue cos-map 1 0 1
Switch(config)# wrr-queue cos-map 2 2 4
Switch(config)# wrr-queue cos-map 3 3 6 7
Switch(config)# wrr-queue cos-map 4 5
```

Certkiller.com

What does the IOS configuration displayed in the exhibit accomplish on a Catalyst

2900 switch?

- A. It enables frames with a CoS 0 or CoS 1 marking to be serviced by WRR (Weight Round Robin) queuing with a weighting value of 1.
- B. It enables frames with a CoS 5 marking to be serviced by the expedite queue.
- C. It guarantees 10% of the link bandwidth to Queue 1 and 20% to queue 2 and 70% to queue 3. Queue 4 is not used.
- D. It sets up the 3 CoS-to-DSCP mappings and DSCP-to-CoS mappings.
- E. It sets up the WRR queueing where frames with a CoS of 3 or 6 or 7 will have the highest priority.

Answer: E

Explanation:

WRR allows bandwidth sharing at the egress port. This command defines the bandwidths for egress WRR through scheduling weights. Four queues participate in the WRR unless you enable the egress expedite queue. The expedite queue is a strict-priority queue that is used until it is empty before using one of the WRR queues.

There is no order of dependencies for the `wrr-queue bandwidth` command. If you enable the egress priority, the weight ratio is calculated with the first three parameters; otherwise, all four parameters are used.

The WRR weights are used to partition the bandwidth between the queues in the event all queues are nonempty. For example, entering weights of 1:3 means that one queue gets 25 percent of the bandwidth and the other queue gets 75 percent as long as both queues have data.

Entering weights of 1:3 do not necessarily lead to the same results as entering weights at 10:30. Weights at 10:30 mean that more data is serviced from each queue and the latency of packets being serviced from the other queue goes up. You should set the weights so that at least one packet (maximum size) can be serviced from the lower priority queue at a time. For the higher priority queue, set the weights so that multiple packets are serviced at any one time.

To map CoS values to drop thresholds for a queue, use the `wrr-queue cos-map` command. Use the `no` form of this command to return to the default settings.

```
wrr-queue cos-map queue-id threshold-id cos-1 ... cos-n
```

```
no wrr-queue cos-map
```

Syntax Description

<i>queue-id</i>	Queue number; the valid value is 1.
<i>threshold-id</i>	Threshold ID; valid values are from 1 to 4.
<i>cos-1 ... cos-n</i>	CoS value; valid values are from 0 to 7.

Defaults

The defaults are as follows:

- Receive queue 1/drop threshold 1 and transmit queue 1/drop threshold 1: CoS 0 and 1.
- Receive queue 1/drop threshold 2 and transmit queue 1/drop threshold 2: CoS 2 and 3.

- Receive queue 2/drop threshold 3 and transmit queue 2/drop threshold 1: CoS 4 and 6.
 - Receive queue 2/drop threshold 4 and transmit queue 2/drop threshold 2: CoS 7.
-

QUESTION 98

You are implementing NAT (Network Address Translation) on the Certkiller network. Which of the following are features and functions of NAT? (Choose all that apply)

- A. Dynamic network address translation using a pool of IP addresses.
- B. Destination based address translation using either route maps or extended accesslists.
- C. NAT overloading for many to one address translations.
- D. Inside and outside source static network translation that allows overlapping network address spaces on the inside and the outside.
- E. NAT can be used with HSRP to provide for ISP redundancy.
- F. All of the above.

Answer: A, B, C, and D

Explanation:

A, B, C, D all describe various methods of implementing NAT.

Incorrect Answers:

E. With HSRP, the standby router would not have the NAT entries of the primary router, so when the fail-over occurs, connections will time out and fail.

Reference:

http://www.cisco.com/en/US/partner/tech/CK648/CK361/technologies_white_paper09186a0080091cb9.shtml

http://www.cisco.com/en/US/partner/tech/CK648/CK361/technologies_q_and_a_item09186a00800e523b.shtml

QUESTION 99

Which attributes should a station receive from a DHCP server?

- A. IP address, network mask, MAC address and DNS server
- B. IP address, DNS, default gateway and MAC address
- C. IP address, network mask, default gateway and host name
- D. IP address, network mask, default gateway and MAC address
- E. None of the above

Answer: E

Explanation:

DHCP servers can supply the following information to hosts on the LAN:

IP address

Subnet mask

Primary DNS server

Secondary DNS servers

Default gateway.

Incorrect Answers:

A, B, D. MAC addresses are burned in addresses that are obtained from the NIC hardware of a PC, not a DHCP server.

C. DHCP servers do not supply host names (such as those used in NetBIOS) or MAC addresses.

QUESTION 100

Which Cisco specific method should be configured on routers to support the need for a single default gateway for LAN hosts when there are two gateway routers providing connectivity to the network?

- A. DHCP
- B. RIP
- C. OSPF
- D. HSRP
- E. VRRP

Answer: D

Explanation:


Hot Standby Routing Protocol enables a virtual gateway on LAN networks, and enables the ability to provide a single default gateway for hosts to use, even though multiple routers are in use. This can provide for a level of load balancing, along with automatic failover capability for redundancy.

Additional Information on HSRP follows:

Hot Standby Routing Protocol (HSRP)

Hot Standby Routing Protocol (HSRP) is a Cisco proprietary protocol that brings routing functionality to end devices that would otherwise be incapable of taking advantage of redundant network connections. HSRP enables a pair of Cisco routers to work together to present the appearance of a single virtual default-gateway to end devices on a LAN segment. When you configure HSRP, the administrator assigns the virtual IP address whereas the Cisco IOS chooses a MAC address that falls within Cisco's MAC address block.

HSRP uses a priority scheme that enables routers within the same *standby group* to determine which is the *Active router* and which is the *Standby router*. The router with the highest priority is designated as the Active router; this would be the router that will forward all traffic. The mode of the router (Active/Standby) is communicated among routers within the same HSRP group through the HSRP Hello Protocol. A router that is a member of an HSRP group assumes it is in the Active mode until it hears an HSRP Hello that contains a priority that is higher than that configured on its interface. By default, the HSRP Hellos are sent out every 3 seconds and the hold timer is 10 seconds. If an HSRP router in Standby mode misses three consecutive HSRP Hellos, the router will assume that the Active router is finished and will transition into Active mode.



Router A Configuration

```
interface ethernet 0
ip address 192.168.10.2 255.255.255.0
standby 1 priority 200
standby 1 preempt
standby 1 ip 192.168.10.1
```

Router B Configuration

```
interface ethernet 0
ip address 192.168.10.3 255.255.255.0
standby 1 priority 100
standby 1 preempt
standby 1 ip 192.168.10.1
```

Figure 8.9

Simple HSRP example.

Incorrect Answers:

A. DHCP is the Dynamic Host Configuration Protocol, used to provide IP addressing, default gateway, and DNS information to LAN hosts.

B, C. These are routing protocols, and do not provide a means for allowing 2 or more routers to act as a single default gateway.

E. VRRP is the Virtual Router Redundancy Protocol, which is very similar to HSRP.

The major difference between the two is that HSRP is Cisco proprietary, while VRRP is an industry standard. This question asked for a Cisco specific solution.