



Exam : 350-001

Title : Cisco Certified Internetworking Expert

Ver : 07-22-08

QUESTION 1

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 2

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 3

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 4

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:

Refer to the following example:

Let's look at three routes that have just been installed in the routing table, and see how they look on the router.

```
router# show ip route
```

```
....
```

```
D 192.168.32.0/26 [90/25789217] via 10.1.1.1
```

```
R 192.168.32.0/24 [120/4] via 10.1.1.2
```

```
O 192.168.32.0/19 [110/229840] via 10.1.1.3
```

```
....
```

If a packet arrives on a router interface destined for 192.168.32.1, which route would the router choose? It depends on the prefix length, or the number of bits set in the subnet mask. Longer prefixes are always preferred over shorter ones when forwarding a packet.

In this case, a packet destined to 192.168.32.1 is directed toward 10.1.1.1, because

192.168.32.1 falls within the 192.168.32.0/26 network (192.168.32.0 to 192.168.32.63). It is chosen based on the longest match, not the fact that it has a lower AD. It also falls within the other two routes available, but the 192.168.32.0/26 has the longest prefix within the routing table (26 bits versus 24 or 19 bits).

Likewise, if a packet destined for 192.168.32.100 arrives on one of the router's interfaces, it's forwarded to 10.1.1.2, because 192.168.32.100 doesn't fall within 192.168.32.0/26 (192.168.32.0 through 192.168.32.63), but it does fall within the 192.168.32.0/24 destination (192.168.32.0 through 192.168.32.255). Again, it also falls into the range covered by 192.168.32.0/19, but 192.168.32.0/24 has a longer prefix length.

Incorrect Answers:

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

Reference:

www.cisco.com/en/US/tech/CK365/technologies_tech_note09186a0080094823.shtml

QUESTION 5

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.
- F. None of the above are true

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 version use hop counts as the metric.
- E. Both RIP versions are distance vector routing protocols.

QUESTION 6

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.
E. None of the above

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 7

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 be 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 8

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 9

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 10

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, E, F

Explanation:

Both RIP version 1 and RIP version 2 are distance vector protocols. EIGRP is an enhanced distance vector protocol, relying on the Diffused Update Algorithm (DUAL) to calculate the shortest path to a destination within a network

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.

QUESTION 11

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 12

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:

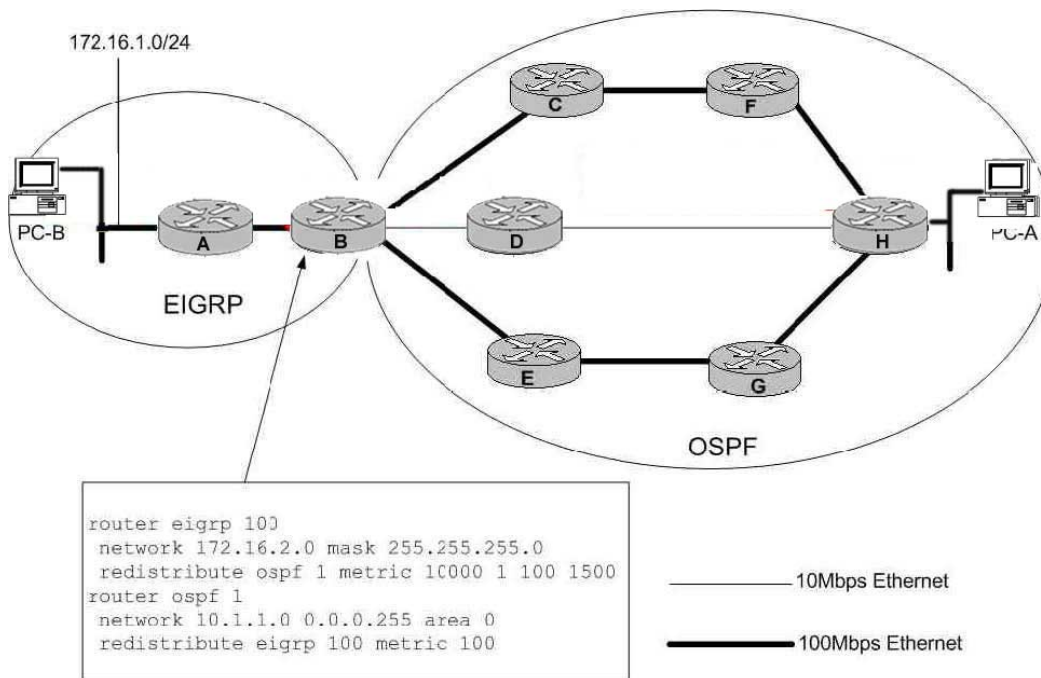
When routes are redistributed into OSPF, only routes that are not subnetted are redistributed if the subnets keyword is not specified. It is generally a good idea to include the "subnets" keyword at all times when redistributing routes from other protocols into OSPF.

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 13

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). This is true even though the redistributed routes are external type-2 routes. 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 14

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.

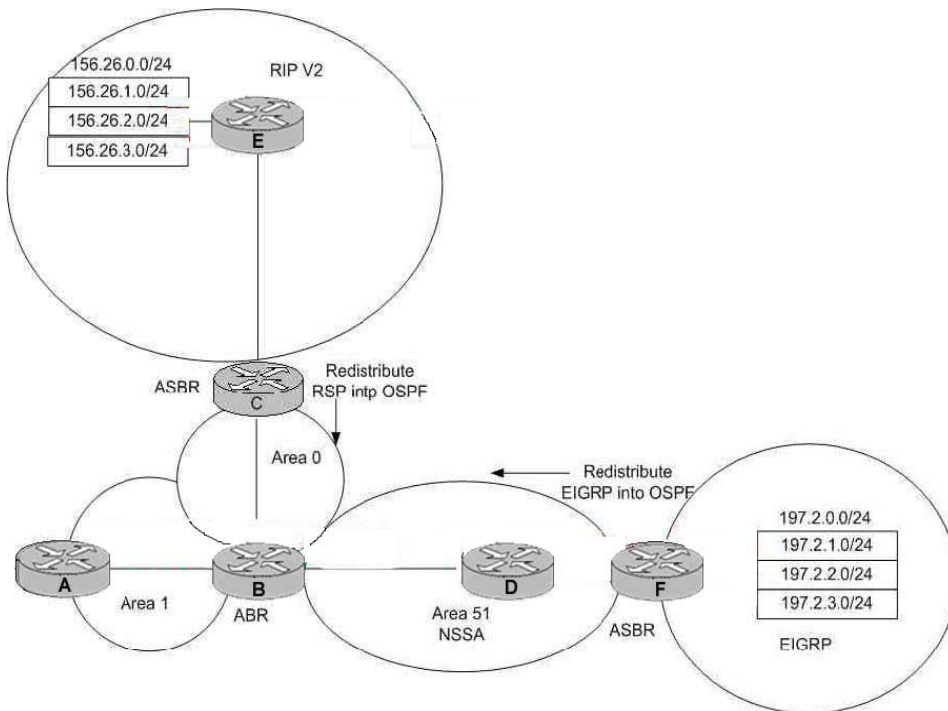
Answer: B

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 15

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 16

Four Certkiller routers are connected together as shown below:



Please study the exhibit above carefully. The Certkiller network is using a classful routing protocol. Subnet 10.1.1.0/24 is sourced by Router Certkiller 1 and Advertised to Router Certkiller 2. Router Certkiller 2 the sends updates to Router Certkiller 3, who forwards updates to Router Certkiller 4, which propagates routing information beyond. With regards to only the 10.1.1.0/24 subnet, what does Router Certkiller 4 advertise out for its 10.1.4.0/24 interface?

- A. 10.0.0.0/8
- B. 10.0.0.0
- C. No update is sent regarding the 10.1.1.0/24 subnet
- D. 10.1.1.0
- E. 10.1.1.0/24
- F. None of the above

Answer: C

Explanation:

Cisco routers running a classful routing protocol will automatically summarize at network boundaries. In this case, Certkiller 2 will advertise the summarized 10.0.0.0/8 network to Certkiller 3. Since Certkiller 3 will have a more specific route to the 10/0.0.0/8 network, he will not advertise this route to Certkiller 4, so Certkiller 4 will not receive either the 10.1.1.0/24 or the 10.0.0.0/8 route.

QUESTION 17

You want to improve the performance of the routers in the Certkiller network.

Which of the following could you use to minimize router resource requirements and improves manageability? (Select two)

- A. Resource Reservation Protocol
- B. CPU optimization
- C. Simple Network Management Protocol (SNMP)

- D. Auto-summarization
- E. Prefix Aggregation

Answer: D, E

Explanation:

Automatic route summarization and prefix aggregation is always a recommended best design practice whenever possible, as it means less routing table entries for the router to store. For example, many subnets can be hidden behind a single routing table entry, making these entries smaller, and routing more efficient).

QUESTION 18

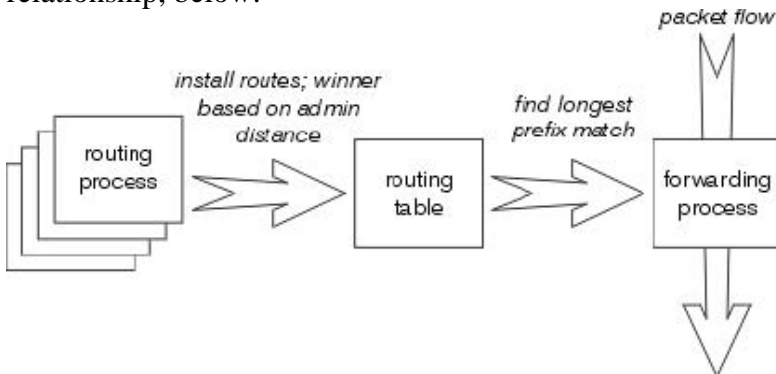
When a router within the Certkiller network comes to make a forwarding decision, the _____ always wins among the routes installed in the routing table.

- A. Administrative Distance
- B. Router ID
- C. Longest prefix match
- D. Routing process ID
- E. Peer ID
- F. None of the above

Answer: C

Explanation:

Making a forwarding decision actually consists of three sets of processes: the routing protocols, the routing table, and the actual process which makes a forwarding decision and switches packets. These three sets of processes are illustrated, along with their relationship, below:



The longest prefix match always wins among the routes actually installed in the routing table, while the routing protocol with the lowest administrative distance always wins when installing routes into the routing table.

Reference:

[www.cisco.com/en/US/tech/ CK3 65/technologies_tech_note09186a0080094823.shtml](http://www.cisco.com/en/US/tech/CK365/technologies_tech_note09186a0080094823.shtml)

QUESTION 19

Refer to the following exhibits:


```
! Certkiller4 Partial Running-Config
!
interface Serial2/0.101 point-to-point
 ip address 150.1.11.17 255.255.255.240
 ip summary-address rip 192.168.1.80 255.255.255.252
 frame-relay interface-dlci 101
!
router rip
 version 2
 network 150.1.0.0
 network 192.168.1.0
 no auto-summary
!
! Output Omitted
```

Certkiller4 #show ip route connected

```
192.168.1.0/32 is subnetted, 5 subnets
C    192.168.1.81 is directly connected, Loopback2
C    192.168.1.80 is directly connected, Loopback1
C    192.168.1.83 is directly connected, Loopback3
C    192.168.1.82 is directly connected, Loopback6
C    192.168.1.84 is directly connected, Loopback4
150.1.0.0/28 is subnetted, 2 subnets
C    150.1.11.16 is directly connected, Serial2/0.101
```

Based upon the partial Certkiller router configuration and the "show ip route connected" output shown in the exhibits above, which RIPv2 updates will be sent out of the Serial2/0.101 sub-interface from router Certkiller 4? (Select all that apply)

- A. 192.168.1.83/32
- B. 192.168.1.84/30
- C. 192.168.1.0/24
- D. 192.168.1.80/30
- E. 192.168.1.84/32
- F. 192.168.1.82/32
- G. 192.168.1.80/32
- H. 192.168.1.81/32
- I. None of the above

Answer: D, E

Explanation:

By default, RIP version 2 summarizes networks automatically. In the configuration example above, automatic summarization has been disabled. However, the "IP summary address" configuration statement takes precedence over automatic network summary, so the individual host loopback addresses will be summarized into one 192.168.1.80/30 network route. This will summarize the 192.168.1.80, 192.168.1.81, 192.168.1.82, and 192.168.1.83 networks into one route, leaving only the 192.168.1.84 network. This single host route will then also be advertised, since the automatic summarization feature was disabled.

QUESTION 20

What option is the best way to apply Classless Inter-Domain Routing (CIDR) if Certkiller, Inc wants to summarize the following addresses: 200.1.0.0/16, 200.2.0.0/16, 200.3.0.0/16, 200.5.0.0/16, 200.6.0.0/16, 200.7.0.0/16?

- A. 200.0.0.0/14, 200.4.0.0/15, 200.6.0.0/16, 200.7.0.0/16
- B. 200.0.0.0/16
- C. 200.4.0.0/14, 200.2.0.0/15, 200.2.0.0/16, 200.1.0.0/16
- D. 200.4.0.0/14, 200.2.0.0/15, 200.1.0.0/16
- E. 200.0.0.0/18
- F. None of the above

Answer: D

Explanation:

The Network 200.4.0.0/14 will encompass the 200.5.0.0, 200.6.0.0 and 200.7.0.0 networks. The second summarization, 200.2.0.0/15 will take care of both the 200.2.0.0 and 200.3.0.0 networks. Finally, the last network is needed in order to include the only remaining network, which is 200.1.0.0/16. This will summarize all 6 networks using only 3 statements.

Incorrect Answers:

- A. Although this answer will also fulfill the needs of summarizing all 6 networks, it is not the most efficient way as 4 network entries are needed here, instead of only 3 in answer choice D.
- B. This will mean that only the 200.0.0.0/16 network is advertised, which is not even one of the networks that need to be summarized.
- C. This is also not the most efficient choice, as the third statement (200.2.0.0/16) is redundant, since this network is already included in the 200.2.0.0/15 summarized route.
- E. This network mask would not include all of the needed networks.

QUESTION 21

Certkiller has a frame-relay network with 2 sites; a headquarters site and a remote site, each with a PVC connecting the 2 sites. The network is running RIPv2. Certkiller is now expanding and adding another remote site in the frame relay network and has ordered a second PVC between the new remote site and the headquarters site. All frame-relay interface IP addresses are in a single subnet. Certkiller has configured frame-relay DLCI mappings and can successfully ping from the new remote to the headquarters site as well as the other remote site. However, the new router does not have a route in its route table to the other remote site's LAN, and cannot ping the LAN interface or any hosts on that LAN. What is most likely causing the problem?

- A. Neighbor statements are not configured on the two remote sites, pointing to all other sites.

- B. The headquarters site router has split-horizon enabled on the frame-relay interface.
- C. The frame-relay IP to DLCI mappings are incorrectly configured.
- D. RIP cannot propagate routing updates over a partial mesh frame-relay configuration, so another routing protocol should be selected.
- E. Triggered updates should be configured on the headquarters router, to directly forward routing updates between the two remote sites.
- F. None of the above

Answer: B

Explanation:

RIP version 2 is a distance vector routing protocols, and by default all distance vector routing protocols utilize the split horizon rule to avoid routing loops. The split horizon rule blocks routing updates to be sent over the same interface that the route was learned from. In this case, the routes from the remote frame relay sites will not be sent to the other remote locations. In a hub and spoke topology such as this, the only way to ensure full connectivity between all locations using RIPv2 is to use sub-interfaces, or to disable the use of split horizons on the physical serial interface.

QUESTION 22

Layer 6 of the OSI model is responsible for which of the following functions?

- 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 23

Part of the configuration of router CK1 is shown below:

```
interface Ethernet0/0
ip address 10.1.1.1 255.255.255.0
ip summary-address rip 10.2.0.0 255.255.0.0
```

half-duplex

!!

router rip

network 10.0.0.0

You have configured RIPv2 summarization on the CK1 interface Ethernet 0/0 but the routes are still not being summarized. Based on the information provide above, what could be causing the problem?

- A. You need also to enable the auto summerazation under the RIP process.
- B. You need also to disable the auto summerazation under the RIP process.
- C. RIP does not support summarization on interface basis.
- D. Split horizon is enabled on the Ethernet 0/0 interface
- E. The mask configured on the "ip summary-address" command must be /24 bits.
- F. None of the above

Answer: D

Explanation:

Cisco routers can summarize routes in two ways:

1. Automatically, by summarizing subprefixes to the classful network boundary when crossing classful network boundaries (autosummary)
2. As specifically configured, advertising a summarized local IP address pool on the specified interface.

Autosummary will override the configured summary-address feature on a given interface except when both of the following conditions are true:

1. The configured interface summary-address and the IP address of the configured interface share the same major network (the classful, nonsubnetted portion of the IP address).
2. Split horizon is not enabled on the interface.

Note: If split horizon is enabled, neither an autosummary address nor the interface summary-address is advertised.

In the following example configuration, the major network is 10.0.0.0. The 10 in the address defines a Class A address space, allowing space for 0.x.x.x unique hosts where x defines unique bit positions in the addresses for these hosts. The summary of the major net defines the prefix as implied by the class (A, B, or C) of the address, without any network mask. The summary address 10.2.0.0 overrides the autosummary address of 10.0.0.0, 10.2.0.0 is advertised out interface E1, and 10.0.0.0 is not advertised:

int Ethernet 0/0

ip address 10.1.1.1 255.255.255.0

ip summary-address rip 10.2.0.0 255.255.0.0

no ip split-horizon

router rip

network 10.0.0.0

The above configuration is what should have been configured on router CK1 , by disabling split horizons.

Incorrect Answers:

A, B: By default, automatic summarization is already enabled. In this example, we need to disabled it. Automatic summarization is not the problem, however, since the manually configured summary address will override the automatically summarized address.

C: To configure IP summary addressing, use the "ip summary-address rip ip_address ip_network_mask" under the interface configuration:

E: The following subnet restrictions apply:

Supernet advertisement (advertising any network prefix less than its classful major network) is not allowed in RIP route summarization, other than advertising a supernet learned in the routing tables. Supernets learned on any interface that is subject to configuration are still learned. For example, the following summarization is invalid:
interface E1

```
..  
ip summary-address rip 10.0.0.0 252.0.0.0 (invalid supernet summarization)
```

Each route summarization on an interface must have a unique major net, even if the subnet mask is unique. For example, the following is not permitted:

```
int E1
```

```
...  
ip summary-address rip 10.1.0.0 255.255.0.0  
ip summary-address rip 10.2.0.0 255.255.0.0 (or different mask)
```


However, the subnet mask used does not need to be a /24.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1831/products_configuration_guide_chapter09186a00800d

QUESTION 24

The Certkiller multi-protocol network is displayed below:



Please study the exhibit carefully. Assume all necessary configurations are already correct for routing. Subnet 10.1.1.0/24 is sourced by Router Certkiller 1 and advertised via BGP, OSPF and EIGRP. Eventually, Router Certkiller 7 learns for this subnet. What is the routing protocol and administrative distance that Router Certkiller 7 used to reach subnet 10.1.1.0/24?

A. EIGRP, AD 5

- B. OSPF, AD 110
- C. EIGRP, AD 90
- D. EIGRP, AD 170
- E. BGP, AD 20
- F. BGP, AD 200
- G. None of the above


Answer: E

Explanation:

In this example, router Certkiller 7 will learn this route via OSPF, EIGRP, and IBGP, and external BGP. Since the administrative distance of EBGP is 20, this will be the preferred route to the 10.1.1.0/24 network.

QUESTION 25

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 26

Cisco Express Forwarding (CEF) has been enabled on all Certkiller routers. What table contains a mirror image of the forwarding information in the IP routing table used in CEF switching?

- A. Field information base
- B. Forwarding information base
- C. Filed information based
- D. Forwarding information based
- E. Forwarding instant based
- F. None of the above

Answer: B

Explanation:

Forwarding Information Base:

CEF uses a Forwarding Information Base (FIB) to make IP destination prefix-based switching decisions. The FIB is conceptually similar to a routing table or information base. It maintains a mirror image of the forwarding information contained in the IP routing table. When routing or topology changes occur in the network, the IP routing table is updated, and those changes are reflected in the FIB. The FIB maintains next-hop address information based on the information in the IP routing table.


Because there is a one-to-one correlation between FIB entries and routing table entries, the FIB contains all known routes and eliminates the need for route cache maintenance that is associated with switching paths such as fast switching and optimum switching.

Reference:

http://www.cisco.com/en/US/docs/ios/12_1/switch/configuration/guide/xcdef.html

QUESTION 27

Network Topology Exhibit:



Certkiller 1 configuration exhibit:

```
hostname Certkiller1
!
Router is
network 192.168.3.0
network 10.0.0.0
offset list 2 in 2 serial0
!
access-list 2 permit 10.1.0.0
```

Certkiller 3 configuration exhibit:

```
hostname Certkiller1

router rip
network 192.168.1.0
network 10.0.0.0
offset-list 1 in 2 serial0

access-list 1 permit 10.1.0.0
```

You work as a network engineer at Certkiller . Based on the information provided above, what statement is correct regarding the configurations in the figures?

- A. 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.
- B. The RIP metric cannot be changed and Load sharing will done between the two paths that exist from Certkiller 1 to 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.
- C. The RIP metric between Certkiller 1 and Certkiller 3 is 2 because the offset-list command is changing the metric to 2.
- D. The RIP metric between Certkiller 1 and Certkiller 3 remains 1.
- E. None of the above

Answer: A

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 28

You have issues the "clear ip route *" command on router CK1 . What would occur as a result of this command being issued?

- A. A router would recalculate its entire routing table and re-establish its neighbor relationships.

- B. Only its neighbor relationships would be re-established.
- C. Only link state routing protocols would be recalculated and only those neighbor relationships re-established.
- D. A router would recalculate its entire routing table but its neighbor relationships would not be affected.
- E. None of the above


Answer: 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 29

The Certkiller multi-protocol network is shown below:



Part of the router configurations for Certkiller 2 is also shown below:

```
Interface GigabitEthernet0/1
ip address 172.16.2.1 255.255.255.0
!
interface Serial0/0/0
ip address 172.16.3.1 255.255.255.0
!
router grg p 20
 redistribute rip
 network 172.16.0.0
 default-metric 10000 100 255 1 1500
 auto-summary
!
router ospf 30
 log-adjacency-changes
 network 172.16.0.0 0.0.255.255 area 0
!
router rip
 passive-interface GigabitEthernet0/1
 network 172.16.0.0

CertKiller2 #show ip eigrp topology 170.170.0.0 255.255.0.0
% IP-EIGRP (AS 20): Route not in topology table

CertKiller2 #show ip route 170.170.0.0 255.255.0.0
% Subnet

CertKiller2 #show ip route
<...>
 170.170.0.0/32 is subnetted, 1 subnets
 O   170.170.2.2 [110/2] via 172.16.2.2, 01:40:24,
GigabitEthernet0/1
 172.16.0.0/24 is subnetted, 3 subnets
 C   172.16.1.0 is directly connected, GigabitEthernet0/0
 C   172.16.2.0 is directly connected, GigabitEthernet0/1
 C   172.16.3.0 is directly connected, Serial0/0/0
```

You work as a network engineer at Certkiller .com. Please study the exhibits carefully. Router Certkiller 2 wants to redistribute all the routes in the RIP domain into the EIGRP domain but the network 170.170.0.0/16 is not being installed into Router Certkiller 2 routing table. What could fix the problem?

- A. Redistribute RIP into EIGRP 20 and filter route 170.170.0.0
- B. Redistribute RIP into OSPF 30
- C. Redistribute OSPF into EIGRP 20
- D. Filter route 170.170.0.0 from RIP
- E. Filter route 170.170.0.0. from OSPF
- F. None of the above

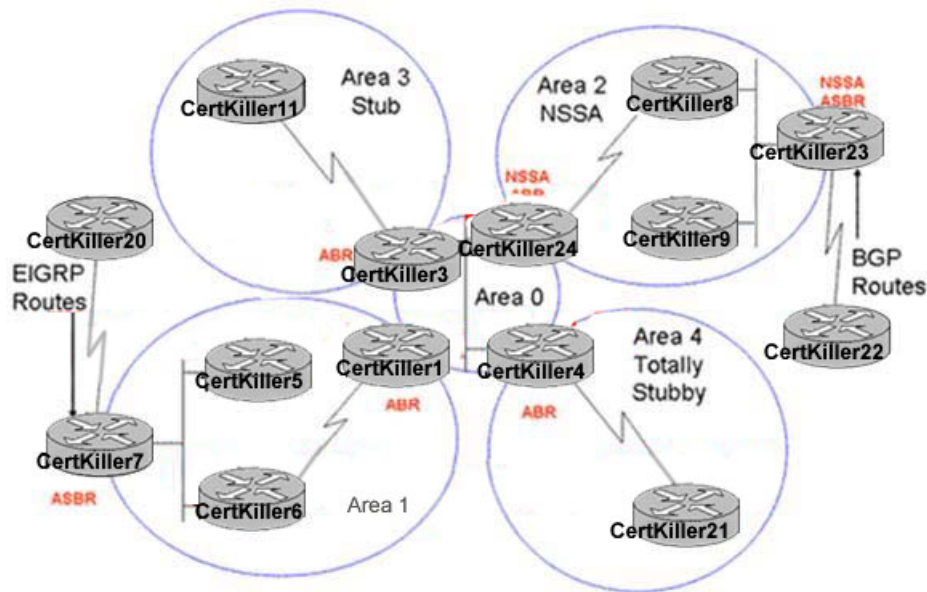
Answer: E

Explanation:

Since OSPF has a lower administrative distance than RIP, Certkiller 2 will install the 172.17.0.0/16 route learned via OSPF, and not RIP. Since router Certkiller 2 is only redistributing RIP routes into EIGRP, this route will not be installed included. Filtering this 172.17.0.0/16 route from OSPF will force the router to accept the RIP learned route and then redistribute it into EIGRP.

QUESTION 30

The Certkiller multi-protocol network topology is displayed below:



Please study the exhibit above carefully. Router Certkiller 23 (in area 2) is redistributing routes learned from BGP into the OSPF process. Certkiller 24 will generate what OSPF LSA types for advertisement to area 0? (Choose three)

- A. Type 4 - ASBR Summary LSAs
- B. Type 5 - AS External LSAs
- C. Type 1 - Router LSAs
- D. Type 3 - Network Summary LSAs
- E. Type 2 - Network LSAs

Answer: A, B, D

Explanation:

In this example, router Certkiller 24 is an area border router (ABR) as it connects areas 2 and 0. It also learned of routes from Certkiller 23, which is an ASBR since it redistributes BGP routes into OSPF so these routes will also need to be advertised into area 0.

The OSPF LSA Types are as follows:

Type 1: Router link advertisements generated by each router for each area it belongs to. Flooded to a single area only.

Type 2: Network link advertisements generated by designated routers describing the set of routers attached to a particular network. Flooded to the area that contains the network.

Type 3/4: Summary link advertisements generated by ABRs describing inter-area routes. Type 3 describes routes to networks and is used for summarization. Type 4 describes routes to the ASBR.

Type 5: Generated by the ASBR and describes links external to the Autonomous System (AS). Flooded to all areas except stub areas.

Type 6: Group membership link entry generated by multicast OSPF routers.


Type 7: NSSA external routes generated by ASBR. Only flooded to the NSS

A. The ABR

converts LSA type 7 into LSA type 5 before flooding them into the backbone (area 0).

QUESTION 31

The Certkiller network topology is shown in the following exhibit:



Network configuration exhibit:

Certkiller 1 configuration:

```
router eigrp 100
network 0.0.0.0 0.0.0.0
```

...

```
router ospf 100
network 0.0.0.0 0.0.0.0 area 1
redistribute eigrp 100 cost 10 subnets
```

Certkiller2 Configuration:

```
router eigrp 100
network 0.0.0.0 0.0.0.0
```

In the network illustrated Certkiller 3, Certkiller 4 are configured to run all connected links in OSPF area 1. The network administrator is complaining that traffic destined to 192.168.1.0/24 is being routed to Certkiller 2, even though Certkiller 2 is not running OSPF. Is this correct OSPF behavior or is something wrong with this network?

- A. The next hop towards 192.168.1.0/24 at Certkiller 4 should be 10.1.1.1, since Certkiller 1 is redistributing the route from EIGRP into OSPF. Certkiller 3 is forwarding traffic incorrectly
- B. Certkiller 4 would not have a route towards 192.168.1.0/24, so the network administrator is wrong in thinking any traffic is being forwarded there
- C. The next hop towards 192.168.1.0/24 at Certkiller 4 should be 10.1.1.2 which is

Certkiller 2

D. The next hop towards 192.168.1.0/24 at Certkiller 4 would be 10.1.2.2, which is Certkiller 3. Certkiller 3 should be load sharing between Certkiller 1 and Certkiller 2 for its next hop

E. None of the above


Answer: C

Explanation:

Since OSPF and EIGRP are being redistributed between Certkiller 1 and Certkiller 2, the route will appear to Certkiller 4 as an external route, with the next hop being the IP address at Certkiller 2.

QUESTION 32

The Certkiller network is depicted below:



Router Certkiller 6 is configured as shown below:

```
router rip
  version 2
  redistribute eigrp 10
  passive-interface default
  no passive-interface serial0/0.103
  network 10.0.0.0
  Default-metric
  no auto-summary
router eigrp 10
  redistribute rip
  passive-interface default
  no passive-interface fastethernet0/0
  no passive-interface fastethernet0/1
  network 10.0.0.0
  default-metric 10000 0 255 1 1500
  no auto-summary
!
access-list 1 deny 10.5.5.3
access-list 1 deny 10.5.5.3 0.0.0.3
access-list 1 deny 10.7.7.0 0.0.0.15
access-list 1 deny 10.50.50.0 0.0.0.7
access-list 1 permit any
```

You are required to configure redistribution of IGP protocols to ensure full IP visibility between all routers. As a safety precaution you must ensure that Certkiller 6 can not learn EIGRP routes it previously advertised into the RIP domain back from Certkiller 4.

What should you do in this scenario?

- A. Apply a distribute-list command to the FastEthernet and serial interfaces
- B. Apply a distribute-list command to the router rip area with the serial 0/0.103 interface only
- C. Apply a distribute-list command to the router EIGRP area with the serial interfaces
- D. Apply a route-map to the FastEthernet interfaces
- E. Apply a route-map and distribute-list command to complete the configuration

Answer: B

Explanation:

In order to prevent the EIGRP subnet routes from being advertised back to router Certkiller 6, we need to apply a distribute list command to the RIP routing process. The distribute list command should specify the routes that were configured in access-list 1. This will prevent the EIGRP subnets from being advertised back in via RIP. Since interface serial 0/0.103 is used as the connection to router Certkiller 4, the distribute list should be applied to this interface only. The other serial link to router Certkiller 1 does not need to be included, since this interface is specified as passive, by the "passive-interface default" configuration line.

Incorrect Answers:

A: Applying a distribute list to the fast Ethernet interfaces would result in lost connectivity between the EIGRP routers.

C: The distribute list needs to be applied to the RIP routing process, not the EIGRP process since you want to filter the incoming networks from the RIP network on the frame relay network.

D, E: It is not necessary to use route-maps for simply filtering network subnets.

QUESTION 33

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.

F. None of the above

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 34

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 35

While troubleshooting a LAN issue on the Certkiller network, you notice a number

of unicast frames. Having multiple unknown unicast frames in a switch would most likely deplete which of the following resources?

- A. MAC Addresses available in the system
- B. Power Consumption
- C. Memory available for frame buffering
- D. TCAM entries
- E. Available bandwidth
- F. None of the above

Answer: E

Explanation:

LAN switches use forwarding tables (Layer 2 (L2) tables, Content Addressable Memory (CAM) tables) to direct traffic to specific ports based on the VLAN number and the destination MAC address of the frame. When there is no entry corresponding to the frame's destination MAC address in the incoming VLAN, the (unicast) frame will be sent to all forwarding ports within the respective VLAN, which causes flooding.

Limited flooding is part of the normal switching process. There are situations, however, when continuous flooding can cause adverse performance effects on the network.

Large amounts of flooded traffic might saturate low-bandwidth links causing network performance issues or complete connectivity outage to devices connected across such low-bandwidth links.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps700/products_tech_note09186a00801d0808.shtml

QUESTION 36

You need to troubleshoot a Spanning Tree issue on one of your LAN segments, and you want to manipulate the path cost. Spanning tree protocol calculates path cost based on what?

- A. Interface Bandwidth
- B. Hop-Count
- C. Interface delay
- D. Interface Bandwidth and delay
- E. Bridge Priority
- F. None of the above

Answer: A

Explanation:

STP calculates the path cost based on the media speed (bandwidth) of the links between switches and the port cost of each port forwarding frame. Spanning tree selects the root port based on the path cost. The port with the lowest path cost to the root bridge becomes the root port. The root port is always in the forwarding state.

Reference:

http://www.cisco.com/en/US/tech/CK389/CK621/technologies_configuration_example09186a008009467c.shtml

QUESTION 37

Certkiller is utilizing Unidirectional Link Detection (UDLD) in their LAN. What are the characteristics and benefits of UDLD? (Choose three)

- A. UDLD prevents spanning tree loops caused by one way link conditions
- B. UDLD detects wiring mistakes when receive and transmit twisted pairs are not connected to the correct pinouts
- C. UDLD detects wiring mistakes when receive and transmit fibers are not connected to the same port on the remote side
- D. UDLD prevents spanning tree loops caused by link speed and duplex configuration mismatches
- E. UDLD protects against a situation where light is received on both sides of a fiber GE link (link up) but the fiber optic hardware is not communicating correctly

Answer: A, C, E

Explanation:

The UDLD protocol allows devices connected through fiber-optic or copper Ethernet cables (for example, Category5 cabling) to monitor the physical configuration of the cables and detect when a unidirectional link exists. When a unidirectional link is detected, UDLD shuts down the affected port and alerts the user. Unidirectional links can cause a variety of problems, including spanning-tree topology loops.

UDLD is a Layer2 protocol that works with Layer 1 mechanisms such as autonegotiation to determine the physical status of a link. At Layer1, autonegotiation handles physical signaling and fault detection. UDLD also performs tasks that autonegotiation cannot perform such as detecting the identities of neighbors and shutting down misconnected ports. When both autonegotiation and UDLD are enabled, Layer1 and Layer2 detection features can work together to prevent physical and logical unidirectional connections and malfunctioning of other protocols.

A unidirectional link occurs whenever traffic transmitted by the local device over a link is received by the neighbor but traffic transmitted from the neighbor is not received by the local device. For example, if one of the fiber strands in a pair is disconnected, as long as autonegotiation is active the link does not stay up. In this situation, the logical link is undetermined, and UDLD does not take any actions. If both fibers are working normally at Layer1, then UDLD at Layer 2 determines whether those fibers are connected correctly and whether traffic is flowing bidirectionally between the correct neighbors. This check cannot be performed by autonegotiation, because autonegotiation is a Layer1 feature.

The switch periodically transmits UDLD messages (packets) to neighbor devices on ports with UDLD enabled. If the messages are echoed back to the sender within a specific time frame and they are lacking a specific acknowledgment (echo), the link is flagged as unidirectional and the port is shut down. Devices on both ends of the link must support UDLD in order for the protocol to successfully identify and disable unidirectional links.

Reference:

<http://www.cisco.com/en/US/docs/switches/lan/catalyst5000/catos/5.x/configuration/guide/udld.html#wp101984>

QUESTION 38


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

- A. BPDUs 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.



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

The TCN is a very simple BPDU that contains absolutely no information that a bridge sends out every hello_time seconds (this is locally configured hello_time, not the

hello_time specified in configuration BPDUs). The designated bridge acknowledges the TCN by immediately sending back a normal configuration BPDU with the topology change acknowledgement (TCA) bit set. The bridge 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:

1. 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.
2. 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/CK6>

[21/technologies_tech_note09186a0080094797.shtml#portfastcommand](http://www.cisco.com/en/US/tech/CK389/CK6/technologies_tech_note09186a0080094797.shtml#portfastcommand)

QUESTION 39

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	

1. Protocol ID - indicates that the packet is a BPDU.
2. Version - the version of the BPDU being used.
3. Message Type - the stage of the negotiation.
4. Flags - two bits are used to indicate a change in topology and to indicate acknowledgement of the TCN BPDU.
5. Root ID - the root bridge priority (2 bytes) followed by the MAC address (6 bytes).
6. Root Path Cost - the total cost to from this particular bridge to the designated root bridge.
7. Bridge ID - the bridge priority (2 bytes) followed by the MAC address (6 bytes), lowest value wins! The default bridge priority is 0x8000 (3276810).
8. 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.

9. Message Age - timers for aging messages (only has effect on the network if the root bridge is configured with this parameter).

10. 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.

11. 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.

12. 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 40

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 41

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 42

The Certkiller network is shown in the following exhibit:



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

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.

D: After the "set span root 1" command is set, CK1 will send a PDU to CK2 and a new STP election will occur. During this time, the port will transition into the blocking, listening, and learning states and during this time the forwarding of traffic will be temporarily halted.

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

QUESTION 43

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

!

bridge1 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  
!
```

```
bridge2 protocol ieee  
bridge2 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 44

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

A. STP Timer

B. Hold Timer

C. Hello Timer

D. Max Age Timer

E. Forward Delay Timer

Answer: C

Explanation:

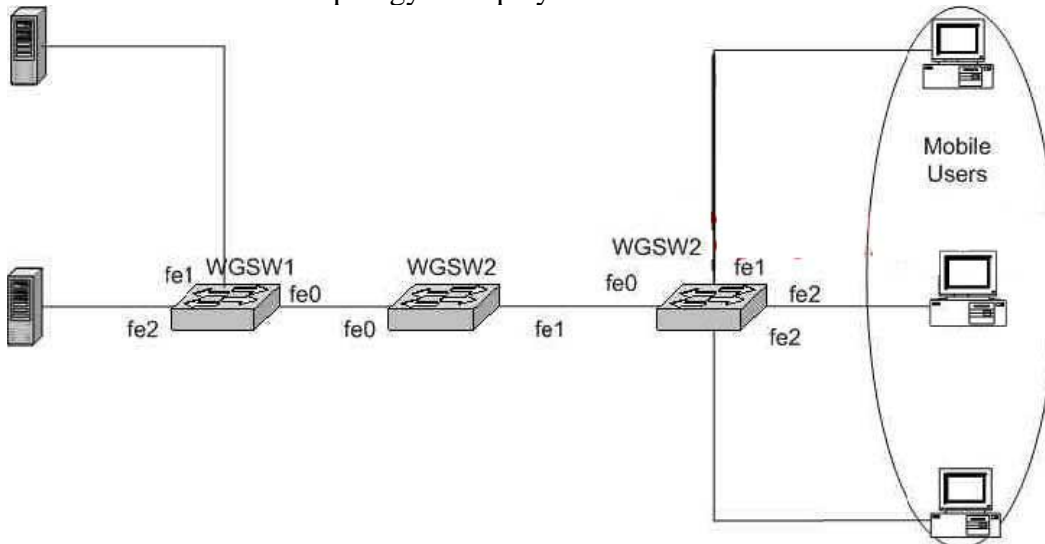
The STP Hello Time is the time between each Bridge Protocol Data Unit (BPDU) that is sent on a port. This is equal to two seconds by default, but can be tuned to be between one and ten seconds.

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 45

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 46

The following was executed on switch Certkiller C:

```
CertkillerC # sh spanning-tree vlan 2 detail
VLAN0002 is executing the IEEE compatible Spanning Tree protocol
Bridge Identifier has priority 32768, group 2, address 0009.7c00.2280
Configured hello time 2, max age 20, forward delay 15
Current root has priority 32770, address 0007.ab49.7100
Root port is 3 (FastEthernet0/3), cost of root path is 38
Topology change flag not set, detected flag not set
Number of topology changes 7 last change occurred 3w5d ago
from FastEthernet0/2
Times: hold 1, topology change 35, notification 2
hello 2, max age 20, forward delay 15
Timers: hello 0, topology change 0, notification 0, aging 300
...
```



If all switches in the Certkiller network run the same type of spanning tree, what is the total number of spanning tree topology changes that occurred in this network?

- A. 7
- B. 35
- C. Can not be determined. Only the root bridge tracks the complete amount of topology changes
- D. 0
- E. 2

Answer: A


Explanation:

The role of the TC mechanism is to correct L2 forwarding tables after the forwarding

topology has changed. This is necessary to avoid a connectivity outage because, after a TC, some MAC addresses previously accessible through particular ports might become accessible through different ports. TC shortens the forwarding table aging time on all switches in the VLAN where the TC occurs; thus, if the address is not relearned, it will age-out and flooding will occur to ensure packets reach the destination MAC address. TC is triggered by the change of a port's STP state to or from the STP forwarding state. Note: With Rapid STP or Multiple STP (IEEE 802.1w and IEEE 802.1s), TC is triggered by a change of the port's state to forwarding, as well as the role change from designated to root. With Rapid STP, the L2 forwarding table is immediately flushed, as opposed to 802.1d, which shortens the aging time. The immediate flushing of the forwarding table restores connectivity faster, but will cause more flooding. In the output shown above, the total number of topology changes for VLAN 2 is shown to be 7, with the last change occurring 3 weeks and 5 days ago.

QUESTION 47

The Certkiller switched LAN is shown below:



Due to hardware failure on the link between switches Certkiller A and Certkiller B, Spanning Tree BPDUs from switch Certkiller A are no longer received by switch Certkiller B, but the link remains up (see the drawing) Provided LoopGuard feature is configured on all ports, which port will be put into 'Loop-inconsistent' state?

- A. Port on switch Certkiller C connecting to switch Certkiller B
- B. Port on switch Certkiller B connecting to switch Certkiller C
- C. LoopGuard would not detect any issue in this scenario
- D. Port on switch Certkiller A connecting to switch Certkiller B and port on switch Certkiller B connecting to switch Certkiller A
- E. Port on switch Certkiller B connecting to switch Certkiller C

Answer: C


Understanding How Loop Guard Works:

Loop guard helps prevent bridging loops that could occur because of a uni-directional link failure on a point-to-point link. When enabled globally, the loop guard applies to all point-to-point ports on the system. Loop guard detects root ports and blocked ports and ensures that they keep receiving BPDUs from their designated port on the segment. If a

loop guard enabled root or blocked port stop a receiving BPDUs from its designated port, it transitions to the loop-inconsistent blocking state, assuming there is a physical link error on this port. The port recovers from this loop-inconsistent state as soon as it receives a BPDU.

You can enable loop guard on a per-port basis. When you enable loop guard, it is automatically applied to all of the active instances or VLANs to which that port belongs. When you disable loop guard, it is disabled for the specified ports. Disabling loop guard moves all loop-inconsistent ports to the listening state.

Example:



The figure above illustrates the following configuration:

Switches A and B are distribution switches.

Switch C is an access switch.

Loop guard is enabled on ports 3/1 and 3/2 on Switches A, B, and C.

Enabling loop guard on a root switch has no effect but provides protection when a root switch becomes a nonroot switch.

These caveats apply to loop guard:

Spanning tree always chooses the first operational port in the channel to send the BPDUs.

If that link becomes unidirectional, loop guard blocks the channel, even if other links in the channel are functioning properly.

If a set of ports that are already blocked by loop guard are grouped together to form a channel, spanning tree loses all the state information for those ports and the new channel port may obtain the forwarding state with a designated role.

Note

You can enable UniDirectional Link Detection (UDLD) to help isolate the link failure. A loop may occur until UDLD detects the failure, but loop guard will not be able to detect it.

Loop guard has no effect on a disabled spanning tree instance or a VLAN.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a0080160

QUESTION 48

On a Certkiller bridge running the rapid spanning tree protocol, which port will send a BPDU with the proposal flag?

- A. Designated port in forwarding state
- B. Designated port in non-forwarding state or the Root port in forwarding state
- C. Root port in blocking state
- D. Alternate port
- E. None of the above

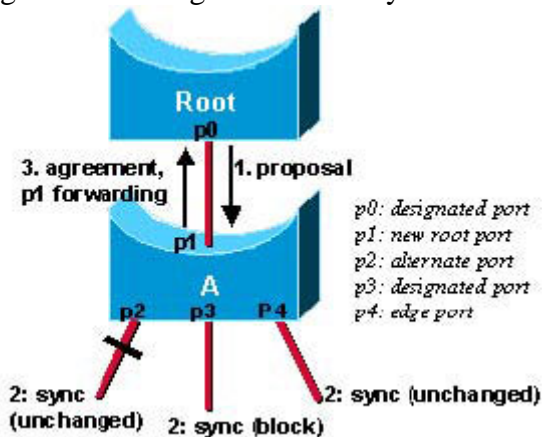
Answer: B

Explanation:

Proposal/Agreement Sequence:

When a port has been selected by the STA to become a designated port, 802.1d still waits twice <forward delay> seconds (2x15 by default) before transitioning it to the forwarding state. In RSTP, this condition corresponds to a port with a designated role but a blocking state. The diagrams below illustrate how fast transition is achieved step-by-step. Suppose a new link is created between the root and Switch

A. Both ports on this link are put in a designated blocking state until they receive a BPDU from their counterpart.




When a designated port is in a discarding or learning state (and only in this case), it sets the proposal bit on the BPDUs it sends out. This is what happens for port p0 of the root bridge, as shown in Step 1 of the diagram above. Because Switch A receives superior information, it immediately knows that p1 is going to be its new root port. Switch A then starts a sync to ensure that all of its ports are in-sync with this new information. A port is in-sync if it meets either of the following criteria:

1. The port is in blocking state (which means discarding, in a stable topology).
2. The port is an edge port.

In order to illustrate the effect of the sync mechanism on different kind of ports, suppose there exists an alternate port p2, a designated forwarding port p3, and an edge port p4 on Switch


A. Notice that p2 and p4 already meet one of the criteria listed above. In order to be in sync (Step 2 of the diagram above), Switch A just needs to block port p3, assigning it the discarding state. Now that all of its ports are in sync, Switch A can now unblock its newly selected root port p1 and reply to the root by sending an agreement message (Step 3). This message is a copy of the proposal BPDU, with the agreement bit set instead of the proposal bit. This ensures that port p0 knows exactly to which proposal the agreement

it receives corresponds to.



Once p0 receives that agreement, it can immediately transition to forwarding. This is Step 4 of the figure above. Notice that port p3 was left in a designated discarding state after the sync. In Step 4, that port is in the exact same situation as was port p0 during Step 1. It then starts proposing to its neighbor, attempting to quickly transition to forwarding.

1. The proposal agreement mechanism is very fast, as it does not rely on any timers. This wave of handshakes propagates quickly towards the edge of the network, and quickly restores connectivity after a change in the topology.
2. If a designated discarding port does not receive an agreement after having sent a proposal, it slowly transitions to the forwarding state, falling back to the traditional 802.1d listening-learning sequence. This could happen for instance if the remote bridge doesn't understand RSTP BPDUs, or if the remote bridge's port is blocking.
3. Cisco introduced an enhancement to the sync mechanism that allows a bridge to put only its former root port in the discarding state when syncing. Detailing the way this mechanism works is beyond the scope of this document. However, one can safely assume that it will be invoked in most common convergence cases. The scenario described in the Convergence with 802.1w section of this document now becomes extremely efficient, as only the ports on the path to the final blocked port are temporarily confused.



Reference: <http://www.cisco.com/warp/public/473/146.html#agree>

QUESTION 49

In RSTP (Rapid Spanning Tree Protocol) what is a port that provides an alternate

path to the leaves of the Spanning Tree and what state is this port in when it is not in the active topology?

- A. Root port and listening
- B. Alternate port and forwarding
- C. Alternate port and learning
- D. Designated port and learning
- E. Backup port and discarding
- F. None of the above

Answer: E

Explanation:

The RSTP provides rapid convergence of the spanning tree by assigning port roles and by determining the active topology. The RSTP builds upon the IEEE 802.1D STP to select the switch with the highest switch priority (lowest numerical priority value) as the root switch. Then the RSTP assigns one of these port roles to individual ports:

1. Root port-provides the best path (lowest cost) when the switch forwards packets to the root switch.

2.

Designated port-connects to the designated switch, which incurs the lowest path cost when forwarding packets from that LAN to the root switch. The port through which the designated switch is attached to the LAN is called the designated port.

3. Alternate port-offers an alternate path toward the root switch to that provided by the current root port.

4. Backup port-acts as a backup for the path provided by a designated port toward the leaves of the spanning tree. A backup port can exist only when two ports are connected together in a loopback by a point-to-point link or when a switch has two or more connections to a shared LAN segment.

5. Disabled port-has no role within the operation of the spanning tree.

A port with the root or a designated port role is included in the active topology. A port with the alternate or backup port role is excluded from the active topology.


In a stable topology with consistent port roles throughout the network, the RSTP ensures that every root port and designated port immediately transition to the forwarding state while all alternate and backup ports are always in the discarding state (equivalent to blocking in 802.1D).

Reference:

http://www.cisco.com/en/US/partner/products/hw/switches/ps628/products_configuration_guide_chapter09186a

QUESTION 50

Exhibit:



In the diagram shown above, Switch Certkiller 1 is the Root of Spanning Tree. If there is a Unidirectional link failure between switches Certkiller 1 and Certkiller 3, and Switch Certkiller 3 stops receiving BPDUs from Switch Certkiller 1, it will transition its blocked port to the forwarding state and we can have a Spanning Tree loop.

What feature can we use to prevent this from happening? (Select all that apply)

- A. Portfast
- B. Portfast BPDU filter
- C. Portfast BPDU guard
- D. UDLD
- E. Loopguard
- F. None of the above

Answer: D, E


Explanation:

Loop guard helps prevent bridging loops that could occur because of a uni-directional link failure on a point-to-point link. When enabled globally, the loop guard applies to all point-to-point ports on the system. Loop guard detects root ports and blocked ports and ensures that they keep receiving BPDUs from their designated port on the segment. If a loop guard enabled root or blocked port stop a receiving BPDUs from its designated port, it transitions to the loop-inconsistent blocking state, assuming there is a physical link error on this port. The port recovers from this loop-inconsistent state as soon as it receives a BPDU.

You can enable loop guard on a per-port basis. When you enable loop guard, it is automatically applied to all of the active instances or VLANs to which that port belongs. When you disable loop guard, it is disabled for the specified ports. Disabling loop guard moves all loop-inconsistent ports to the listening state.

If you enable loop guard on a channel and the first link becomes unidirectional, loop guard blocks the entire channel until the affected port is removed from the channel.

Example:



In this example:

Switches A and B are distribution switches.

Switch C is an access switch.

Loop guard is enabled on ports 3/1 and 3/2 on Switches A, B, and C.

Enabling loop guard on a root switch has no effect but provides protection when a root switch becomes a nonroot switch.

Note:

You can enable UniDirectional Link Detection (UDLD) to help isolate the link failure. A loop may occur until UDLD detects the failure, but loop guard will not be able to detect it.

Loop guard has no effect on a disabled spanning tree instance or a VLAN.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a0080160

QUESTION 51

Troubleshooting STP convergence errors within the Certkiller LAN reveals that it has multiple bridging loops. What Cisco IOS switching feature, when improperly implemented, is most likely the cause of these errors?

- A. Port Fast
- B. Uplink Fast
- C. Backbone Fast
- D. Dot1q Trunking
- E. Fast EtherChannel
- F. None of the above

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 52

Switch CK1 is running the rapid spanning tree protocol (RSTP). Upon a topology change, what happens to dynamic entries in the L2 forwarding table?

- A. Only entries behind port where TC was received are removed
- B. All entries are removed except for those behind edge ports and the port where TC was received
- C. All entries are removed (purged)
- D. All entries are removed except for entries behind edge ports
- E. Aging timer is set to 15 seconds, so idle entries age out
- F. None of the above

Answer: B

Explanation:

Topology Change Detection

In RSTP, only non-edge ports moving to the forwarding state cause a topology change. This means that a loss of connectivity is not considered as a topology change any more, contrarily to 802.1d (that is, a port that moves to blocking no longer generates a TC).

When a RSTP bridge detects a topology change, the following happens:

1. It starts the TC While timer with a value equal to twice the hello time for all its non-edge designated ports and its root port if necessary.
2. It flushes the MAC addresses associated with all these ports.

Topology Change Propagation

When a bridge receives a BPDU with the TC bit set from a neighbor, the following happens:

1. It clears the MAC addresses learnt on all its ports except the one that received the topology change.
2. It starts the TC While timer and sends BPDUs with TC set on all its designated ports and root port (RSTP no longer uses the specific TCN BPDU, unless a legacy bridge needs to be notified).

Reference: <http://www.cisco.com/warp/public/473/146.html#agree>

QUESTION 53

On a Certkiller bridge running the Rapid Spanning-tree protocol (RSTP), BPDU information on the port will be aged when?

- A. After MaxAge time
- B. 15 seconds
- C. RSTP does not age out BPDU information on ports
- D. After BPDU Age will reach MaxAge or after 3 hello times -which ever occurs first
- E. After 6 seconds
- F. None of the above

Answer: D

Explanation:

On a given port, if hellos are not received for three consecutive times, protocol information can be immediately aged out (or if max_age expires). Because of the

previously mentioned protocol modification, BPDUs are now used as a keep-alive mechanism between bridges. A bridge considers that it has lost connectivity to its direct neighboring root or designated bridge if it misses three BPDUs in a row. This fast aging of the information allows quick failure detection. If a bridge fails to receive BPDUs from a neighbor, it is certain that the connection to that neighbor has been lost, as opposed to 802.1d where the problem could have been anywhere on the path to the root.

Reference:

http://www.cisco.com/en/US/partner/tech/CK389/CK621/technologies_white_paper09186a0080094cfa.shtml

QUESTION 54

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 Certkiller switch to prevent these network connectivity problems? (Select two)

- 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)# no spanning-tree backbonefast
- G. Certkiller 2950(config)# no spanning-tree uplinkfast

Answer: B, C

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/CK389/CK621/technologies_tech_note09186a008009482f.shtml

QUESTION 55

On the Certkiller switched LAN, which of the following are true regarding Unidirectional Link Detection (UDLD)? (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 56

After properly configuring multiple VLANs, a The Certkiller network has decided to increase the security of its VLAN environment. Which of the following can be done on a switched network to enhance security measures? (Choose all that apply).

- A. If a port is connected to a "Foreign" device, make sure to disable CDP, DTP, PagP, UDLD, and any other unnecessary protocol, and to enable Uplinkfst/BPDU guard on it.
- B. Enable the rootguard feature to prevent a directly or indirectly connected STP-capable device to affect the location of the root bridge.
- C. Configure the VTP domains appropriately or turn off VTP altogether if you want to limit or prevent possible undesirable protocol interactions with regard to network-wide VLAN configuration.
- D. Disable all unused ports and place them in an unused VLAN to avoid unauthorized access.
- E. Set the native VLAN ID to match the port VLAN ID (PVID) of any 802.1Q trunks to prevent spoofing from one VLAN to another.

Answer: B, C, D

Explanation:

The root guard feature is designed to provide a way to enforce the root bridge placement in the network, and to prevent unauthorized devices from becoming the root.

Turning off VTP if it is not used is generally a good idea, as a new switch with a higher ID value that is inserted into the VTP domain can be used to modify and delete all of the VLANs in an existing network.

It is also a best practice to disable and isolate all unused ports, as this will prevent unauthorized users from entering the LAN, and plugging into the network via an unused port.

Incorrect Answers:

A. UDLD is a useful feature that provides no security risks. It is recommended to have this feature enabled. BPDU guard and root guard are similar, but their impact is different. BPDU guard disables the port upon BPDU reception if portfast is enabled on the port. This effectively denies devices behind such ports to participate in STP.

E. If a user's native VLAN ID is the same as the port VLAN ID (PVID) of the 802.1Q trunk, then the user can send frames from his VLAN and have them "hop" to other VLANs. This weakness is part of the 802.1Q specification and does not apply to Cisco ISL trunking ports.

The workaround for this threat is to ensure that every 802.1Q trunking port has a PVID, or native VLAN ID, that is unique throughout the campus network.

QUESTION 57

802.1s is being utilized in the Certkiller LAN. 802.1s defines deployment for which of the following?

- A. One STP instance per set of Bridges

- B. One global instance for all VLANs
- C. One STP instance for each VLAN
- D. One STP instance per set of VLANs
- E. None of the above

Answer: D

Explanation:

802.1s for MST is an amendment to 802.1Q. MST extends the IEEE 802.1w rapid spanning tree (RST) algorithm to multiple spanning trees. This extension provides both rapid convergence and load balancing in a VLAN environment. MST converges faster than PVST+. MST is backward compatible with 802.1D STP, 802.1w (rapid spanning tree protocol [RSTP]), and the Cisco PVST+ architecture.

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

Reference:

<http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.1E/native/configuration/guide/spantree.html>
#

QUESTION 58

In comparing STP protocols used within the CK LAN, how are BPDUs handled by 802.1w compared to 802.1D?

- A. 802.1w bridges only relay BPDUs received from the root
- B. 802.1d bridges do not relay BPDUs
- C. 802.1d bridges only relay BPDUs received from the root
- D. 802.1w bridges do not relay BPDUs
- E. None of the above

Answer: C

Explanation:

With 802.1D, BPDUs are sent every hello-time, and not simply relayed anymore. With 802.1D, a non-root bridge only generates BPDUs when it receives one on the root port. In fact, a bridge relays BPDUs more than it actually generates them. This is not the case with 802.1w. A bridge now sends a BPDU with its current information every <hello-time> seconds (2 by default), even if it does not receive any from the root bridge.

Reference: <http://www.cisco.com/warp/public/473/146.html>

QUESTION 59

What type of ports on switch CK1 would Cisco spanning-tree Port-Fast BPDU-Guard be most appropriate for?

- A. Alternate ports
- B. MST ports
- C. Designated ports
- D. Root ports
- E. None of the above

Answer: C

Explanation:

Understanding How PortFast BPDUs Work:

To prevent loops from occurring in a network, the PortFast mode is supported only on nontrunking access ports because these ports typically do not transmit or receive BPDUs. The most secure implementation of PortFast is to enable it only on ports that connect end stations to switches. Because PortFast can be enabled on nontrunking ports connecting two switches, spanning tree loops can occur because BPDUs are still being transmitted and received on those ports.

The PortFast BPDUs guard feature prevents loops by moving a nontrunking port into an errdisable state when a BPDU is received on that port. When the BPDU guard feature is enabled on the switch, spanning tree shuts down PortFast-configured interfaces that receive BPDUs, instead of putting them into the spanning tree blocking state. In a valid configuration, PortFast-configured interfaces do not receive BPDUs. If a PortFast-configured interface receives a BPDU, an invalid configuration exists, such as connection of an unauthorized device. The BPDU guard feature provides a secure response to invalid configurations because the administrator must manually put the interface back in service.

QUESTION 60

RSPT is defined as which of the following?

- A. RSPT is the 802.1w standard that provides faster spanning-tree convergence over 802.1D 1988 after a topology change and include feature equivalent to Cisco Port-Fast, Uplink-fast and Backbone-fast
- B. RSPT is the 802.1s/w standard and provides
- C. RSPT is the 802.1w standard version of cisco's PVST+
- D. RSPT is the 802.1w standard that provides faster spanning-tree convergence over 802.1D. 1988 after a topology change and, it is also includes features equivalent cisco BPDU-Guard, Root-Guard and Loop-Guard
- E. None of the above

Answer: C

Explanation:

The Rapid per-VLAN Spanning-Tree plus (Rapid-PVST+) is based on IEEE 802.1W Rapid Spanning Tree Protocol (RSPT) for rapid convergence of the spanning tree upon network failure and topology changes.

Reference:

http://www.cisco.com/en/US/docs/switches/lan/catalyst3750/software/release/12.1_14_ea1/release/notes/OL421

QUESTION 61

Loops can cause network failures because of excessive traffic. What CAT6500 feature can be used to prevent this excessive traffic within the Certkiller LAN? (Select 2 answers)

- A. BPDU-Guard
- B. Storm control
- C. Storm suppression
- D. Loop guard
- E. Broadcast suppression

Answer: B, D

Explanation:

Loop Guard: The STP loop guard feature provides additional protection against Layer 2 forwarding loops (STP loops). An STP loop is created when an STP blocking port in a redundant topology erroneously transitions to the forwarding state. This usually happens because one of the ports of a physically redundant topology (not necessarily the STP blocking port) no longer receives STP BPDUs. In its operation, STP relies on continuous reception or transmission of BPDUs based on the port role. The designated port transmits BPDUs, and the non-designated port receives BPDUs.

When one of the ports in a physically redundant topology no longer receives BPDUs, the STP conceives that the topology is loop free. Eventually, the blocking port from the alternate or backup port becomes designated and moves to a forwarding state. This situation creates a loop.

The loop guard feature makes additional checks. If BPDUs are not received on a non-designated port, and loop guard is enabled, that port is moved into the STP loop-inconsistent blocking state, instead of the listening / learning / forwarding state. Without the loop guard feature, the port assumes the designated port role. The port moves to the STP forwarding state and creates a loop.

Storm Control: A traffic storm occurs when packets flood the LAN, creating excessive traffic and degrading network performance. The traffic storm control feature prevents LAN ports from being disrupted by a broadcast, multicast, or unicast traffic storm on physical interfaces.


Traffic storm control (also called traffic suppression) monitors incoming traffic levels over a 1-second traffic storm control interval and, during the interval, compares the traffic level with the traffic storm control level that you configure. The traffic storm control level is a percentage of the total available bandwidth of the port. Each port has a single traffic storm control level that is used for all types of traffic (broadcast, multicast, and unicast).

Traffic storm control monitors the level of each traffic type for which you enable traffic storm control in 1-second traffic storm control intervals. Within an interval, when the

ingress traffic for which traffic storm control is enabled reaches the traffic storm control level that is configured on the port, traffic storm control drops the traffic until the traffic storm control interval ends.

QUESTION 62

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 63

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 64

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_chapter09186a

QUESTION 65

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 66

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 67

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 version1 and VTP version2 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 version2 unless every network device in the VTP domain supports version2.

Reference:

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

QUESTION 68

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 Catalyst6500 series switch is in VTP server mode and is in the no-management domain state until the switch receives an advertisement for a domain over a trunk link or you configure a management domain.

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 Catalyst6500 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 69

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 70

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 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 provide 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 71

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 72

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 73

Certkiller is using extended VLANs (VLAN IDs 1006-4094) on their switches. What should the 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 74

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. 802.1Q 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.

For 802.1Q trunking, one VLAN is not tagged. This VLAN is called native VLAN. The native VLAN is used for untagged traffic when the port is in 802.1Q trunking mode.

While configuring 802.1Q trunking, it is very important to keep in mind that the native VLAN must be configured the same on each side of the trunk link.

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. The ISL header supports 10 bits for ISL identification. Each bit can be one of two values, allowing for the 1024 unique VLANs.

C. ISL encapsulation adds 30 bytes to the entire frame, while the 802.1Q tag is only 4 bytes in length.

QUESTION 75

The Certkiller network is illustrated below:



In the shown diagram, Server Certkiller 1's default gateway points to Router Certkiller 1's VLAN1 interface and Server Certkiller 2's default gateway points to Router Certkiller 2's VLAN2 interface. Between Switch Certkiller 1 and Certkiller 2, both VLANs 1 and 2 are being forwarded over a trunk. When there is data transfer between the servers workstations, WS Certkiller 1 and WS Certkiller 2 see a lot of input traffic.

How can we limit this problem?


- A. Increase aging time on routers
- B. Disable MAC address aging time on the switches
- C. Disable ARP timeout on routers
- D. Reduce MAC address aging time on the switches
- E. Bring ARP aging time on Routers and MAC address aging time on switches close to each other

Answer: E

Explanation:

The problem described in this question is related to asymmetric routing, due to each workstation having different default gateways. The default ARP cache aging time on a router is 4 hours. The default aging time of the switch content-addressable memory (CAM) entry is 5 minutes. The ARP aging time of the host workstations is not significant for this discussion. However, the example sets the ARP aging time to 4 hours.

This diagram illustrates this issue. This topology example includes Catalyst 6500s with Multilayer Switch Feature Cards (MSFCs) in each switch. The switches are interconnected via a trunk which carries traffic for VLAN 1 and VLAN 2.



Consequences of Asymmetric Routing:

Consider the case of the continuous ping of host B by host

A. Remember that host A

sends the echo packet to MSFC1, and host B sends the echo reply to MSFC2, which is in an asymmetric routing state. The only time that Switch 1 learns the source MAC of host B is when host B replies to an ARP request from MSFC1. This is because host B uses MSFC2 as its default gateway and does not send packets to MSFC1 and, consequently, Switch 1. Since the ARP timeout is 4 hours by default, Switch 1 ages the MAC address of host B after 5 minutes by default. Switch 2 ages host A after 5 minutes. As a result, Switch 1 must treat any packet with a destination MAC of host B as an unknown unicast. The switch floods the packet that comes from host A and is destined for host B out all ports. In addition, because there is no MAC address entry host B in Switch 1, there is no MLS entry as well.

The echo reply packets that come from host B experience the same issue after the MAC address entry for host A ages on Switch 2. Host B forwards the echo reply to MSFC2, which in turn routes the packet and sends it out on VLAN 1. The switch does not have an entry host A in the MAC address table and must flood the packet out all ports in VLAN 1.

Asymmetric routing issues do not break connectivity. However, asymmetric routing can cause excessive unicast flooding and MLS entries that are missing. There are three configuration changes that can remedy this situation:

1. Adjust the MAC aging time on the respective switches to 14,400 seconds (4 hours) or longer.
2. Change the ARP timeout on the routers to 5 minutes (300 seconds).
3. Change the MAC aging time and ARP timeout to the same timeout value.

Reference:

http://www.cisco.com/en/US/tech/CK648/CK362/technologies_tech_note09186a0080094afd.shtml

QUESTION 76

You need to monitor a remote Certkiller switch using a sniffer. What feature shown below can be used to transport monitoring session traffic from a Catalyst switch in this Certkiller network across an IP cloud to a Sniffer on a remote site?


- A. ERSPAN
- B. Protocol filtering
- C. SPAN
- D. None of the above
- E. RSPAN
- F. None of the above

Answer: A

Explanation:

ERSPAN supports source ports, source VLANs, and destinations on different switches, which provides remote monitoring of multiple switches across your network (see Figure below). ERSPAN uses a GRE tunnel to carry traffic between switches.

ERSPAN Configuration:



Reference:

<http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/span.html#wp1059>

QUESTION 77

All Certkiller switches are configured as VTP transparent mode. What is a key advantage of configuring all switches in the Certkiller network to this mode?

- A. It reduces the total number of VLANs required in the enterprise network.
- B. It allows for more rapid deployment of VLANs throughout the enterprise.
- C. It reduces the size of the spanning tree network, and as a result, improves STP convergence time.
- D. It prevents a network administrator from accidentally deleting VLAN information from all switches.
- E. It ensures consistency between VLAN numbering for all switches in the switched network.
- F. None of the above

Answer: D

Explanation:

When inserting a VTP client or server with a higher configuration revision number, the other switches will delete their configuration information and take the VLAN information from the inserted switch. The only way to get the deleted information back is to add the missing VLANs and delete the unwanted VLANs. To avoid this you should set the switch you're inserting into the network to transparent mode because that resets the configuration number, then switch it back to client or server mode. Another way of resetting the configuration number is to change the domain name to something else, like " Certkiller " and then change it back.

QUESTION 78

Certkiller is a service provider that wants to offer service for transporting dot1q trunk traffic between remote customer locations. Certkiller has Catalyst switches in its network with ISL trunks in the core. What feature can the service provider use with current setup to provide the service to the customer over a single VLAN?

- A. Dot1q Tunneling
- B. VLAN mapping
- C. None of the above
- D. Layer 2 Protocol Tunneling
- E. VLAN translation
- F. None of the above

Answer: A

Explanation:

Understanding 802.1Q Tunneling:

The VLAN ranges required by different customers in the same Service Provider network might overlap, and customer traffic through the infrastructure might be mixed. Assigning a unique range of VLAN IDs to each customer would restrict customer configurations and could easily exceed the VLAN limit (4096) of the 802.1Q specification.

802.1Q tunneling enables Service Providers to use a single VLAN to support customers who have multiple VLANs, while preserving customer VLAN IDs and keeping traffic in different customer VLANs segregated.

A port configured to support 802.1Q tunneling is called a tunnel port. When you configure tunneling, you assign a tunnel port to a VLAN ID that is dedicated to tunneling. Each customer requires a separate Service Provider VLAN ID, but that Service Provider VLAN ID supports VLANs of all the customers.

Customer traffic tagged in the normal way with appropriate VLAN IDs comes from an 802.1Q trunk port on the customer device and into a tunnel port on the Service Provider edge switch. The link between the customer device and the edge switch is asymmetric because one end is configured as an 802.1Q trunk port, and the other end is configured as a tunnel port. You assign the tunnel port interface to an access VLAN ID that is unique to each customer.

Reference:

<http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/20ew/configuration/guide/tunnel.html#wp102>

QUESTION 79

In order to increase the security in the Certkiller network, port security was configured on all switches. What statement shown below is true about port security?

- A. When you enable port security on a port, any static or dynamic CAM entries associated with the port are treated as secure.
- B. If a security violation occurs, the Link LED for that port turns orange, and a link-down trap is sent to the Simple Network Management Protocol (SNMP) manager.
- C. Port security can be configured on a SPAN source port.
- D. Port security can be configured on a SPAN destination port
- E. Port security can be configured on a trunk port.
- F. None of the above

Answer: B

Explanation:

Port Security Configuration Guidelines:

When a secure port receives a packet, the source MAC address of the packet is compared to the list of secure source addresses that were manually configured or autoconfigured (learned) on the port. If a MAC address of a device attached to the port differs from the list of secure addresses, the port either shuts down permanently (default mode), shuts down for the time you have specified, or drops incoming packets from the insecure host. The port's behavior depends on how you configure it to respond to a security violation. If a security violation occurs, the Link LED for that port turns orange, and a link-down trap is sent to the Simple Network Management Protocol (SNMP) manager. An SNMP trap is not sent if you configure the port for restrictive violation mode. A trap is sent only if you configure the port to shut down during a security violation.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps700/products_configuration_guide_chapter09186a008007f

Incorrect Answers:

A, C, D, E. These incorrect answers can be summarized in the following statements:

1. You cannot configure port security on the trunk port of a 6500 with Cat OS.
2. You cannot enable port security on a SPAN destination port of the 6500 with Cat OS.
3. You cannot configure dynamic, static, or permanent CAM entries on a secure port.
4. When you enable port security on a port, any static or dynamic CAM entries associated with the port are cleared; any currently configured permanent CAM entries are treated as secure.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a008007f

QUESTION 80

When an IP packet is Layer 3-switched from a source in one VLAN in the Certkiller LAN to a destination in another VLAN, what field in a packet will be rewritten?

- A. Layer 3 source address
- B. Layer 3 TTL
- C. Layer 2 TTL
- D. Layer 3 Transport Protocol
- E. Layer 3 destination address
- F. None of the above

Answer: B

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 Header		Layer 3 IP Header				Data	FCS
	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:

1. Destination MAC address: The MAC address of the next hop must be written to the rewritten frame.
2. Source MAC address: The source MAC address must be written to the MAC address of the router.
3. IP TTL: This must be decremented by one, as per the normal rules of IP routing.
4. 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 81

Part of the Certkiller network is shown below:



Your work as a network technician at Certkiller .com. Your boss, Mrs. Certkiller, asks you to study the exhibit carefully. She tells you that both PC and server are using HSRP virtual address for default gateway in their respective vlans. You connected computer Certkiller B to Switch Certkiller 3 and captured some packets in vlan 100. You have noticed that unicast packets from the Server Certkiller C vlan 20 to computer Certkiller in vlan 100 are constantly being flooded affecting the performance of other devices in vlan 100. What is the most appropriate way to fix this issue?


- A. Configure the MAC address table aging and ARP aging timers to match on switches Certkiller 3 and Certkiller 4.
- B. Configure the MAC address of the PC in VLAN 100 as static on switch Certkiller 4.
- C. Configure the MAC address of the server Certkiller C in VLAN 100 as static on Switch Certkiller 3.
- D. Disable HSRP on switch Certkiller 3.
- E. Configure static ARP entry for the PC address in VLAN 100 on switch Certkiller 3.
- F. None of the above

Answer: A

Explanation:

The default ARP cache aging time on a router is 4 hours. The default aging time of the switch content-addressable memory (CAM) entry is 5 minutes. The ARP aging time of the host workstations is not significant for this discussion. However, the example sets the ARP aging time to 4 hours.

This diagram illustrates this issue. This topology example includes Catalyst 6500s with Multilayer Switch Feature Cards (MSFCs) in each switch. Although this example uses MSFCs, you can use any router instead of the MSFC. Example routers that you can use include the Route Switch Module (RSM), Gigabit Switch Router (GSR), and Cisco 7500. The hosts are directly connected to ports on the switch. The switches are interconnected via a trunk which carries traffic for VLAN 1 and VLAN 2.



Consequences of Asymmetric Routing

Consider the case of the continuous ping of host B by host

A. Remember that host A

sends the echo packet to MSFC1, and host B sends the echo reply to MSFC2, which is in an asymmetric routing state. The only time that Switch 1 learns the source MAC of host B is when host B replies to an ARP request from MSFC1. This is because host B uses MSFC2 as its default gateway and does not send packets to MSFC1 and, consequently, Switch 1. Since the ARP timeout is 4 hours by default, Switch 1 ages the MAC address of host B after 5 minutes by default. Switch 2 ages host A after 5 minutes. As a result, Switch 1 must treat any packet with a destination MAC of host B as an unknown unicast. The switch floods the packet that comes from host A and is destined for host B out all ports. In addition, because there is no MAC address entry host B in Switch 1, there is no MLS entry as well.

The echo reply packets that come from host B experience the same issue after the MAC address entry for host A ages on Switch 2. Host B forwards the echo reply to MSFC2, which in turn routes the packet and sends it out on VLAN 1. The switch does not have an entry host A in the MAC address table and must flood the packet out all ports in VLAN 1.

Asymmetric routing issues do not break connectivity. However, asymmetric routing can cause excessive unicast flooding and MLS entries that are missing. There are three configuration changes that can remedy this situation:

1. Adjust the MAC aging time on the respective switches to 14,400 seconds (4 hours) or longer.
2. Change the ARP timeout on the routers to 5 minutes (300 seconds).
3. Change the MAC aging time and ARP timeout to the same timeout value.

Reference:

http://www.cisco.com/en/US/tech/ CK6 48/ CK3 62/technologies_tech_note09186a0080094afd.shtml

QUESTION 82

In the Certkiller switched LAN, 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
- F. None of the above

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 83

When considering the use of DTP in the Certkiller LAN, what is true regarding it?

- A. DTP is not supported on private VLAN ports or tunnel ports
- B. It is a point-to-multipoint
- C. You always want to disable DTP
- D. It supports auto negotiation in both ISL/802.1Q
- E. It is a point-to-point protocol
- F. None of the above

Answer: A

Explanation:

From the Cisco Catalyst 3750 configuration guide:

Trunk negotiation is managed by the Dynamic Trunking Protocol (DTP), which is a Point-to-Point Protocol. However, some internetworking devices might forward DTP frames improperly, which could cause misconfigurations.

To avoid this, you should configure interfaces connected to devices that do not support DTP to not forward DTP frames, that is, to turn off DTP.

If you do not intend to trunk across those links, use the switchport mode access interface configuration command to disable trunking.

To enable trunking to a device that does not support DTP, use the switchport mode trunk and switchport nonegotiate interface configuration commands to cause the interface to become a trunk but to not generate DTP frames. Use the switchport trunk encapsulation isl or switchport trunk encapsulation dot1q interface to select the encapsulation type on the trunk port.

You can also specify on DTP interfaces whether the trunk uses ISL or IEEE 802.1Q encapsulation or if the encapsulation type is autonegotiated. The DTP supports autonegotiation of both ISL and IEEE802.1Q trunks.

Note:

DTP is not supported on private-VLAN ports or tunnel ports.

Reference:

http://www.cisco.com/en/US/docs/switches/lan/catalyst3750/software/release/12.2_35_se/configuration/guide/sw

QUESTION 84

A new Certkiller user is complaining about poor network performance from her PC. Poor Performance, collisions and intermittent communication problems between a PC and a switch port may be a result of:

- A. The wrong wire category being used
- B. Mismatching duplex modes
- C. The port being errdisabled on the switch
- D. Mismatching speeds
- E. All of the above

Answer: B

Explanation:

Issues with autonegotiation of duplex generally do not result in link establishment issues. Instead, autonegotiation issues mainly result in performance-related issues. A duplex mismatch can result in performance issues, intermittent connectivity, and loss of communication. When you troubleshoot NIC issues, verify that the NIC and switch use a valid configuration.

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

QUESTION 85

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:

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 c	100 Mbps, Half-duplex	Duplex Mismatch 1
AUTO	100 Mbps, Full-duplex	100 Mbps, Half-duplex	100 Mbps, Full-duplex	Duplex Mismatch 1
100 Mb	100 Mb	100 Mb	100 Mb	C t

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 86

During routine maintenance, you 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 87

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 88

You are connecting a new PC with a 10/100 NIC to a Certkiller 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
- F. None of the above

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 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.

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 89

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 90

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 91

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, 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 LSBfirst.

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 92

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 93

If a Certkiller LAN switch Gigabit Ethernet or 10-Gigabit Ethernet port's receive buffer becomes full, what protocol can be used to request the remote port to delay sending frames for a specified time?

- A. 802.1U
- B. 802.3Z
- C. 802.1D
- D. 802.3
- E. 802.3AF
- F. None of the above

Answer: B

Explanation:

802.3Z defines the standard for Gigabit Ethernet. Included in this is a flow control mechanism.

IEEE 802.3Z Flow Control:

Gigabit Ethernet ports on 802.3Z compliant switch ports use flow control to inhibit the transmission of packets to the port for a period of time; other Ethernet ports use flow control to respond to flow control requests.

If a Gigabit Ethernet port receive buffer becomes full, the port transmits a "pause" packet that tells remote ports to delay sending more packets for a specified period of time. All Ethernet ports (1000Mbps, 100Mbps, and 10Mbps) can receive and act upon "pause" packets from other devices.

Incorrect Answers:

- A: This draft defined some technical and editorial changes to the 802.1 standard.
- C: This defines the Spanning Tree Protocol
- D: This defines the CSMA/CD Ethernet standard.
- E: This defines power over Ethernet standards commonly used for VOIP applications.

QUESTION 94

You are troubleshooting problems on the Certkiller LAN and you suspect duplex mismatches. Which switch port errors are in indication of a duplex mismatch on a 10/100/1000 IEEE 802.3u gigabit Ethernet port? (Select 5)

- A. Alignment Errors
- B. FCS Errors
- C. Multiple Collisions
- D. Runts
- E. Excessive Collisions
- F. Late Collisions

Answer: A, B, D, E, F

Explanation:

The causes for the most common errors found in a 10/100/1000 LAN environment is found in the following table (Note the answers to this question in bold).

Counter	Description
---------	-------------

Alignment Errors	These are the result of collisions at half-duplex, duplex mismatch, bad hardware (NIC, cable, or port), or a connected device that generates frames that do not end with on an octet and have a bad FCS.
FCS	These are the result of collisions at half-duplex, duplex mismatch, bad hardware (NIC, cable, or port), or a connected device that generates frames with bad FCS.
Single Collisions	This is an indication of a half-duplex configuration.
Multiple Collisions	This is an indication of a half-duplex configuration.
Late Collisions	This is an indication of faulty hardware (NIC, cable, or switch port) or a duplex mismatch.
Excessive Collisions	This is an indication of overutilization of the switch port at half-duplex or duplex mismatch.
Runts	This is an indication of the result of collisions, duplex mismatch, IEEE 802.1Q (dot1q), or an Inter-Switch Link Protocol (ISL) configuration issue.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps700/products_tech_note09186a00800a7af0.shtml

QUESTION 95

While troubleshooting a LAN issue on the Certkiller network, you notice a large amount of alignment errors, FCS Errors, and late collisions. This may indicate:

- A. A half-duplex connection between the switch and an end-point on a 10/100/1000Base-T Ethernet link
- B. These errors are normal under most circumstances
- C. There is a duplex mismatch on a 1000Base-LX/LH
- D. Duplex Mismatch on a 10/100/1000Base-T Ethernet link
- E. None of the above

Answer: D

Explanation:

Many different modes of operations for Ethernet over twisted pair (10/100/1000 Base T), and most network adapters are capable of different modes of operations. In 1995, a standard was released for allowing two network adapters connected to each other to negotiate the best possible shared mode of operation. The autonegotiation standard contained a mechanism for detecting the speed but not the duplex setting of Ethernet peers that did not use autonegotiation.

When two linked interfaces are set to different duplex modes, the effect of this duplex mismatch is a network that functions much slower than its nominal speed. The primary rule for avoiding this is to avoid setting one end of a connection to full duplex and the other end to autonegotiation.

Duplex mismatch may be inadvertently caused when an administrator configures an interface to a fixed mode (e.g 100 Mbit/s full duplex) and fails to configure the remote interface, leaving it set to autonegotiate. Then, when the autonegotiation process fails, half duplex is assumed by the autonegotiating side of the link.

The resulting duplex mismatch results in a dramatically slow network, in which many collisions, and especially late collisions occur on the interface set to half-duplex, and FCS errors are seen on the full-duplex side.

Gigabit Ethernet standards require autonegotiation to be on in order to operate.

Reference: <http://en.wikipedia.org/wiki/10BASE-T>

QUESTION 96

A workstation named Certkiller 3 has been connected to a workgroup Ethernet switch using a Category 5e cable. Certkiller 3 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. Mismatched duplex settings between the workstation and the switch
- B. Port has erroneously been configured as an 802.1q trunk port
- C. Mismatched speed settings between the workstation and the switch
- D. Bad cable between the workstation and the switch
- E. Bad Network Interface Card on the workstation
- F. None of the above

Answer: A

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.

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

QUESTION 97

You want to optimize an ether-channel trunk between two Certkiller switches. What is the Cisco recommended best practice PaGP setting for ports on Ether-channel trunks?

- A. desirable - on
- B. auto - auto
- C. desirable - desirable
- D. on - on
- E. desirable - auto
- F. None of the above

Answer: C

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:

1. A port in desirable mode can successfully form an EtherChannel with another port that is in desirable or auto mode.
2. A port in auto mode can form an EtherChannel with another port in desirable mode.
3. A port in auto mode cannot form an EtherChannel with another port that is also in auto

mode, since neither port initiates negotiation.

4. 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.

5. A port in off mode cannot form a channel with any port.

Reference:

<http://www.cisco.com/en/US/tech/CK389/CK2>

[13/technologies_tech_note09186a00800949c2.shtml#pagptoconfig](http://www.cisco.com/en/US/tech/CK389/CK213/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

The following is 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 98

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 in this network?

- A. Auto-auto
- B. Desirable-on
- C. On-auto
- D. Auto-desirable
- E. None of the above

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 99

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
- F. None of the above

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:

First Side of cable goesto Second Side of cable

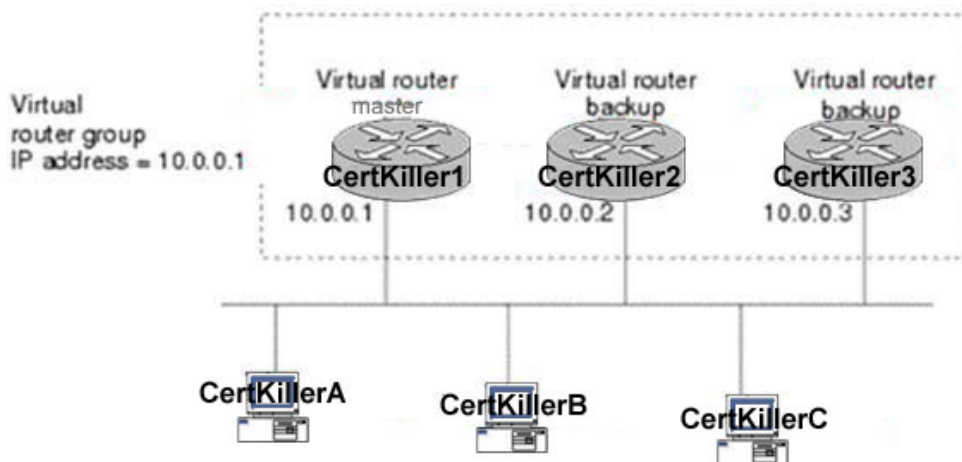
First Side of cable		goesto		Second Side of cable	
Color	Name	Pin	Pin	Name	Color
White/Orange	TX+	1	3	RX+	White/Orange
Orange	TX-	2	6	RX-	Orange
White/Green	RX+	3	1	TX+	White/Green
Green	RX-	6	2	TX-	Green

Blue	Extra Pins	4	4	Optional Pins (can be connected straight to same color pins). But not used in transmission	Blue
White/Blue		5	5		White/Blue
White/Brown		7	7		White/Brown
Brown		8	8		Brown

Reference: <http://www.infonewsindia.com/pinout/pinoutnetwork.html>

QUESTION 100

The Certkiller network is providing gateway redundancy according to the diagram shown below:



According to the diagram shown above, what protocol is being used to implement the First Hop Redundancy Connection under a multi-vendor solution?


- A. Gateway Load Balancing Protocol (GLBP)
- B. MSHRP
- C. Hot Standby Router Protocol (HSRP)
- D. Virtual Router Redundancy Protocol (VRRP)
- E. None of the above

Answer: D

Explanation:

Figure1 below shows a LAN topology in which VRRP is configured. In this example, Routers A, B, and C are VRRP routers (routers running VRRP) that comprise a virtual router. The IP address of the virtual router is the same as that configured for the Ethernet interface of Router A (10.0.0.1).

Figure1 Basic VRRP Topology



Because the virtual router uses the IP address of the physical Ethernet interface of Router A, Router A assumes the role of the virtual router master and is also known as the IP address owner. As the virtual router master, Router A controls the IP address of the virtual router and is responsible for forwarding packets sent to this IP address. Clients 1 through 3 are configured with the default gateway IP address of 10.0.0.1.


Routers B and C function as virtual router backups. If the virtual router master fails, the router configured with the higher priority will become the virtual router master and provide uninterrupted service for the LAN hosts. When Router A recovers, it becomes the virtual router master again.

Reference:

http://www.cisco.com/en/US/products/ps6350/products_configuration_guide_chapter09186a008042fbd9.html

QUESTION 101

The Certkiller network topology is shown in the following exhibit, where routers Certkiller 1 and Certkiller 2 are configured with HSRP:



Log exhibit:

```
Oct 12 13:15:41: %STANDBY-3-DUPADDR: Duplicate address 10.25.0.1
on Vlan25, sourced by 0000.0c07.ac19
Oct 13 16:25:41: %STANDBY-3-DUPADDR: Duplicate address 10.25.0.1
on Vlan25, sourced by 0000.0c07.ac19
Oct 15 22:31:02: %STANDBY-3-DUPADDR: Duplicate address 10.25.0.1
on Vlan25, sourced by 0000.0c07.ac19
Oct 15 22:41:01: %STANDBY-3-DUPADDR: Duplicate address 10.25.0.1
on Vlan25, sourced by 0000.0c07.ac19
```

Please study the exhibit carefully. Based on the information provided above, what type of issue does the error log indicate if the IP Address in the error log is located off Router Certkiller 1 WAN?

- A. HSRP standby Configuration Error
- B. HSRP Secondary Address Configuration Error
- C. HSRP Burned in Address Error
- D. Not an HSRP Problem, But rather an STP Error or Router/Switch Configuration issue
- E. None of the above

Answer: D

Explanation:

From the Cisco Troubleshoot HSRP Case Studies:

Case Study #1: HSRP Standby IP Address Is Reported as a Duplicate IP Address

These error messages can appear:

```
Oct 12 13:15:41: %STANDBY-3-DUPADDR: Duplicate address 10.25.0.1
on Vlan25, sourced by 0000.0c07.ac19
Oct 13 16:25:41: %STANDBY-3-DUPADDR: Duplicate address 10.25.0.1
on Vlan25, sourced by 0000.0c07.ac19
Oct 15 22:31:02: %STANDBY-3-DUPADDR: Duplicate address 10.25.0.1
on Vlan25, sourced by 0000.0c07.ac19
Oct 15 22:41:01: %STANDBY-3-DUPADDR: Duplicate address 10.25.0.1
on Vlan25, sourced by 0000.0c07.ac19
```

These error messages do not necessarily indicate an HSRP problem. Rather, the error messages indicate a possible Spanning Tree Protocol (STP) loop or router/switch configuration issue. The error messages are just symptoms of another problem.

In addition, these error messages do not prevent the proper operation of HSRP. The duplicate HSRP packet is ignored. These error messages are throttled at 30-second intervals. But, slow network performance and packet loss can result from the network instability that causes the STANDBY-3-DUPADDR error messages of the HSRP address.

Reference:

http://www.cisco.com/en/US/tech/CK648/CK362/technologies_tech_note09186a0080094afd.shtml

QUESTION 102

You are attempting to properly subnet the IP space of one of the Certkiller locations. For the network "200.10.10.0" there is a need for 3 loopback interfaces, 2 point to point links, one Ethernet with 50 stations and one Ethernet with 96 stations. What option below would be the most efficient (for saving IP addresses)?

- A. 200.10.10.1/32, 200.10.10.2/32, 200.10.10.3/32
200.10.10.4/30, 200.10.10.8/30
200.10.10.64/26, 200.10.10.128/25
- B. 200.10.10.0/32, 200.10.10.1/32, 200.10.10.2/32
200.10.10.4/32, 200.10.10.8/31
200.10.10.64/26, 200.10.10.128/25
- C. 200.10.10.1/32, 200.10.10.2/32, 200.10.10.3/32
200.10.10.4/31, 200.10.10.8/32
200.10.10.64/27, 200.10.10.128/27
- D. D. 200.10.10.1/31, 200.10.10.2/31, 200.10.10.3/31
200.10.10.4/30, 200.10.10.8/30
200.10.10.64/27, 200.10.10.128/26
- E. There is not enough address available on that network for these subnets.

Answer: A

Explanation:

Choice A will provide the necessary subnetting to achieve all of the necessary network/host combinations that are needed at this location. Since each loopback interface is used only as an internal network to the router, no actual hosts are needed so the host subnet mask of a /32 will be sufficient for all three loopback interfaces. Point to point networks commonly use the /30 subnet mask, since in any point to point link, only two hosts are needed (one for the serial interface of the router at each end). Finally, the subnet mask of /26 will provide for 62 useable addresses and the final subnet mask of /25 will provide for 126 useable IP hosts on the second Ethernet network.

Incorrect Answers:

B. This answer includes /32 network masks for the Pt-Pt links, which can not be used for the two point to point links, since they do not provide any useable IP addresses.

C. This answer includes /32 network masks for the Pt-Pt links, which can not be used for the two point to point links, since they do not provide any useable IP addresses. In addition, the subnet mask of /27 provides for only 30 useable IP addresses for the two Ethernet segments.

D. The /27 subnet mask will provide for only 30 useable IP addresses for one of the Ethernet networks, which is insufficient.

Note: Until IOS version 12.2, a /31 address could not be used for point to point links because it does not provide useable IP addresses. However, /31 addressing for point to point links has been an option in Cisco IOS since version 12.2 and is an IEEE standard defined in RFC 3021 <http://www.faqs.org/rfcs/rfc3021.html>

QUESTION 103

What option is the best way to apply CIDRI if a service provider wants to summarize the following addresses: 202.1.0.0/16, 202.2.0.0/16, 202.3.0.0/16, 202.5.0.0/16, 202.6.0.0/16, 202.7.0.0/16?

- A. 202.0.0.0/14, 202.4.0.0/15, 202.6.0.0/16, 202.7.0.0/16

- B. 202.0.0.0/16
- C. 202.4.0.0/14, 202.2.0.0/15, 202.2.0.0/16, 202.1.0.0/16
- D. 202.4.0.0/14, 202.2.0.0/15, 202.1.0.0/16
- E. 202.0.0.0/18

Answer: D

Explanation:

The Network 200.4.0.0/14 will encompass the 200.5.0.0, 200.6.0.0 and 200.7.0.0 networks. The second summarization, 200.2.0.0/15 will take care of both the 200.2.0.0 and 200.3.0.0 networks. Finally, the last network is needed in order to include the only remaining network, which is 200.1.0.0/16. This will summarize all 6 networks using only 3 statements.

Incorrect Answers:

- A. Although this answer will also fulfill the needs of summarizing all 6 networks, it is not the most efficient way as 4 network entries are needed here, instead of only 3 in answer choice D.
- B. This will mean that only the 200.0.0.0/16 network is advertised, which is not even one of the networks that need to be summarized.
- C. This is also not the most efficient choice, as the third statement (200.2.0.0/16) is redundant, since this network is already included in the 200.2.0.0/15 summarized route.
- E. This network mask would not include all of the needed networks.

QUESTION 104

What should be configured on redundant routers to support the need for a default gateway on LAN network hosts when there are two gateway routers providing connectivity to the rest of the network?

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

Answer: D

Explanation:

The Hot Standby Router Protocol (HSRP) provides network redundancy for IP networks, ensuring that user traffic immediately and transparently recovers from default gateway failures in the network. It is implemented on networks where there are two or more gateway routers that provide connectivity to the rest of the network with the standby router acting as an automatic failover should the primary router fail.

Incorrect Answers:

A, E: DHCP and BOOTP are used to provide IP addressing, DNS, and default gateway information to end user hosts.

B, C: RIP and OSPF are routing protocols, and will not provide for automatic default gateway redundancy for PC hosts.

QUESTION 105

Which Network Address Translation type describes the internal network that uses private network addresses?

- A. Inside local
- B. Inside global
- C. Outside local
- D. Outside global
- E. None of the above
- F. None of the above

Answer: A

Explanation:

Cisco uses the term inside local for the private IP addresses and inside global for the public IP addresses. The enterprise network that uses private addresses, and therefore that needs NAT, is the "inside" part of the network. The Internet side of the NAT function is the "outside" part of the network. A host that needs NAT has the IP address it uses inside the network, and it needs an IP address to represent it in the outside network.

Incorrect Answers:

- B. The inside global address is a legitimate IP address assigned by the NIC or service provider that represents one or more inside local IP addresses to the outside world.
 - C. The outside local address is the IP address of an outside host as it appears to the inside network. Not necessarily a legitimate address, it is allocated from an address space routable on the inside.
 - D. The outside global address the IP address assigned to a host on the outside network by the host's owner. The address is allocated from a globally routable address or network space.
-

QUESTION 106

A network administrator of Certkiller .com is using a private IP address space for the company network with many to one NAT to allow the users to have access to the Internet. Shortly after this, a web server is added to the network. What must be done to allow outside users access to the web server via the Internet?

- A. Use a dynamic mapping with the reversekeyword.
- B. Place the server's internal IP address in the external NAT records.
- C. There must be a static one to one NAT entry for the web server's address.
- D. Nothing more needs to be done as dynamic NAT is automatic.
- E. Place the server's IP address into the NAT pool.


Answer: C

Explanation:

Without a static NAT mapping, the server will be NATed out of the NAT pool. Since many to one NAT (PAT) uses dynamic port mapping, no outside stations will be able to reach the server consistently.

QUESTION 107

The Certkiller LAN is displayed below:



Users in VLAN X behind the Access Layer switch complain that they cannot access the Internet when both layer 3 links in the San Jose switch fail. When only one of the L3 links in San Jose fail, users are still able to get to the Internet. Which command should be used to ensure connectivity to the Internet, even if both L3 San Jose links fail?

- A. Standby track
- B. Standby timer
- C. Standby authentication
- D. Standby use-bia
- E. Standby priority

Answer: A

Explanation:

Interface tracking allows you to specify another interface on the router for the HSRP process to monitor in order to alter the HSRP priority for a given group. If the specified interface's line protocol goes down, the HSRP priority of this router is reduced, allowing another HSRP router with higher priority to become active.

Incorrect Answers:

- B. Standby timer is used to set the hello time between HSRP routers.
- C. Standby authentication is used as a security measure between HSRP routers, using a password authentication process.
- D. By default, HSRP uses the preassigned HSRP virtual MAC address on Ethernet and FDDI, or the functional address on Token Ring. To configure HSRP to use the interface's

burnt-in address as its virtual MAC address, instead of the default, use the standby use-bia command.

E. The priority is used to determine which router will be the active one.

Reference:

<http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/cs009.htm>

QUESTION 108

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:

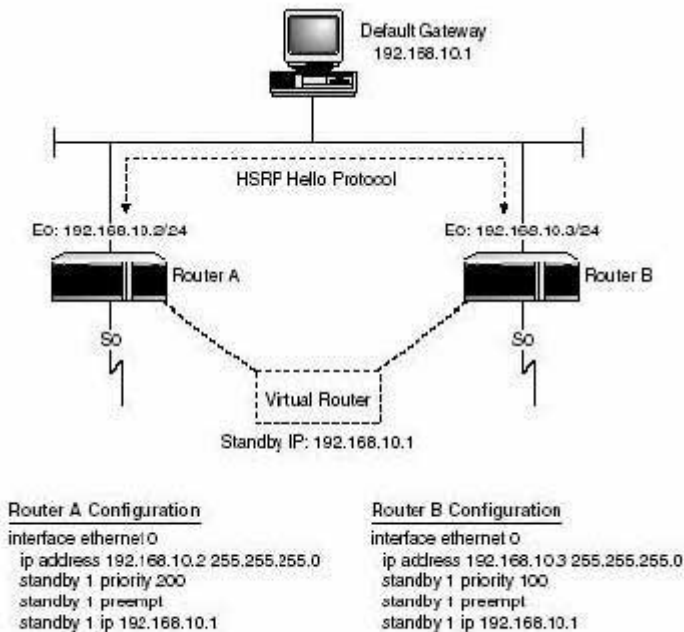
Hot Standby Routing Protocol enables a virtual gateway on LAN networks, and enables the ability to provide a single default gateway for hosts to use, even though multiple routers are in use. This can provide for a level of load balancing, along with automatic failover capability for redundancy.

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.

**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.

QUESTION 109

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 access-lists.
- 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, 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 110

The Certkiller network is displayed below:



The log output of router Certkiller A shows:

```
Jan 9 08:00:42.623: %STANDBY-6-STATECHANGE: Standby: 49: Vlan149
state Active -> Speak
Jan 9 08:00:56.011: %STANDBY-6-STATECHANGE: Standby: 49: Vlan149
state Speak -> Standby
Jan 9 08:01:03.011: %STANDBY-6-STATECHANGE: Standby: 49: Vlan149
state Speak -> Standby
Jan 9 08:01:29.427: %STANDBY-6-STATECHANGE: Standby: 49: Vlan149
state Standby -> Active
Jan 9 08:01:36.808: %STANDBY-6-STATECHANGE: Standby: 49: Vlan149
state Active -> Speak
Jan 9 08:01:43.808: %STANDBY-6-STATECHANGE: Standby: 49: Vlan149
state Speak -> Standby
```

Refer to the exhibits shown above.

What type of issue does the error log on router Certkiller A indicate?

- A. A Physical Layer Problem
- B. An HSRP standby router Configuration Error
- C. An HSRP router interfaces are in the wrong VLAN
- D. Port fast is enables on both HSRP routers
- E. None of the above

Answer: A

Explanation:

Case Study: HSRPStateContinuously Changes (Active, Standby, Speak)

These error messages can appear:

```
Jan 9 08:00:42.623: %STANDBY-6-STATECHANGE: Standby: 49: Vlan149
```

```
state Standby -> Active
```

```
Jan 9 08:00:56.011: %STANDBY-6-STATECHANGE: Standby: 49: Vlan149
```

```
state Active -> Speak
```

```
Jan 9 08:01:03.011: %STANDBY-6-STATECHANGE: Standby: 49: Vlan149
```

```
state Speak -> Standby
```

```
Jan 9 08:01:29.427: %STANDBY-6-STATECHANGE: Standby: 49: Vlan149
```

```
state Standby -> Active
```

```
Jan 9 08:01:36.808: %STANDBY-6-STATECHANGE: Standby: 49: Vlan149
```


Vlan149 state Active -> Speak

Jan 9 08:01:43.808: %STANDBY-6-STATECHANGE: Standby: 49:

Vlan149 state Speak -> Standby

These error messages describe a situation in which a standby HSRP router did not receive three successive HSRP hello packets from its HSRP peer. The output shows that the standby router moves from the standby state to the active state. Shortly thereafter, the router returns to the standby state. Unless this error message occurs during the initial installation, an HSRP issue probably does not cause the error message. The error messages signify the loss of HSRP hellos between the peers. When you troubleshoot this issue, you must verify the communication between the HSRP peers. A random, momentary loss of data communication between the peers is the most common problem that results in these messages.

There are several possible causes for the loss of HSRP packets between the peers. The most common problems are physical layer problems or excessive network traffic caused by spanning tree issues.

Reference:

http://www.cisco.com/en/US/tech/CK648/CK362/technologies_tech_note09186a0080094afd.shtml

QUESTION 111

New Cisco cache engines are being implemented into the Certkiller network. Which two statements are FALSE regarding the Web Cache Communication Protocol (WCCP) version 2?

- A. It allows redirection of traffic other than HTTP, including a variety of UDP and TCP traffic
- B. Only one router can redirect contents requests
- C. Multiple routers can redirect contents requests
- D. It works only with IP networks
- E. The Cache Engine defines one central "home router", and stores it in its memory.

Answer: B, E

Explanation:

WCCP transparently redirects Hypertext Transfer Protocol (HTTP) requests going to the intended server to a Cache Engine. End users do not know that the page came from the Cache Engine rather than the originally requested web server.

WCCP Version 2 now contains the following new features:

Multiple router support

Improved security

Faster throughput

Redirection of multiple TCP port-destined traffic

Load distributing applications capability

Client IP addressing transparency

Multirouter Support:

WCCP Version 2 enables a series of Cache Engines, called a Cache Engine cluster, to connect to multiple routers. This feature provides redundancy and a more distributed

architecture for instances when a Cache Engine needs to connect to a large number of interfaces. This strategy also has the benefit of keeping all the Cache Engines in a single cluster, avoiding unnecessary duplication of web pages across several clusters.

Reference:

http://www.cisco.com/en/US/products/sw/conntsw/ps547/products_user_guide_chapter09186a008009f1ae.html

QUESTION 112

Select the mode that NTP servers can associate with each other: (Select all that apply)

- A. Client and Server
- B. Peer
- C. Broadcast/Multicast

Answer: B

Explanation:

NTP synchronizes timekeeping among a set of distributed time servers and clients. This synchronization allows events to be correlated when system logs are created and other time-specific events occur. An NTP association can be a peer association (meaning that this system is willing to either synchronize to the other system or to allow the other system to synchronize to it), or it can be a server association (meaning that only this system will synchronize to the other system, and not the other way around). An NTP server can only be configured as a peer to another NTP server.

QUESTION 113

A customer wants to install a new frame-relay router in their network. One goal is to ensure that the new router has the correct configuration to maintain a consistent time and date, like the other routers in the network. The customer wants to configure the new router to periodically poll a UNIX server that has a very reliable and stable clock for the correct time. This will synchronize the new router's clock with the UNIX server. What command should be configured on the new router to synchronize its clock with a centralized clock service?

- A. ntp master
- B. ntp server
- C. ip ntp clock
- D. ntp peer
- E. sntp master
- F. All of the above

Answer: B

Explanation:


To allow the system clock to be synchronized by a time server, use the ntp server global configuration command.

Incorrect Answers:

- A. This command will configure the router to act as the NTP master server, providing time information to other devices.
- C. This is an invalid command.
- D. Use this command if you want to allow the router to synchronize with an NTP peer, or vice versa.
- E. SNTP is a simpler version of NTP used by lower end Cisco devices. If SNTP were to be used, the correct syntax would be "sntp server" and not "sntp master."

QUESTION 114

Your network is shown in the following exhibit:



DHCP Client

You want all PC's at CK1 to be able to obtain their IP address dynamically from the DHCP server that resides at CK2 . Currently, the hosts are not able to obtain an IP address, and they are receiving error messages saying that the DHCP server is busy or is unavailable. What must be done to enable these PC's to obtain dynamic IP addresses.

- A. Enable the command "ip helper-address 10.1.1.100" under the S0 interface on Router CK1 .
- B. Enable the command "ip helper-address 10.1.1.100" under the E0 interface on Router CK1 .
- C. Enable the command "ip helper-address 255.255.255.255" under the E0 interface on Router CK1 .
- D. Enable the command "ip helper-address 255.255.255.255" under the S0 interface on Router CK2 .
- E. Enable the command "ip helper-address 10.1.1.100" in global configuration mode on router CK1 .
- F. Enable the command "ip helper-address 10.1.1.100" in global configuration mode on router CK2 .

Answer: B

Explanation:

By default, routers drop all broadcast packets sent through them. Because DHCP clients use BOOTP packets, which are broadcasted to all hosts (255.255.255.255), they will be dropped by router CK1 . The "ip helper-address" command enables the router to forward these BOOTP broadcast packets to a specific host, as specified by the address following

the "ip helper-address" command. Note that this command must be placed on the router's interface that is receiving the broadcast packets from the hosts, which is E0 of the CK1 router.

QUESTION 115

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 116

A new Certkiller office needs to use their Cisco router as a DHCP server. When configuring the Cisco IOS DHCP Server to handle DHCP in this Certkiller office, what two statements are required to get DHCP to work? (Choose two)

- A. Configure Manual Bindings
- B. Configure a DHCP address pool
- C. Configure a DHCP server boot file
- D. Exclude IP Address not to be used in DHCP scopes
- E. Configure the timeout value for Ping Packets

Answer: B, D

Explanation:

To configure the Cisco IOS DHCP Server feature, first configure a database agent or disable conflict logging, then configure IP addresses that the DHCP server should not assign (excluded addresses) and should assign (a pool of available IP addresses) to requesting clients. These configuration tasks are explained in the following sections.

Each task in the following list is identified as required or optional.

Configuring a DHCP Database Agent or Disabling DHCP Conflict Logging (Required)

Excluding IP Addresses (Required)

Configuring a DHCP Address Pool (Required)

Configuring Manual Bindings (Optional)

Configuring a DHCP Server Boot File (Optional)

Configuring the Number of Ping Packets (Optional)

Configuring the Timeout Value for Ping Packets (Optional)

Enabling the Cisco IOS DHCP Server Feature (Optional)

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1830/products_feature_guide09186a008008743b.html

QUESTION 117

Router CK1 has been configured with the "ntp peer" command. What does the use of the NTP peer statement imply when configured in this Certkiller router?

- A. Static Client
- B. Symmetric Active Mode
- C. Static Server
- D. NTP Broadcast Client
- E. None of the above

Answer: B

Explanation:

When a networking device is operating in the symmetric active mode, it polls its assigned time-serving hosts for the current time and it responds to polls by its hosts. Because this is a peer-to-peer relationship, the host also retains time-related information about the local networking device that it is communicating with. This mode should be used when there is a number of mutually redundant servers that are interconnected via diverse network paths. Most stratum 1 and stratum 2 servers on the Internet today adopt this form of network setup. Use the peer command to individually specify the time-serving hosts that you want your networking device to consider synchronizing with and to set your networking device to operate in the symmetric active mode.

Reference:

http://www.cisco.com/en/US/docs/ios_xr_sw/iosxr_r3.3/system_management/configuration/guide/yc33ntp.html

QUESTION 118

From the list of applications below, identify the TCP port numbers used for: 443, 389, 137, 110, and 23 in the proper sequence:

- A. BGP, POP3, SNMP, TFTP, Telnet
- B. LDAP, SNMP, TFTP, POP3, Telnet
- C. HTTPS, SNMP, POP3, DNS, Telnet
- D. Finger, DHCP Server, NetBios Name Server, POP3, Telnet
- E. HTTPS, LDAP, NetBios Name Server, POP3, Telnet

F. None of the above

Answer: E

Explanation:

The following shows the TCP port numbers used with the associated applications:

HTTPS (secure WWW): 443

LDAP: 389 on the directory server

NetBios Name Server: 137

POP3: 110

Telnet: 23

Incorrect Answers:

A. BGP uses TCP port 179.

B, C. SNMP uses TCP port 161.

D. Finger uses TCP port 79 while DHCP uses 67 (BOOTP)

A complete list of TCP port numbers and their assignments can be found here:

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

QUESTION 119

Certkiller users transfer a large number of files using TFTP. Upon which protocol or protocols does this protocol rely on?

A. ICMP and UDP

B. IP and TCP

C. NFS

D. FTP

E. UDP

F. None of the above

Answer: E

Explanation:

The 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). TFTP uses UDP port 69.

QUESTION 120

A Certkiller branch office uses Telnet and FTP to access an application at the main office over a T1 HDLC link. You wish to increase the performance over this link through the use of a compression algorithm. What compression type will provide the best performance improvement on this link?

A. Compressed Real-time Transport Protocol

B. TCP header compression

C. Stacker compression

- D. Predictor compression
- E. None of the above

Answer: C

Explanation:

You can configure point-to-point software compression for all LAPB, PPP, and HDLC encapsulations using either the predictor or stacker compression methods. Compression reduces the size of frames via lossless data compression. HDLC encapsulations support the Stacker compression algorithm. PPP and LAPB encapsulations support both predictor and Stacker compression algorithms.

When compression is performed in software installed in the router's main processor, it might significantly affect system performance.

Compression requires that both ends of the serial link be configured to use compression.

QUESTION 121

You are in charge of security management within the Certkiller service provider network. You have identified a Denial of Service (DOS) attack against an edge router connecting to another network. Which options would be most useful to deploying to protect the router?

- A. AutoSecure
- B. Netflow
- C. AAA
- D. IOS IPS
- E. CBAC (Context Based Access Control)
- F. CoPP (Control Plane Policing)
- G. None of the above

Answer: F

Explanation:

The Control Plane Policing feature allows users to configure a quality of service (QoS) filter that manages the traffic flow of control plane packets to protect the control plane of CiscoIOS routers and switches against reconnaissance and denial-of-service (DoS) attacks. In this way, the control plane (CP) can help maintain packet forwarding and protocol states despite an attack or heavy traffic load on the router or switch.

Configuring the Control Plane Policing feature on your Cisco router or switch provides the following benefits:

Protection against DoS attacks at infrastructure routers and switches

QoS control for packets that are destined to the control plane of Cisco routers or switches

Ease of configuration for control plane policies

Better platform reliability and availability

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1838/products_feature_guide09186a008052446b.html

QUESTION 122

What is the standard transport protocol and port used for SYSLOG messages that get sent from Cisco devices to the Certkiller management station?

- A. UDP 514
- B. TCP 520
- C. UDP 530
- D. TCP 540
- E. UDP 535
- F. None of the above

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

Reference:

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

QUESTION 123

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 124

Which types of SNMPv1 messages are sent from the NMS (Network Management Station) using SNMP version 1 to the Agent?

- A. Trap, Get and Set
- B. Get, Set and Getnext
- C. Get, Set, Getnext and GetBulk
- D. Get, Set and GetBulk
- E. Trap only
- F. None of the above

Answer: B

Explanation:

SNMP itself is a simple request/response protocol, and the SNMPv1 operations used by the NMS are defined as below.

Get: Allows the NMS to retrieve an object variable from the agent.

GetNext: Allows the NMS to retrieve the next object variable from a table or list within an agent. In SNMPv1, when a NMS wants to retrieve all elements of a table from an agent, it initiates a Get operation, followed by a series of GetNext operations.

Set: Allows the NMS to set values for object variables within an agent.

Incorrect Answers:

A, E. SNMP traps are used by the agent to inform the NMS of some events.

C, D. GetBulk is used in SNMPv2, not version 1. SNMPv2 defines two new operations: GetBulk and Inform. The GetBulk operation is used to efficiently retrieve large blocks of data. The Inform operation allows one NMS to send trap information to another NMS and to then receive a response. In SNMPv2, if the agent responding to GetBulk operations cannot provide values for all the variables in a list, it provides partial results.

QUESTION 125

What is the difference between the community formats of SNMPv1 and SNMPv2c?

- A. With SNMPv1, communities are sent as clear text and on SNMPv2c they are encrypted.
- B. On SNMPv1 communities are encrypted and on SNMPv2c they are sent as clear text.
- C. There is no difference because both versions send encrypted communities.
- D. There is no difference because both versions send communities as clear text.
- E. SNMPv2c does not use communities.
- F. None of the above

Answer: D

Explanation:

The original Internet standard Network Management Framework, described in RFCs 1155, 1157, and 1213, is called the SNMP version 1 (SNMPv1) framework. Relevant portions of the proposed framework for version 2C of the Simple Network Management Protocol (SNMPv2C) are described in RFCs 1901 through 1908.

SNMPv1 and SNMPv2c use a community string match for user authentication.

Community strings provided a weak form of access control in earlier versions of SNMP. SNMPv3 provides much improved access control using strong authentication and should be preferred over SNMPv1 and SNMPv2c wherever it is supported. Both versions send communities as clear text messages.

QUESTION 126

Network management tools use Management Information Base (MIB) information to monitor and manage networks. Which of the following is NOT part of the MIB-2 specification, as defined in RFC 1213? (Choose all that apply)

- A. The System Group
- B. The TCP Group
- C. The Transmission Group
- D. The Enterprises Group
- E. The RMON Group
- F. The ICMP Group

Answer: D, E

Explanation:

RFC 1213 defines the "Management Information Base for Network Management of TCP/IP-based internets: MIB-II" specification. It defines all of the following groups: System, Interfaces, Address Translation, IP, ICMP, TCP, UDP, EGP, Transmission, and SNMP. The RMON group is not part of RFC 1213, nor is the Enterprises Group

QUESTION 127

Which statements are true about the purpose and functionality between SNMP and MIBs? (Select three)

- A. A Management Information Base (MIB) is a collection of information that is organized hierarchically.
- B. A Management Information Base (MIB) is a collection of network device information that is organized in a bulk transfer mode to the management station.
- C. MIBs are accessed using a network-management protocol such as SNMP.
- D. MIBs are accessed using a network-management protocol such as TCP.
- E. MIBs are comprised of managed objects and are identified by the object identifiers.
- F. MIBs are comprised of managed objects and are identified by the lmhosts table.

Answer: A, C, E

Explanation:

The Cisco MIB variables are accessible via the Simple Network Management Protocol (SNMP), which is an application-layer protocol designed to facilitate the exchange of management information between network devices. The SNMP system consists of three parts: SNMP manager, SNMP agent, and MIB.

The MIB structure is logically represented by a tree hierarchy. The root of the tree is unnamed and splits into three main branches: Consultative Committee for International Telegraph and Telephone (CCITT), International Organization for Standardization (ISO), and joint ISO/CCITT.

Finally, each group of MIB variables is accompanied by an illustration that indicates the specific object identifier for each variable.

QUESTION 128

Which options are true regarding the privacy capability using cryptography and the authentication method for SNMPv1, SNMPv2c and SNMPv3? (Choose all that

apply)

- A. SNMPv1 has no privacy and uses community for authentication.
- B. SNMPv2c has privacy and uses community for authentication.
- C. SNMPv2c has privacy and uses usernames for authentication.
- D. SNMPv3 has privacy and use community for authentication.
- E. SNMPv3 has privacy and uses usernames for authentication.

Answer: A, E

Explanation:

SNMPv1 and SNMPv2 use the notion of communities to establish trust between managers and agents. An agent is configured with three community names: read-only, read-write, and trap. The community names are essentially passwords; there's no real difference between a community string and the password you use to access your account on the computer. The three community strings control different kinds of activities. As its name implies, the read-only community string lets you read data values, but doesn't let you modify the data. For example, it allows you to read the number of packets that have been transferred through the ports on your router, but doesn't let you reset the counters. The read-write community is allowed to read and modify data values; with the read-write community string, you can read the counters, reset their values, and even reset the interfaces or do other things that change the router's configuration. Finally, the trap community string allows you to receive traps (asynchronous notifications) from the agent.

SNMPv3 not only encrypts all transmissions but also enables the responder (usually an SNMP agent) to authenticate the user generating the request, guarantee the integrity of the message using a digital signature, and apply complex and granular access-control rules to each request. It also lets the administrator specify these levels of protection in varied combinations (unsecured, authenticated and authenticated with encryption). In addition, any number of access-control rules can be applied at the SNMP agent or manager. While this level of security was completely impractical in hardware 10 years ago, today's infrastructure devices have enough RAM and CPU cycles to support not only this advanced SNMP security but also full-fledged Web management services--all in firmware.

QUESTION 129

Which security features are defined in SNMPv3? (Select all that apply)

- A. Authentication
- B. Domain checking
- C. Accounting
- D. Privacy

Answer: A, D

Explanation:

SNMP Version 3 (SNMPv3) adds security and remote configuration capabilities to the previous versions. The SNMPv3 architecture introduces the User-based Security Model (USM) for message security and the View-based Access Control Model (VACM) for access control.

The principal security enhancements defined in SNMP version 3 is authentication, privacy, and access control.

Incorrect Answers:

B, C. SNMP version 3 provides no defines no mechanisms for checking the domain or accounting.

QUESTION 130

What SNMP message type reports events to the NMS reliably?

- A. Get
- B. Response
- C. Inform
- D. Trap
- E. Get Bulk
- F. None of the above

Answer: C


Explanation:

SNMP notifications can be sent as traps or inform requests. Traps are unreliable because the receiver does not send acknowledgments when it receives traps. The sender cannot determine if the traps were received. However, an SNMP entity that receives an inform request acknowledges the message with an SNMP response PDU. If the sender never receives the response, the inform request can be sent again. Thus, informs are more likely to reach their intended destination.

However, informs consume more resources in the agent and in the network. Unlike a trap, which is discarded as soon as it is sent, an inform request must be held in memory until a response is received or the request times out. Also, traps are sent only once, while an inform message may be retried several times. The retries increase traffic and contribute to a higher overhead on the network.

QUESTION 131

The Certkiller LAN is displayed below:



What should the Cisco IOS commands look like in the Certkiller 1 router to perform the

exhibit?

- A. logging source-interface fastethernet 0/0
logging 16.0.1.4
logging facility sys9
logging on
- B. logging 16.0.1.4
logging trap debugging
logging facility sys9
logging source-interface serial0
logging on
- C. logging 16.0.1.4
logging trap debugging
logging facility sys9
logging source-interface ethernet1
logging on
- D. logging 16.0.1.1
logging trap 7
logging source-interface serial 1
logging origin-id ip

Answer: C

Explanation:

In the example displayed above, the syslog server resides at 16.0.1.4 so we will want to send all SNMP traps to this IP address. In addition, the source interface information that should be sent to this server is the ethernet 1 interface, since this is the address used for all messages sent to the server.

Incorrect Answers:

- A. In this example the wrong interface source is used. In addition, the logging level information that should be sent to the server is not specified.
- B. Here the wrong interface is configured as the logging source
- D. This choice specified the wrong source interface, as well as the wrong IP address of the syslog server.

QUESTION 132

SNMP version 3 has been implemented throughout the Certkiller network. On SNMPv3 which message types are classified as Unconfirmed Class PDU? Select all that apply.

- A. Get
- B. Trap
- C. Inform
- D. Report
- E. Response

Answer: B, D, E

Explanation:

SNMP PDU Classes

SNMPv1 originally defined six PDUs. The number of PDUs was expanded and some changes made to their name and use in SNMPv2 and SNMPv3. The current SNMP Framework categorizes the PDUs into different classes. These classes describe both the function of each message type and the kind of communication they use to perform their task (polling versus interrupting).

Table 210 shows the main SNMPv2/SNMPv3 PDU classes, describes them, and shows which PDUs are in each class in SNMPv2/SNMPv3. These classes were not used in SNMPv1 but for clarity I also show which messages from SNMPv1 fall into the classes conceptually:

Table 210: SNMP PDU (Message) Classes			
SNMPv3 PDU Class	Description	SNMPv1 PDUs	SNMPv2/SNMPv3 PDUs
Read	Messages that read management information from a managed device using a polling mechanism.	GetRequest-PDU, GetNextRequest-PDU	GetRequest-PDU, GetNextRequest-PDU GetBulkRequest-PDU
Write	Messages that change management information on a managed device to affect the device's operation.	SetRequest-PDU	SetRequest-PDU
Response	Messages sent in response to a previous request.	GetResponse-PDU	Response-PDU
Notification	Messages used by a device to send an interrupt-like notification to an SNMP manager.	Trap-PDU	Trapv2-PDU, InformRequest-PDU

--	--	--	--

The GetBulkRequest-PDU and InformRequest-PDU messages are new in SNMPv2/v3. The GetResponse-PDU message was renamed just Response-PDU (since it is in fact a response and not a message that "gets" anything), and the new Trapv2-PDU replaces Trap-PDU.

There are three other "special" classes defined by the current SNMP Framework. The Internal class contains a special message called Report-PDU defined for internal SNMP communication. The SNMP standards also provide two classes called Confirmed and Unconfirmed, used to categorize the messages in my table above based on whether or not they are acknowledged. The Report-PDU, Trapv2-PDU, and Response-PDU messages are considered Unconfirmed and the rest are Confirmed.

Reference:

http://www.tcpipguide.com/free/t_SNMPProtocolGeneralOperationCommunicationMethodsan-2.htm

QUESTION 133

On what lower level transport protocol does SNMP rely and why?

- A. TCP, because SNMP requires the reliability of TCP, which ensures packets are transmitted reliably, in event that a packet is lost in the network.
- B. UDP, because SNMP is an application that does not require the reliability provided by TCP.
- C. IP, because SNMP requires the reliability of IP packets, which can detect lost packets and retransmit them if required.
- D. UDP, because SNMP is an application that requires the reliability of UDP and UDP's ability to detect lost packets and retransmit them.
- E. TCP, because SNMP is an application that does not require detection and retransmission of lost packets.

Answer: B

Explanation:

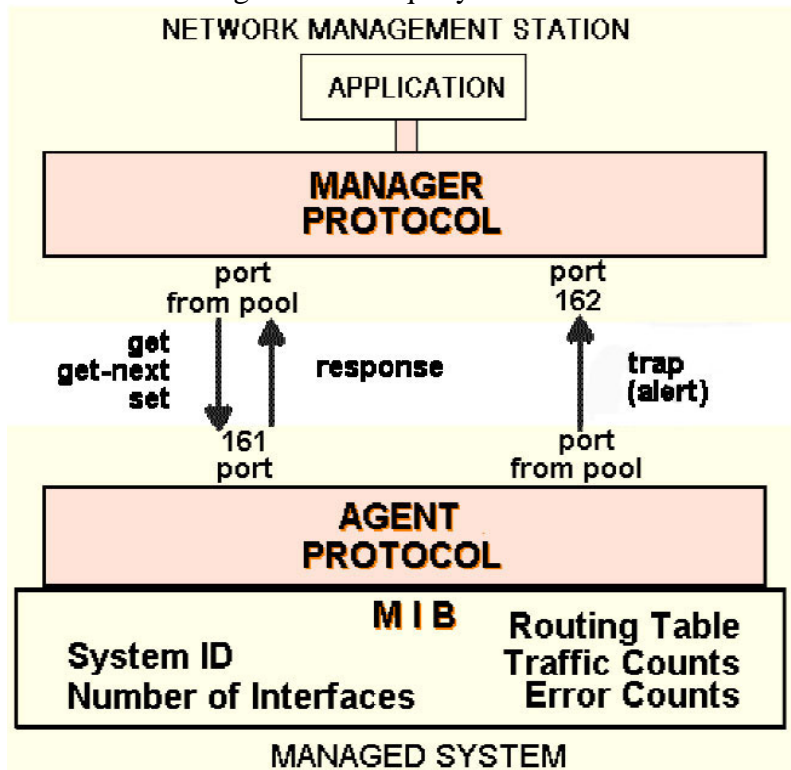
SNMP uses the User Datagram Protocol (UDP) as the transport protocol for passing data between managers and agents. UDP, defined in RFC 768, was chosen over the Transmission Control Protocol (TCP) because it is connectionless; that is, no end-to-end connection is made between the agent and the NMS when datagrams (packets) are sent back and forth. This aspect of UDP makes it unreliable, since there is no acknowledgment of lost datagrams at the protocol level. It's up to the SNMP application to determine if datagrams are lost and retransmit them if it so desires. This is typically accomplished with a simple timeout. The NMS sends a UDP request to an agent and waits for a response. The length of time the NMS waits depends on how it's configured. If the timeout is reached and the NMS has not heard back from the agent, it assumes the packet was lost and retransmits the request. The number of times the NMS retransmits packets is also configurable.

At least as far as regular information requests are concerned, the unreliable nature of UDP isn't a real problem. At worst, the management station issues a request and never receives a response. For traps, the situation is somewhat different. If an agent sends a trap

and the trap never arrives, the NMS has no way of knowing that it was ever sent. The agent doesn't even know that it needs to resend the trap, because the NMS is not required to send a response back to the agent acknowledging receipt of the trap.

The upside to the unreliable nature of UDP is that it requires low overhead, so the impact on your network's performance is reduced. SNMP has been implemented over TCP, but this is more for special-case situations in which someone is developing an agent for a proprietary piece of equipment. In a heavily congested and managed network, SNMP over TCP is a bad idea. It's also worth realizing that TCP isn't magic, and that SNMP is designed for working with networks that are in trouble -- if your network never failed, you wouldn't need to monitor it. When a network is failing, a protocol that tries to get the data through but gives up if it can't is almost certainly a better design choice than a protocol that will flood the network with retransmissions in its attempt to achieve reliability.

SNMP uses the UDP port 161 for sending and receiving requests, and port 162 for receiving traps from managed devices. Every device that implements SNMP must use these port numbers as the defaults, but some vendors allow you to change the default ports in the agent's configuration. If these defaults are changed, the NMS must be made aware of the changes so it can query the device on the correct ports.



SNMP use of UDP port numbers.

QUESTION 134

While auditing the Certkiller network you come across the following messages:

```
event manager applet IPSNMPWD
 event ioswdsysmon sub1 cpu-proc taskname "IP SNMP" op ge val 50
 period 10
 action 1.0 publish-event sub-system 798 type 1 arg1 "IP SNMP" arg2
 "50"
 action 2.0 snmp-trap intdata1 50 strdata "IP SNMP CPU error"
 snmp-server enable traps event-manager
```

Based on the information shown above, what three statements are true? (Choose three)

- A. Event Manager applet is configured incorrectly
- B. The IP SNMP process is monitored every 10 seconds
- C. The IP SNMP process is monitored every 10 minutes
- D. When cpu-process is greater than 50% an event is generated
- E. SNMP Trap type 50 messages are sent to the event managers
- F. Publish event to well known user 798 with an SNMP trap message

Answer: B, D, F

Explanation:

B, D: Example:

The following example shows how to configure three EEM applets to demonstrate how the Cisco IOS watchdog system monitor (IOSWDSysMon) event detector works:

Watchdog System Monitor Sample1 Policy

The first policy triggers an applet when the average CPU usage for the process named IP Input is greater than or equal to 1 percent for 10 seconds:

```
event manager applet IOSWD_Sample1
```

```
event ioswdsysmon sub1 cpu-proc taskname "IP Input" op ge val 1 period 10
```

```
action 1.0 syslog msg "IOSWD_Sample1 Policy Triggered"
```

```
F: action publish-event
```

To specify the action of publishing an application-specific event when the event specified for an Embedded Event Manager (EEM) applet is triggered, use the action publish-event command in applet configuration mode. To remove the action of publishing an application-specific event, use the no form of this command.

action label publish-event sub-system sub-system-id type event-type arg1 argument-data [arg2 argument-data] [arg3 argument-data] [arg4 argument-data]

no action label publish-event

Syntax Description

label ✖	Unique identifier that can be any string value. Actions are sorted and run in ascending alphanumeric key sequence using the label as the sort key. If the string contains embedded blanks, enclose it in double quotation marks. ✖
sub-system ✖	Specifies an identifier for the subsystem named in the sub-system-id argument that will publish the application event. ¶ <ul style="list-style-type: none">sub-system-id—Identifier of the subsystem. Number in the range from 1 to 4294967295. If the event is to be published by an EEM policy, the sub-system-id reserved for a customer policy is 798. ✖

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps5207/products_command_reference_chapter09186a00801a

QUESTION 135

Part of the configuration of router CK1 is shown below:

```
aaa new-model
username myname password abc123
aaa authentication ppp access-dot1x local
aaa authentication login access1 local
aaa authentication dot1x default radius
dot1x system-auth-control
tacacs-server host 192.168.1.15 key quert123
radius server host 192.168.2.27 key poiuy098
!
interface fastethernet 5/1
  dot1x port-control auto
```

Based on the information shown above, what is the effect of the configuration on Certkiller users attempting to access the LAN attached to interface FE 5/1 on CK1 ?

- A. They will be authenticated via ppp using the local database.
- B. They will be authenticated via ppp using the server at IP address 192.168.1.15.
- C. They will be authenticated via ppp using the server at IP address 192.168.2.27
- D. They will be authenticated via 802.1x using the local database.
- E. They will be authenticated via 802.1x using the server at IP address 192.168.1.150.
- F. They will be authenticated via 802.1x using the server at IP address 192.168.2.27.

Answer: F

Explanation:

When you enable 802.1X port-based authentication, note the following syntax information:

To create a default list that is used when a named list is not specified in the authentication command, use the default keyword followed by the methods that are to be used in default situations. The default method list is automatically applied to all interfaces.

Enter at least one of these keywords:

group radius-Use the list of all RADIUS servers for authentication.

none-Use no authentication. The client is automatically authenticated by the switch without using the information supplied by the client.

This example shows how to enable AAA and 802.1X on Fast Ethernet port 5/1:

Router# configure terminal

CK1 (config)# aaa new-model

CK1 (config)# aaa authentication dot1x default group radius

CK1 (config)# dot1x system-auth-control

CK1 (config)# interface fastethernet 5/1

CK1 (config-if)# dot1x port-control auto

CK1 (config-if)# end

In this example, the default 802.1x authentication method is configured to be RADIUS, and the RADIUS server is located at IP address 192.168.2.27.

QUESTION 136

Part of the configuration for router CK1 is shown below:

```
enable secret 483924
!
aaa new-model
username myname password Certkiller24
aaa authentication login default enable
aaa authentication login access1 local
aaa authentication login access2 radius tacacs+
aaa authentication login access3 tacacs+ local
tacacs-server host 192.168.1.15 key qwert123
radius-server host 192.168.2.27 key poiuy098
!
Line console 0
  Login authentication access3

line vty 0 4
  password certkiller1
  login
```

Based on the configuration output shown above, what method is being used to secure the console port of CK1 ?

- A. Authentication is being done using the local database.
- B. Authentication is being done using the login password Certkiller 1.
- C. Authentication is being done using the enable password as a default
- D. Authentication is being done using the server at IP address 192.168.1.15. If a connection to that server fails, the local database will be used.
- E. Authentication is being done using the server at IP address 192.168.2.27
- F. None of the above

Answer: D

Explanation:

The router is using the keyword access3 for authentication for the console port. Access3 points to two different methods for authentication; the first is TACACS+ which is located at 192.168.1.15. If the authentication connection to the server fails, then the local database will be used as a backup.

Incorrect Answers:

- A. Based on the configuration file above, TACAS+ is the primary authentication method and the local database is to only be used as a backup method.
- B. This is the password that is to be used for Telnet access, not the console password.
- C. The enable password is not used, since the login authentication information is taken

from the "access3" keyword.

E. This is the IP address of the RADIUS server, not the TACACS+ server.

QUESTION 137

You need to enhance the security of the Certkiller network. Which configuration 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

E. None of the above

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

no snmp-server user

Syntax Description:

username	The name of the user on the host that connects to the agent.
groupname	(Optional) The name of the group to which the user is associated.
remote	(Optional) Specifies the remote copy of SNMP on the router.
ip-address	(Optional) The IP address of the device that contains the remote copy of SNMP.
udp-port	(Optional) Specifies a UDP port of the host to use.

port	(Optional) A UDP port number that the host uses. The default is 162.
v1	(Optional) The least secure of the possible security models.
v2c	(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.
v3	(Optional) The most secure of the possible security models.
encrypted	(Optional) Specifies whether the password appears in encrypted format (a series of digits, masking the true characters of the string).
auth	(Optional) Initiates an authentication level setting session.
md5	(Optional) The HMAC-MD5-96 authentication level.
sha	(Optional) The HMAC-SHA-96 authentication level.
auth-password	(Optional) A string (not to exceed 64 characters) that enables the agent to receive packets from the host.
priv	(Optional) The option that initiates a privacy authentication level setting session.
des56	(Optional) The CBC-DES privacy authentication algorithm.
priv password	(Optional) A string (not to exceed 64

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 138

Router CK1 is configured for OSPF. Under the OSPF process, you type in the "area 1 range" command. Which LSA types will be acted upon (summarized) as a result? (Choose all that apply)

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

E. Type 5

Answer: A, B

Explanation:

Area range command is used for summarizing routes on the boundary of two OSPF areas.

The information to be summarized is contained in two types of LSAs: Type 1 and Type2.

Type 1 LSAs are Router LSAs and are generated by each router in an OSPF network.

Type 2 LSAs are network LSAs, which are generated by the DR.

Both Type1 and Type 2 LSAs are flooded within the originating area only. Only when the information needs to be conveyed to another area in a summarized form area-range command is used, which acts on the information provided by these two LSAs.

Reference:

CCIE Professional Development Routing TCP/IP Volume I by Jeff Doyle page 471.

Incorrect Answers:

C. Type 3 LSA are the result of type 1 and type 2 summaries that are created by the area range command.

D. Type 4 LSAs are ASBR summary LSAs

E. Type 5 LSAs are AS External LSAs

QUESTION 139

A change in the topology of the Certkiller OSPF network causes the flooding operation. Which OSPF packet types are used in this LSA Flooding?

A. Hello

B. Link State Update

C. Link State Request

D. Database description

E. Link State Acknowledgement

Answer: B, E

Explanation:

A change in the OSPF network topology is represented as a change in one or more of the OSPF Link State Advertisements (LSAs). Flooding is the process by which changed or new LSAs are sent throughout the network, and are used to ensure that the database of every OSPF router is updated and an identical database is maintained. This flooding makes use of two OSPF packet types: Link State Update packets (type 4), and Link State Acknowledgement packets (type 5).

Reference: Jeff Doyle, "Routing TCP/IP volume 1" page 451.

QUESTION 140

Router CK1 is configured for OSPF and is connected to two areas: area 0 and area 1. You then configure area 1 as a stub area. Which LSAs will now operate inside of area 1?

A. Type 7

- B. Type 1 and 2
- C. Type 1, 2, and 5
- D. Type 3 and 4
- E. Type 1, 2 and 3

Answer: E

Explanation:

Only type 1, 2, and 3 LSAs will be allowed inside of a stub area.

Incorrect Answers:

A. Type 7 LSAs are used for NSSA, not stubby areas.

B. Network Summary LSAs (Type 3) are also allowed.

Reference:

CCIE Professional Development Routing TCP/IP Volume I by Jeff Doyle page 479.

QUESTION 141

Study the Exhibits below carefully:

The following exhibit is an illustration of the output from an ASBR:

ASBBR#show ip ospf database external

OSPF Router with ID (15.33.4.2) (Process ID 10)

Type-5 AS External Link States

LS age: 15

Options: (No TOS-capability, DC)

LS Type: AS External Link

Link State ID: 10.10.1.0 (External Network Number)

Advertising Router: 15.33.4.2

LS Seq Number: 80000002

Checksum: 0x513

Length: 36

Network Mask: /24

Metric Type: 1 (Comparable directly to link state metric)

TOS: 0

Metric: 10

Forward Address: 0.0.0.0

External Route Tag: 0

And this exhibit is an illustration from a router in the network:

Router CK1 #show ip ospf border-routers

OSPF Process 10 internal Routing Table

Codes: i-intra-area route, I-Inter-area route

15.33.4.2(2) via 30.0.0.1, Serial0/0, ASBR, Area0, SPF 4

Based on this information what is the total metric for the route to subnet 10.10.1.0/24 on Router CK1 ?

- A. 1
- B. 8

- C. 12
- D. 20
- E. 22

Answer: C

Explanation:

The metric of the external link shows 10. Then we need to add 2 from the inter-area metric, for a total of 12.

QUESTION 142

In your OSPF network serial 0 on your router, CK1 , is in area 1. Later, you configure serial 0 as passive. What is the effect of this configuration change?

- A. OSPF will accept the routing updates from neighbors.
- B. OSPF will form all the available adjacencies out of that interface.
- C. OSPF will not insert any of the learned routes in the local routing table.
- D. OSPF will not form any adjacency out of that interface.
- E. None of the above.

Answer: D

Explanation:

With passive-interface, an adjacency will never occur out of that interface, as no hello packets are exchanged out of a passive interface.

Incorrect Answers:

A. Normally, defining an interface as passive will accomplish this. No routes will be sent out, but routes can still be received. OSPF differs because link state protocols need information for the entire network topology. Defining an interface as passive with OSPF means that the adjacency will not be established, therefore, no routes will be able to be received on that interface.

QUESTION 143

You are the network administrator at Certkiller . The Certkiller network contains four Routers named CK1 , CK2 , CK3 , and CK4 . All four routers are connected to a hub via Ethernet interfaces. All four routers have a basic OSPF configuration of a network statement for the Ethernet network. During routine maintenance, you issue the show ip ospf neighbor command on Router CK2 . The output from the show ip ospf neighbor command shows 2WAY/DROTHER for its neighbor, Router CK3 . What conclusions can you draw from this output? (Choose all that apply)

- A. Router CK2 is the DR or BDR.
- B. Router CK3 is not a DR or BDR.
- C. Router CK2 - Router CK3 adjacency is not yet FULL.
- D. Router CK2 is not the DR.

E. Router CK4 is the DR.

Answer: B, D

Explanation:

OSPF routers can have one of three neighbor relationships: Designated Router (DR), Backup Designated Router (BDR), or neither. For neither, the router neighbor relationship will show as 2WAY/DROTHER.


Incorrect Answers:

A, C. 2WAY/DROTHER means that the two routers are neither the DR nor the BDR.

E. Either Router CK1 or Router CK4 is the DR. Based on the information that is provided we cannot be sure which one it is.

QUESTION 144

The following exhibit displays the Certkiller OSPF network:



Router CK2 needs to send a string of packets to router CK4 . How will router CK2 decide the path to take to reach CK4 ?

- A. CK2 will select a path after considering the costs inside Area 1 only.
- B. CK2 will alternate between Router CK1 and Router CK3 if the costs are equal.
- C. CK 2 will always go through Router CK1 with no regard for costs.
- D. CK2 will select a path after considering the costs inside both Area 0 and Area 1.
- E. None of the above.

Answer: A

Explanation:

OSPF prefers Intra Area Path over Inter Area Paths.

Incorrect Answers:

B.

The Answer B is incorrect because OSPF does not conduct ECMP load balancing on multiple paths with equal cost if the respective paths span through more than one area. B is incorrect for several reasons. If a packet has to alternate between two paths that means Per Packet load balancing is in effect. Which is normally in place for links less than 56k. For higher link speeds fast switching (default switching mode) is enabled. In this mode all packets to one destination in a target subnet are sent over one path, since route lookup is not performed for every packet, it is rather performed per flow. So B is totally ruled out.

C. Even though router CK1 is the most direct way to reach area 0, OSPF will always prefer to stay in the same area over traversing multiple areas.

D. OSPF prefers Intra area paths, so only the costs associated with reaching CK4 via area 1 will be considered first.

Reference:

<http://www.riverstonenet.com/support/ospf/interface-costs.htm>

QUESTION 145

Router CK1 is configured for OSPF. Interface serial 0 is configured to be in area 0 and interface serial 1 is configured to be in area 1. Under the OSPF process "area 1 nssa default-information-originate" is configured. Which of the following are true? (Choose all that apply)

- A. CK1 will inject a type 3 default route into area 1.
- B. CK1 will inject a type 7 default route into area 1.
- C. CK1 will inject a type 7 default route into area 0.
- D. CK1 needs a default route in its routing table to inject a default into area 1.
- E. CK1 does not need a default route in its routing table to inject a default into area 1.

Answer: B, D

Explanation:

Type 7 routes are injected into OSPF NSSA areas, and the default information originate command will make CK1 inject type 7 default routes into area 1. As a rule, an OSPF router will need a default route itself before injecting a default route into an area, unless the keyword "always" is used in the configuration. For example, "default-information originate always."

Incorrect Answers:

- A. In a NSSA area, the NSSA area generates the default route with the "default-information originate" command, but unlike other default routes that use type 3 information, NSSA default routes use type 7.
- C. CK1 will inject a type 7 NSSA route into area 1, not area 0. Area 0 can not be an NSSA.
- E. Using the command shown in the question above will not create the route, because a previous default route did not already exist within the routing table. A default route would have been injected only if the keyword "always" was inserted.

Reference:

www.cisco.com/en/US/tech/ CK3 65/technologies_tech_note09186a0080094a74.shtml

QUESTION 146

Which of the following OSPF routers can generate a type 4 ASBR-summary LSA on the Certkiller network? (Choose all that apply)

- A. ABRs
- B. DR
- C. BDR
- D. ASBRs
- E. None of the above

Answer: A

Explanation:

Type 4 LSAs are only put out by ABRs and only in two cases: 1. There is an ASBR that the ABR needs to tell the backbone area about. 2. There is a legacy router that is incapable of demand circuits. These last two are indication LSAs and are put out only by an ABR putting itself in the ASBR position, but it is still not an ASBR. An ASBR would not be responsible for reporting either of these situations.

QUESTION 147

Routers CK1 and CK2 are in the same LAN and both are running OSPF. Which multicast IP address will CK1 and CK2 use for sending routing updates to each other? (Choose all that apply)

- A. 224.0.0.10
- B. 224.0.0.1
- C. 224.0.0.13
- D. 224.0.0.5
- E. 224.0.0.9
- F. 224.0.0.6

Answer: D, F

Explanation:


224.0.0.5 is the all-OSPF routers multicast and 224.0.0.6 is the Designated Routers multicast address.

Incorrect Answers:

- A. 224.0.0.10 is used for IGRP.
 - B. 224.0.0.1 is reserved for all systems on the subnet.
 - D. 224.0.0.13 is used by PIM.
 - E. 224.0.0.9 is reserved for RIP version 2 announcements.
-

QUESTION 148

The Certkiller WAN utilizes OSPF as shown below. The OSPF metric for each link is displayed in the diagram as follows:



What is the OSPF shortest path from Network 2 to Network 1 with the OSPF link costs shown in the exhibit?

- A. H G D B A
- B. H G E C B A
- C. H G F D B A
- D. H G E D B A

Answer: B

Explanation:

Cost of links from Network 2 to Network 1 is:

- A. H G D B A = 1+1 +5 +48 +10 = 65
- B. H G E C B A = 1+1+1+30+1+10= 44
- C. H G F D B A = 1+1+10+48+48+10= 118
- D. H G E D B A = 1+1+1+10+48+10= 71

Therefore, the shortest path is the lowest cost path which is option B. It is important to remember that OSPF uses the total cost of the metrics from a source to a given destination, and the number of hop counts is irrelevant.

QUESTION 149

The Certkiller router CK2 is experiencing OSPF problems with a neighbor across a frame relay network. During troubleshooting, OSPF event debugging was issued as shown below:

```
CK2 #debug ip ospf events
OSPF events debugging is on
```


CK2 #

00:16:22: OSPF: Rcd hello from 192.168.0.6 area 4 from Ethernet 0/0 16.16.26.6

00:16:22: OSPF: End of hello processing

00:16:22: OSPF: Send hello to 244.0.0.5 area 4 on Ethernet0/0 from 116.16.26.2

CK2 #

00:16:28: OSPF: Rcd hello from 192.168.0.3 area 3 from Serial1/0 116.16.32.1

00:16:28: OSPF: Mismatched hello parameters from 116.16.32.1

00:16:28: OSPF: Dead R 40 C 120, Hello R 10 C 30 Mask R 255.255.255.252 C 255.255.255.252

CK2 :

00:16:32: OSPF: Rcd hello from 192.168.0.6 area 4 from Ethernet0/0 116.16.26.6

00:16:32: OSPF: End of hello processing

00:16:32: OSPF: Send hello to 224.0.0.5 area 4 on Ethernet0/0 from 116.16.26.2

Based on the information above, what is the most likely reason for the OSPF problems across the frame relay link?

- A. This router is in area 4 while its neighbor is configured to be in area 3.
- B. There is mismatch between the OSPF frame-relay parameters configured on this router and those configured on its neighbor.
- C. The OSPF network mode configured on this router is not the same as the mode configured on its neighbor.
- D. This router has a frame-relay interface DLCI statement that is using the broadcast mode. While its neighbor is using a point-to-point mode.
- E. None of the above.

Answer: C

Explanation:

The default timers for a broadcast network (LAN) are: Hello 10 seconds, Dead 40 seconds

The default timers for an NBMA network (Frame Relay) are: Hello 30 seconds, Dead 120 seconds.

The problem above shows that these timers do not match at each end. The "Dead R 40 C 120, Hello R 10 C 30" means that the configured Dead time is 120 seconds locally on this router, but the received update shows it is configured to be 40 seconds. Similarly, the received hello packet shows that it has a hello time of 10 seconds, where router CK2 is configured for 30 seconds. Although the remote router may have had their timers changed manually within the OSPF process, the most likely cause of the problem is that router CK2 is configured with a network type of NBMA and the other router is configured with a network type of broadcast.

Incorrect Answers:

A. It is common for a router with multiple interfaces to be in different OSPF areas. Each network link must be in the same area, but each router can have multiple interfaces, that each belongs to a different area.

B. The Frame relay parameters do not appear to be misconfigured, just the OSPF timer values.

D. The timers on the local router, CK2, is 30 seconds for the Hello and 120 seconds for the dead, so this router is configured with a NBMA or pt-pt type, while the remote router is using a broadcast network type.

QUESTION 150

Which of the following statements are true regarding the SPF calculation? (Select three)

A. The Dijkstra algorithm is run two times.

B. The previous routing table is saved.

C. The present routing table is invalidated.

D. A router calculates the shortest-path cost using their neighbor(s) as the root for the SPF tree.

E. Cisco routers use a default OSPF cost of $10^7/BW$.

Answer: A, B, C

Explanation:

The Dijkstra algorithm code itself is run two times. The first time deals with routers and the second time always deals with networks.

When the Shortest Path First (SPF) algorithm is computed by an OSPF router, the previous routing table is save before the calculation and used in case any problems arise with the new one. It then invalidates the present routing table and performs the calculation using ITSELF as root in the SPF tree.

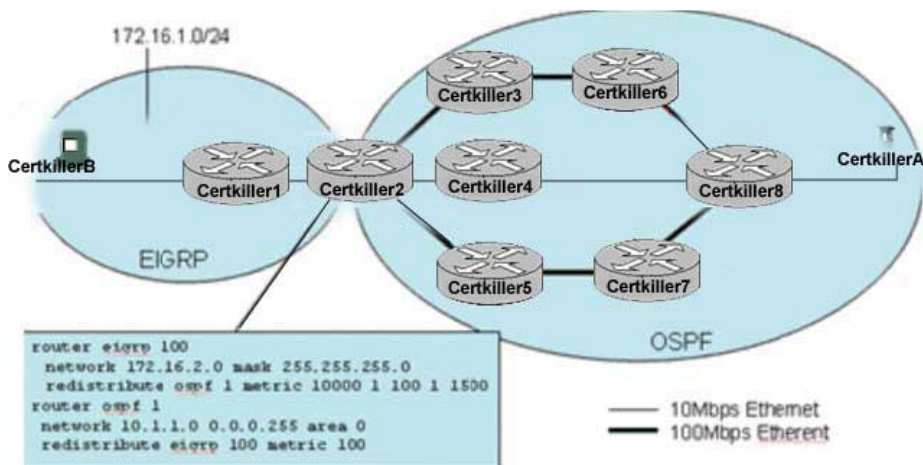
Incorrect Answers:

D. The router itself is the root, not the neighbor. A router periodically advertises its status or link state to its adjacencies. Link state advertisements flood throughout an area ensuring that all routers have exactly the same topological database. This database is a collection of the link state advertisements received from each router belonging to an area. From the information in this database, each router can calculate a shortest path tree with itself designated as the root of the tree.

E. The default OSPF cost of any link is $10^8/Bandwidth$, or $100,000,000/BW$.

QUESTION 151

The Certkiller network is displayed below:



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

- A. Traffic will only flow on the shortest, low-speed path, PC- Certkiller A - Certkiller 8 - Certkiller 4 - Certkiller 2- Certkiller 1 - PC- Certkiller B.
- B. Traffic will flow on both the high speed paths (PC- Certkiller A - Certkiller 8 - Certkiller 6- Certkiller 3 - Certkiller 2 - Certkiller 1 - PC- Certkiller B and PC- Certkiller A - Certkiller 8 - Certkiller 7 - Certkiller 5 - Certkiller 2 - Certkiller 1 - PC- Certkiller 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 Certkiller A to PC- Certkiller B, and uni-directionally on one of the high speed paths from PC- Certkiller B o PC- Certkiller 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:

OSPF uses the bandwidth of the links for the metric, and by default the 100 Mbps links will have an OSPF metric of 1 while the low speed links will have a metric of 10 so only the high speed Ethernet links will be used.

By default, OSPF load balances on up to four equal cost paths. Since both high speed paths will have a metric of 3 (1+1+1) from router Certkiller 8 to Certkiller 2 they traffic will load balance over the two paths.

QUESTION 152

What statement is correct regarding OSPF adjacencies and link-state database synchronization?

- A. Full adjacency occurs when OSPF routers reach the LOADING state.
- B. Adjacency relationship begins in the EXSTART state.
- C. All OSPF neighbors establish adjacencies in the FULL state with all other routers on the broadcast network.

D. The INIT state indicates that a router has received a Hello packet from a neighbor and has seen their own ROUTERID in the Hello packet.

Answer: B

Explanation:

The various states in which a neighbor can be are discussed below.

1. Down - the initial state of a neighbor conversation.
2. Attempt - indicates that an attempt should be made to contact the neighbor.
3. Init - hello packet has been received from the neighbor.
4. 2-Way - communication between two routers is bi-directional.
5. ExStart - first step to creating an adjacency between the two neighboring routers.
6. Exchange - the router is sending data description packets to the neighbor.
7. Loading - Link state request packets are sent to the neighbor.
8. Full - the neighboring routers are fully adjacent.

Incorrect Answers:

A. Full adjacency only occurs after the OSPF router has reached a FULL state.

C. In a broadcast network, all routers only become adjacent with the Designated Router (DR).

D. This state specifies that the router has received a hello packet from its neighbor, but the receiving router's ID was not included in the hello packet. When a router receives a hello packet from a neighbor, it should list the sender's router ID in its hello packet as an acknowledgment that it received a valid hello packet.

QUESTION 153

OSPF is running on the Certkiller network. In OSPF, what LSA type would only cause a partial SPF calculation?

- A. Type 1
- B. Type 2
- C. Type 4
- D. Type 7
- E. Type 9

Answer: D

Explanation:

OSPF Type 7 LSA's are reserved for Not So Stubby Areas (NSSA). This area accepts Type 7 LSAs which are external route advertisements like Type 5s but they are only flooded within the NSS

A. This is usually used when connecting to a branch office running an IGP. Normally this would have to be a standard area since a stub area would not import the external routes. If it was a standard area linking the ISP to the branch office then the ISP would receive all the Type 5 LSAs from the branch which it does not want. Because Type 7 LSAs are only flooded to the NSSA the ISP is saved from the

external routes whereas the NSSA can still receive them. Therefore, when this LSA is generated, only a partial SPF calculation needs to be performed.

QUESTION 154

OSPF is being used as the routing protocol in the Certkiller network. Which two statements regarding the SPF calculation on these OSPF routers are true? (Select two)

- A. The existing routing table is saved so that changes in routing table entries can be identified.
- B. The present routing table is invalidated and is built again from scratch.
- C. The Cisco router calculates the shortest-path cost using their neighbor(s) as the root for the SPF tree.
- D. Cisco routers use a default OSPF cost of $10^7/BW$.

Answer: A, B

Explanation:


When an OSPF router performs a new SPF calculation, the existing routing table is saved and used as a baseline for changes made to the network topology. When any SPF calculation is made, the OSPF neighbor or neighbors is used as the root of the SPF routing tree.

Incorrect Answers:

- C. The root of the SPF tree is always the router itself, not the neighboring router.
- D. The default cost for all OSPF links is $10^8/BW$, or 100,000,000/ configured bandwidth.

QUESTION 155

The Certkiller OSPF/RIPv2 network is displayed below:



Area 1 is an OSPF Not So Stubby Area (NSSA). What type of LSA will Certkiller 3 send out area 0 to indicate the presence of an ASBR in Area 1?

- A. A type 5 because P-bit has been set in the type 4 LSA that was sent from Certkiller 1 to Certkiller 3.
- B. A type 4 because the E-bit was set in the type 7 LSA that was sent from Certkiller 1 to Certkiller 3.
- C. A type 1 because the B-bit was set in the LSA that was propagated from Certkiller 1 to Certkiller 3.
- D. A type 3 because the E-bit was set in the type 1 LSA that was sent from Certkiller 1 to Certkiller 3.

Answer: B

Explanation:

In this case a type 5 LSA would be sent by the ASBR, which is Certkiller 1.

Type 5 Link State advertisements are generated by the ASBR and describe links external to the Autonomous System (AS). This LSAS is flooded to all areas except stub areas.

Here Certkiller 1 is considered to be an ASBR since it is directly connected with the RIP version 2 network.

The E-bit reflects the associated area's External Routing Capability. AS external link advertisements are not flooded into/through OSPF stub areas. The E-bit ensures that all members of a stub area agree on that area's configuration.

LSA type 3 and 4 are summary link advertisements generated by ABRs

describing inter-area routes. Type 3 describes routes to

networks and is used for summarization. Type 4 describes

routes to the ASBR. Since Certkiller 3 needs to advertise the presence of an ASBR, it will send out a type 4 LSA to area 0.

Reference:<http://www.cisco.com/warp/public/104/ospfdb6.html>

QUESTION 156

What statement is accurate regarding OSPF areas?

A. Redistribution is allowed into all types of OSPF areas.

B. When routes are redistributed into an OSPF stub area, they enter as type-5 LSAs.

C. Redistribution is allowed into an OSPF stub area, but not into an OSPF not-so-stubby area.

D. When routes are redistributed into an OSPF not-so-stubby area, they enter as type-5 LSAs.

E. When routes are redistributed into an OSPF not-so-stubby area, they enter as type-7 LSAs.

Answer: E

Explanation:

When routes are redistributed into OSPF, these routes are considered to be external routes. External LSAs are type 5 LSAs. Not so stubby areas allow external routes to be advertised into OSPF network while retaining the characteristics of a stub area. To do this, the ASBR (the one doing the redistributing) in the NSSA will originate a type 7 LSA to advertise the external destinations.

Reference: Jeff Doyle, Routing TCP/IP volume 1, page 483.

Incorrect Answers:

A. Type 5 LSAs are only allowed into Backbone (area 0) and non backbone, non-stub areas.

B. Type 5 LSAs (external LSAs) are not allowed into stub or totally stub areas.

C. The opposite is true. External LSAs are allowed into NSSA, but not stub areas.

D. Type 5 LSAs are not inserted into NSSA for external routes. Type 7 LSAs are created for this purpose.

QUESTION 157

In order to increase network availability in the Certkiller OSPF network, graceful restart was implemented. When using OSPF Graceful Restart, which mechanism is used to continue forwarding packets during a switchover?

- A. Layer 2 forwarding
- B. Hardware Based Forwarding
- C. Reverse Path Forwarding
- D. UDP Forwarding
- E. Forwarding Address
- F. None of the above

Answer: B

Explanation:

Nonstop Forwarding (NSF) for OSPFv2 in Cisco IOS software uses the IETF standardized graceful restart functionality that is described in RFC 3623. Under very specific situations, a router may undergo certain well-known failure conditions that should not affect packet forwarding across the switching platform. NSF capability allows for the forwarding of data packets to continue along routes that are already known, while the routing protocol information is being restored. This capability is useful in cases in which there is a component failure (for example, a Route Processor [RP] crash with a backup RP taking over) or in which there is a scheduled hitless software upgrade. A key element of NSF is packet forwarding. The OSPF protocol depends on CEF to continue forwarding packets during switchover while the routing protocols rebuild the Routing Information Base (RIB) tables. Once OSPF has converged, CEF updates the Forwarding Information Base (FIB) table and removes stale route entries. CEF then updates the line cards with the new FIB information. CEF maintains the FIB and uses the FIB information that was current at the time of a switchover to continue forwarding packets during the switchover. This feature reduces traffic interruption during the switchover.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a00805e8fbd.html

QUESTION 158

Within the Certkiller OSPF network, which statement is true regarding the LSA's contained in the link state database? (Choose all that apply).

- A. The LSRefreshTime is 30 minutes.
- B. LSA's can only be reflooded by the router that originated the LSA.
- C. When an LSA reaches its MaxAge the router will send out a purge message to the other routers within its area.
- D. All LSAs contained in the LSDB expire at the same time unless they are refreshed.

- E. The MaxAge of an LSA is 3600 seconds.
- F. None of the above

Answer: A, E

Explanation:

Each OSPF LSA has an age, which indicates whether the LSA is still valid. Once the LSA reaches the maximum age (one hour), it is discarded. During the aging process, the originating router sends a refresh packet every 30 minutes to refresh the LS

A. Refresh

packets are sent to keep the LSA from expiring, whether there has been a change in the network topology or not. Checksumming is performed on all LSAs every 10 minutes. The router keeps track of LSAs it generates and LSAs it receives from other routers. The router refreshes LSAs it generated; it ages the LSAs it received from other routers.

Incorrect Answers:

B. Each LSA gets refreshed when it is 30minutes old, independent of the other LSAs used by other OSPF routers.

C. Purge messages are not sent to neighboring routers since each router uses its own timers.

D. Global synchronization can be problematic in OSPF networks. This problem is solved by each LSA having its own timer. Each LSA gets refreshed when it is 30minutes old, independent of other LSAs, so the CPU is used only when necessary.

QUESTION 159

To display the routing table of router CK1 , the "show ip route" command was issued. Router CK1 is running OSPF. Which one of the following statements is correct regarding the OSPF information in a routing table?

- A. A routing designated with only an "O" represents either a type-1 or type-2 LSA.
- B. A route that has been redistributed into OSPF can only be represented with either an "E1" or "E2" designation.
- C. Routes that are within an area (intra-area) are marked with an "IA" in the routing table.
- D. Type-7 LSAs display routes redistributed into OSPF from another process, and thus are shown with either an "E1" or "E2" marking.
- E. All LSA types have unique designations in the IP routing table.
- F. None of the above.

Answer: B

Explanation:

The following OSPF codes are used in the IP routing tables of OSPF routers:

O - OSPF

IA - OSPF inter area

N1 - OSPF NSSA external type 1

N2 - OSPF NSSA external type 2

E1 - OSPF external type 1

E2 - OSPF external type 2

When redistributing routes into OSPF, the routes are considered to be learned from an external means, so they will always display as external type 1 or external type 2 routes. These routes will appear as E1 or E2 in the routing table.

Incorrect Answers:

A. A route designated with an "O" only represents a generic OSPF learned route.

C. Routes displayed as "IA" are inter-area, not intra-area.

D. Type 7 routes are displayed as normal OSPF routes.

E. The designations in the routing table are not based on LSA types. They are based only on the type of OSPF route.

QUESTION 160

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

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 161

The Certkiller OSPF topology is shown below:



Based on the metrics shown in the diagram above, what is the OSPF shortest path from Network 1 to Network 2?

- A. Certkiller 1 Certkiller 2 Certkiller 4 Certkiller 5 Certkiller 7 Certkiller 8
- B. Certkiller 1 Certkiller 2 Certkiller 3 Certkiller 5 Certkiller 7 Certkiller 8
- C. Certkiller 1 Certkiller 2 Certkiller 4 Certkiller 5 Certkiller 7 Certkiller 8
- D. Certkiller 1 Certkiller 2 Certkiller 3 Certkiller 5 Certkiller 4 Certkiller 7 Certkiller 8
- E. Certkiller 1 Certkiller 2 Certkiller 4 Certkiller 7 Certkiller 8
- F. None of the above

Answer: E

Explanation:

With OSPF, An area's topology is invisible to entities outside the area. By keeping area topologies separate, OSPF passes less routing traffic than it would if the AS were not partitioned.

Area partitioning creates two different types of OSPF routing, depending on whether the source and the destination are in the same or different areas. Intra-area routing occurs when the source and destination are in the same area; interarea routing occurs when they are in different areas.

The backbone itself is an OSPF area, so all backbone routers use the same procedures and algorithms to maintain routing information within the backbone that any area router would. The backbone topology is invisible to all intra-area routers, as are individual area topologies to the backbone.

In this case, when the traffic reached router Certkiller 2, it would forward it directly to the backbone area, area 0, since this traffic is destined to another area. Since Certkiller 2 is unaware of the metrics inside area 0, it will go directly to Certkiller 4. Note that even though going via Certkiller 3 would result in a lower overall cost, the traffic will be sent directly to area 0 instead of to another router in area 1.

Once Router Certkiller 4 receives the traffic, it will forward it normally using the lowest cost path in area 2, which is via Certkiller 7.

QUESTION 162

The Company network has a number of routers with very high speed interfaces, and you want to ensure that this is reflected in the OSPF process. What IOS command would be used to reset the cost calculation process so that high-speed interfaces can be correctly calculated?

- A. (config-if)#ip ospf cost xxx
- B. (config-if)# ip ospf interface-speed xxx
- C. (config-if)# ip ospf auto-cost reference-bandwidth xxx
- D. (config-router)# ospf auto-cost reference-bandwidth xxx
- E. (config)# ip ospf auto-cost reference-bandwidth xxx

Answer: D

Explanation:

To control how OSPF calculates default metrics for the interface, use the ospf auto-cost command in router configuration mode. In Cisco IOS Release 10.3 and later, 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, while a T1 link will have a metric of 64.

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

If you have multiple links with high bandwidth, you might want to use a larger number to differentiate the cost on those links.

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

Example:

The following example changes the cost of a Fast Ethernet link to 10, while the gigabit Ethernet link remains at a cost of 1. Thus, the link costs are differentiated.

```
router ospf 1
```

```
ospf auto-cost reference-bandwidth 1000
```

Reference:

http://www.cisco.com/en/US/docs/ios/12_0/np1/command/reference/1rospf.html

QUESTION 163

Certkiller uses link-state routing protocols within its network. With these routing protocols, which of the following are true? (Select three)

- A. Each router sends has knowledge of all routers in the flooding domain
- B. Each router sends routing information to all nodes in the flooding domain
- C. Each router individually builds a picture of the entire flooding domain
- D. Each router sends all or some portion of its routing table to neighbor routers

- E. Each router is aware of neighbor routers
- F. Each router install routes direct from the routing updates into the routing table

Answer: A, B, C

Explanation:

Link-State Versus Distance Vector Link-state algorithms (also known as shortest path first algorithms) flood routing information to all nodes in the internetwork. Each router, however, sends only the portion of the routing table that describes the state of its own links. In link-state algorithms, each router builds a picture of the entire network in its routing tables. Distance vector algorithms (also known as Bellman-Ford algorithms) call for each router to send all or some portion of its routing table, but only to its neighbors. In essence, link-state algorithms send small updates everywhere, while distance vector algorithms send larger updates only to neighboring routers. Distance vector algorithms know only about their neighbors.

Because they converge more quickly, link-state algorithms are somewhat less prone to routing loops than distance vector algorithms. On the other hand, link-state algorithms require more CPU power and memory than distance vector algorithms. Link-state algorithms, therefore, can be more expensive to implement and support.

QUESTION 164

In the multi access Certkiller OSPF network, what best describes how neighbor adjacencies are formed?

- A. The router with the highest priority become the DR
- B. The router with the lowest router ID
- C. Election of the DR and BDR begins only after the router that wants to become either DR or BDR enters the "exstart" state
- D. Only those router with the Cisco default priority become the DR or BDR
- E. The router with the highest loopback interface
- F. None of the above

Answer: C

Explanation:

Technically, when a router and its neighbor enter the ExStart state, their conversation is characterized as an adjacency, but the routers have not become fully adjacent yet. ExStart is established using Type 2 database description (DBD) packets, also known as DDPs. The two neighbor routers use hello packets to negotiate who is the "master" and who is the "slave" in their relationship and DBD packets to exchange databases. The router with the highest OSPF router ID "wins" and becomes master (the DR). When the neighbors establish their roles as master and slave (slave being the BDR or DROTHER), they enter the Exchange state and begin sending routing information.

QUESTION 165

Two Certkiller routers named CK1 and CK2 , configured to run BGP have been

connected to a firewall, one on the inside interface and one on the outside interface. BGP has been configured so that these two routers should peer, including configuration of the correct BGP session endpoint addresses and the correct BGP session hop count limit (EBGP multihop). What is a good first test to see if BGP works between CK1 and CK2 across the firewall?

- A. Attempt to telnet from the router connected to the inside of the firewall to the router connected to the outside of the firewall. If telnet works BGP will work, since telnet and BGP both use TCP to transport data
- B. There is no way to make BGP work across a firewall
- C. Ping from the router connected to the inside interface of the firewall to the router connected to the outside interface of the firewall. If you can ping between them BGP should work, since BGP uses IP to transport packets
- D. There is no way to make BGP work across a firewall without special configuration, so there's simple test that would show you if BGP will work or not other than trying to start the peering session
- E. None of the above

Answer: A

Explanation:

Because BGP uses unicast TCP packets on port 179 to communicate with its peers, you can configure the firewall to allow unicast traffic on TCP port 179. This way, BGP peering can be established between the routers that are connected through the firewall. For an example configuration of BGP through PIX firewalls, see the reference link below.

Reference:

http://www.cisco.com/en/US/tech/CK365/technologies_configuration_example09186a008009487d.shtml

QUESTION 166

You are the network administrator at Certkiller . You want to advertise the network 190.72.27.0/27 to an EBGP peer. What command should you use?

- A. network 190.72.27.0
- B. network 190.72.27.0 mask 255.255.255.224
- C. network 190.72.27.0 mask 255.255.225.240
- D. network 190.72.27.0 mask 0.0.0.31.

Answer: B

Explanation:

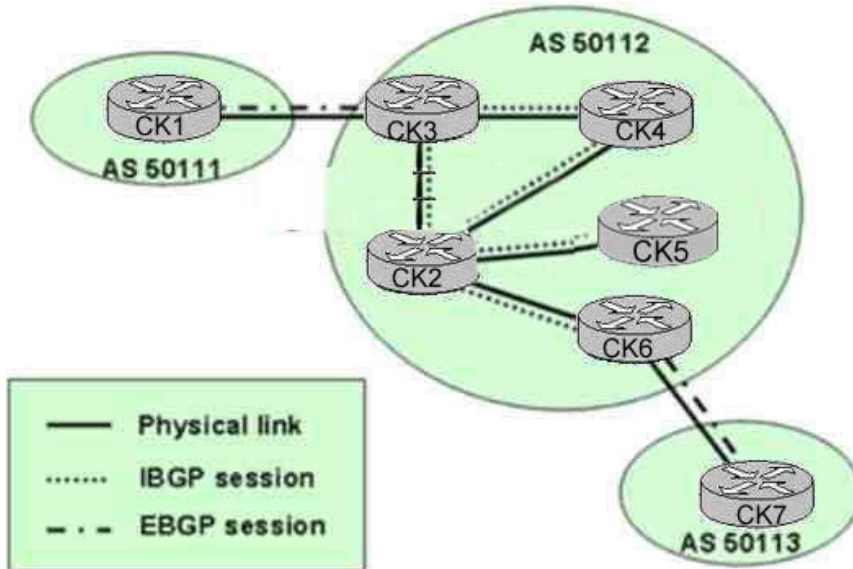
The correct syntax is: network ip-address mask subnet-mask where ip-address is the network address and subnet-mask is the subnet mask. In this case the network address is 190.72.27.0. The subnet mask is a 27 bit subnet mask (11111111.11111111.11111111.11100000) that equates to 255.255.255.224.

Incorrect Answers:

- A. If no mask is specified, the default class mask is used. In the 190.72.27.0 case it would be a /16.
- C. Here the wrong mask is used.
- D. This is the inverse mask, which is normally used by OSPF when specifying the network mask, but not by BGP.

QUESTION 167

The Certkiller BGP network consists of AS 50112 as shown in the diagram below:



Based on the physical connectivity and the IBGP peering shown, what router within the Transit AS 50112 should be setup as the route reflector and which routers should be setup as the clients based on the recommended route reflector design rules?

- A. CK4 should be the route reflector with CK2 and CK5 as its clients.
- B. CK2 should be the route reflector with CK5 and CK6 as its clients.
- C. CK3 should be the route reflector with CK2 and CK4 as its clients.
- D. CK2 should be the route reflector with CK4 and CK5 as its clients.
- E. CK4 should be the route reflector with CK2 and CK3 as its clients.
- F. All of the above are valid options.

Answer: B

Explanation:

Within any BGP autonomous system, every IBGP speaker must have a fully meshed peering arrangement with every other iBGP speaker. This is due to the fact that a BGP speaker will not advertise a route learned via another iBGP speaker to a third iBGP speaker. The use of route reflectors is one way to maintain connectivity throughout the AS without having a fully meshed peering arrangement. By relaxing this restriction a bit and by providing additional control, we can allow a router to advertise (reflect) iBGP

learned routes to other iBGP speakers.

When using route reflectors, the clients need only peer to the route reflector. In the example above, if router CK2 is configured as the route reflector, with routers CK5 and CK6 set up as clients, then 5 and 6 need only peer with CK2 . In doing this, all other routers are fully meshed. No other answer choices will allow us to maintain a fully meshed iBGP configuration.

QUESTION 168

Routers CK1 and CK2 are configured for BGP. Both routers reside in AS 65234. Routes from Router CK2 show up in the BGP table on Router CK1 , but not in the IP routing table.

What could be the cause of this problem?

- A. Synchronization is off.
- B. The BGP peers are down.
- C. BGP multi-hop is disabled on Router CK1 .
- D. Router CK1 is not receiving the same routes via an internal protocol.

Answer: D

Explanation:

BGP Synchronization says: "If your autonomous system is passing traffic from another AS to a third AS, BGP should not advertise a route before all routers in your AS have learned about the route via IGP." Therefore, we can assume that synchronization is on and that the BGP routes have not yet been learned by an IGP.

Incorrect Answers:

- A. If synchronization is off the routes would show up in the IP routing table on CK1 .
- B. If the BGP peers were down, then the routers would not be sending and receiving BGP route information to each other.
- C. BGP multi-hop is only useful for EBGP peers, not IBGP peers.

QUESTION 169

You have a router running BGP for the Internet connections as well as IGRP for use internally. You configure the network backdoor command on this router under the BGP process. What will this do?

- A. It will change the distance of an iBGP route to 20.
- B. It will change the distance of an eBGP route to 200.
- C. It will change the distance of an IGRP route to 20.
- D. It will not change the distance of the route.

Answer: B

Explanation:

Backdoor only makes the IGP learned route the preferred route. To specify a backdoor

route to a BGP border router that will provide better information about the network, use the network backdoor router configuration command. To remove an address from the list, use the no form of this command.

By definition, eBGP updates have a distance of 20 that is lower than the IGP distances.

Default distance is 120 for RIP, 100 for IGRP, 90 for EIGRP, and 110 for OSPF.

By default, BGP has the following distances, but that could be changed by the distance command:

```
distance bgp external-distance internal-distance local-distance
```

```
external-distance:20
```

```
internal-distance:200
```

```
local-distance:200
```

If we want RTA to learn about 160.10.0.0 via RTB (IGP), then we have two options:

- * Change eBGP's external distance or IGP's distance, which is not recommended.

- * Use BGP backdoor.

BGP backdoor makes the IGP route the preferred route

RTA learns 160.10.0.0 from RTB via EIGRP with distance 90, and also learns it from

RTC via eBGP with distance 20. Normally eBGP is preferred, but because of the

backdoor command EIGRP is preferred

References:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1826/products_command_summary_chapter09186a00800d

http://www.cisco.com/en/US/tech/CK365/CK80/technologies_tech_note09186a00800c95bb.shtml#bgpbackdoor

QUESTION 170

You have two routers running BGP to two different ISP's. You wish to influence the way that traffic comes into your network from the Internet, but your company policy prohibits the use of BGP communities. What is the best way to influence this traffic?

- A. Adjust the cost of your routers.
- B. Use MED values.
- C. Increase the weight value on one of your routers.
- D. Decrease the local preference value on one of your routers.
- E. Use AS-path prepending.
- F. Use Metrics.

Answer: E

Explanation:

When influencing incoming traffic from the Internet, the two most widely used methods are AS Path Prepending and Multi-Exit Discriminators (MED). AS Path prepending works by adding AS paths to certain network ranges, making them appear to the Internet to be further away than they really are. MEDs are used to advertise metrics to the neighbor AS to influence the incoming path that traffic should take to reach certain destinations. In this case, AS Path Prepending is preferred over the use of MEDs because

AS path prepending information is distributed to all networks within the Internet. MEDs are only used between neighboring Autonomous Systems. Another advantage to path prepending is that the AS path information is ranked higher in the BGP decision process than the MED information. IN fact, MED information is one of the last things considered in the BGP path decision algorithm.

Note: Although one method of using AS Path prepending requires the use of communities, it is not required to use communities for simply sending prepending information.

Incorrect Answers:

A, C, D, F. These are all methods for influencing traffic going out to the Internet, not coming in.

E. This would be an acceptable way to influence traffic, but would not be the best way.

QUESTION 171

Your router is multi-homed to three different ISP's for Internet access. You then configure "bgp deterministic-med" under the BGP routing process configuration of your router. What effect does this change have on your network?

- A. It configures BGP to compare MEDs between different ASs.
- B. It makes the default metric count the worst possible metric.
- C. It makes the default metric count the best possible metric.
- D. It configures BGP to reorder the entries by neighbor AS.
- E. It configures BGP to reorder the entries by MED.

Answer: D

Explanation:

There is sometimes confusion between the two Border Gateway Protocol (BGP) configuration commands `bgp deterministic-med` and `bgp always-compare-med`. Enabling the `bgp deterministic-med` command ensures the comparison of the MED variable when choosing routes advertised by different peers in the same autonomous system. Enabling the `bgp always-compare-med` command ensures the comparison of the MED for paths from neighbors in different autonomous systems. The `bgp always-compare-med` command is useful when multiple service providers or enterprises agree on a uniform policy for setting MED. Thus, for network X, if Internet Service Provider A (ISP A) sets the MED to 10, and ISP B sets the MED to 20, both ISPs agree that ISP A has the better performing path to X.

When BGP receives multiple routes to a particular destination, it lists them in the reverse order that they were received, from the newest to the oldest. BGP then compares the routes in pairs, starting with the newest entry and moving toward the oldest entry (starting at top of the list and moving down). For example, entry1 and entry2 are compared. The better of these two is then compared to entry3, and so on. The `bgp always-compare-med` command reorders the entries by neighbor AS.

Incorrect Answers:

- A. The router would compare MEDs between different AS numbers if the "bgp

always-comapre-med" was configured, not the bgp deterministic-med command.

B, C. This command does not affect the default BGP metric.

E. This command reorders the entries based on AS number, not MED.

Reference:

http://www.cisco.com/en/US/tech/CK365/technologies_tech_note09186a0080094925.shtml

QUESTION 172

Which of the following attributes are "well known" BGP attributes? (Choose all that apply)

- A. Atomic-aggregate
- B. MED
- C. Next-hop
- D. AS-path
- E. Origin
- F. Weight
- G. Aggregator

Answer: A, C, D, E

Explanation:

The following BGP attributes are all well known:

Well Known, Mandatory attributes: AS_PATH, NEXT-HOP and ORIGIN

Well Known, Discretionary attributes: LOCAL_PREF and ATOMIC_AGGREGATE

Incorrect Answers:

B, E, F. The optional, transitive attributes are AGGREGATOR and COMMUNITY. The optional non-transitive attributes include MULTI_EXIT_DISC (MED, the ORIGINATOR_ID. and CLUSTER_LIST.

Reference:

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm

QUESTION 173

In BGP routing, what does the rule of synchronization mean?

- A. It means that a BGP router can only advertise an iBGP-learned route provided that the route is in the only in the BGP table.
- B. It means that a BGP router can only advertise an eBGP-learned route provided that the route is an IGP route in the routing table.
- C. It means that a BGP router can only advertise an iBGP-learned route provided that the route is in the routing table of all its iBGP neighbors.
- D. It means that a BGP router can only advertise an eBGP-learned route provided that the route is metric 0 in the BGP table.
- E. It means that a BGP router can only advertise an iBGP-learned route provided that the route is an IGP route in the routing table.

Answer: E

Explanation:

The BGP rule of synchronization states that a BGP router should not advertise to external neighbors destinations learned from IBGP neighbors unless those destinations are also known via an IGP.

Incorrect Answers:

B, D. Synchronization is used to ensure that you don't develop black holes by advertising local routes to the rest of the world, when the local routers don't even know how to get to the route in question. That's why synchronization with the IGP is not a concern when you either create a full iBGP mesh, or implement route reflectors, confederations, or both. Therefore, synchronization is implemented only for IBGP routes, not EBGP.

C. The route needs only be in the routing table of its own router, not every neighboring router.

Reference: "Internet Routing Architectures" Sam Halabi page 143, Cisco Press.

QUESTION 174

What is the correct sequence order that BGP routers use when determining the best route to any given destination?

- A. MED, Local preference, AS-path, Weight, Origin Code
- B. Origin Code, MED, Weight, AS Path, Local Preference
- C. Weight, Local Preference, AS-path, Origin Code, MED
- D. Weight, Local Preference, MED, AS-Path, Origin Code
- E. MED, Weight, Local Preference, Origin Code, AS Path

Answer: C

Explanation:

How the Best Path Algorithm Works

BGP assigns the first valid path as the current best path. It then compares the best path with the next path in list, until it reaches the end of the list of valid paths. Following is a list of rules used to determine the best path:

1. Prefer the path with the largest WEIGHT. Note: WEIGHT is a Cisco-specific parameter, local to the router on which it's configured.
2. Prefer the path with the largest LOCAL_PREF.
3. Prefer the path that was locally originated via a network or aggregate BGP subcommand, or through redistribution from an IGP. Local paths sourced by network/redistribute commands are preferred over local aggregates sourced by the aggregate-address command.
4. Prefer the path with the shortest AS_PATH. Note the following:
 5. 1. This step is skipped if bgp bestpath as-path ignore is configured.
 2. An AS_SET counts as 1, no matter how many ASs are in the set.
 3. The AS_CONFED_SEQUENCE is not included in the AS_PATH length.
4. Prefer the path with the lowest origin type: IGP is lower than EGP, and EGP is lower than INCOMPLETE.
5. Prefer the path with the lowest multi-exit discriminator (MED).

QUESTION 175

You are setting up BGP on router CK1 and you wish to simplify the configuration file through the use of BGP peer groups. Which of the following best describes the proper use of BGP peer groups?

- A. They should be used for peers with common community values
- B. They should be used for peers with common inbound announcement policies
- C. They should be used for peers with common outbound announcement policies
- D. They should be used to combine MED inbound policies
- E. They should be used to peers with common transitive AS policies

Answer: C

Explanation:

The major benefit of specifying a BGP peer group is that it reduces the amount of system resources (CPU and memory) used in an update generation, and it also simplifies the BGP configuration. It reduces the load on system resources by allowing the routing table to be checked only once, and updates to be replicated to all peer group members instead of being done individually for each peer in the peer group. Depending on the number of peer group members, the number of prefixes in the table, and the number of prefixes advertised, this can significantly reduce the load. Cisco recommends that you group together peers with identical outbound announcement policies.

QUESTION 176

Router CK1 is being configured for as both an IBGP peer to the other routers within the Certkiller network, and as an EBGP peer to the ISP. Select the BGP attributes that are required to be sent to these BGP neighbors from CK1 :

- A. AS_PATH
- B. MED
- C. NEXT_HOP
- D. LOCAL_PREF
- E. ORIGIN
- F. ROUTER_ID

Answer: A, C, E

Explanation:

AS-PATH, NEXT-HOP, and ORIGIN are all well known, mandatory BGP attributes, which is defined below:

Well known mandatory attributes: These attributes must be recognized by all BGP speakers, and must be included in all update messages. Almost all of the attributes impacting the path decision process, described in the next section, are well known mandatory attributes.

Origin Code:

The ORIGIN is a well known mandatory attribute that indicates the origin of the prefix, or rather, the way in which the prefix was injected into BGP. There are three origin codes, listed in order of preference:

IGP, meaning the prefix was originated from information learned from an interior gateway protocol

EGP, meaning the prefix originated from the EGP protocol, which BGP replaced

INCOMPLETE, meaning the prefix originated from some unknown source

AS Path:

The AS_PATH is a well-known mandatory attribute and is the list of all autonomous systems the prefixes contained in this update have passed through. The local autonomous system number is added by a BGP speaker when advertising a prefix to an eBGP peer.

Next Hop:

The BGP NEXT_HOP is a well-known mandatory attribute. The Next Hop attribute is set when a BGP speaker advertises a prefix to a BGP speaker outside its local autonomous system (it may also be set when advertising routes within an AS, this will be discussed in later sections). The Next Hop attribute may also serve as a way to direct traffic to another speaker, rather than the speaker advertising the route itself.

Incorrect Answers:

B. The MUTLI_EXIT_DISC (MED) is an optional non-transitive attribute that provides a mechanism for the network administrator to convey to adjacent autonomous systems to optimal entry point in the local AS.

D. The LOCAL_PREF attribute is a well-known attribute that represents the network operator's degree of preference for a route within the entire AS. It is not a mandatory attribute and it is not applied to all BGP updates.

F. The router ID is not a well known, mandatory BGP attribute.

QUESTION 177

Assume the following routes are in the BGP routing table of router CK1 .

172.16.0.0/24

172.16.1.0/24

172.16.2.0/24

172.16.3.0/24

Also assume the following commands have been configured:

```
router bgp 1
```

```
neighbor 10.1.1.1 remote-as 2
```

```
aggregate-address 172.16.0.0 255.255.252.0 suppress-map specific
```

```
access-list 1 permit 172.16.0.0 0.0.0.3.255
```

```
route-map specific permit 10
```

```
match ip-address 1
```

Which BGP routes will CK1 advertise?

A. 172.16.0.0/22

B. 172.16.0.0/22, 172.16.2.0/24, 172.16.3.0/24

C. 172.16.0.0/22, 172.16.0.0/24, 172.16.1.0/24

D. 172.16.2.0/24 and 172.16.3.0/24

E. 172.16.0.0/22 and 172.16.1.0/24

Answer: A

Explanation:

BGP allows the aggregation of specific routes into one route using the aggregate-address address mask command. Aggregation applies to routes that exist in the BGP routing table. This is in contrast to the network command, which applies to the routes that exists in IP routing table. Aggregation can be performed if at least one or more of the specific routes of the aggregate address exist in the BGP routing table. In this specific example, the router will summarize the routes into 172.16.0.0/22, as long as at least one of the more specific 172.16 assumed routes actually exist in the routing table. Normally, aggregate addresses are advertised in addition to the more specific subnets. However, in this case the suppress map will filter the more specific routes, advertising only the 172.16.0.0/22 route.

QUESTION 178

A BGP router in the Certkiller network called P1R3 is configured as shown below:

```
!  
hostname P1R3  
!  
! Output omitted  
!  
router bgp 50001  
synchronization  
bgp log-neighbor-changes  
neighbor 10.200.200.11 remote-as 50001  
neighbor 10.200.200.11 update-source loopback0  
neighbor 10.200.200.12 remote-as 20001  
neighbor 10.200.200.12 update-source Loopback0  
neighbor 10.200.200.14 remote-as 50001  
neighbor 10.200.200.14 update-source Loopback0  
no auto-summary  
P1R3#show ip bgp summary  
BGP router identifier 10.200.200.13, local As number 50001  
BGP table version is 1, main routing table version 1  
6 network entries using 606 bytes of memory  
7 path entires using 336 bytes of memory  
4 BGP path attribute entries using 240 bytes of memory  
3 BGP AS-PATH entries using 72 bytes of memory  
0 BGP route-map cache entries using 0 bytes of memory  
0 BGP filter-list cache entries using 0 bytes of memory  
BGP using 1254 total bytes of memory  
BGP activity 6/0 prefixes, 7/2 paths, scan interval 60 secs  
Neighbor V AS MsgRcvd MsgSent TblVer InO OutO Up/Down State/Pfxrcd  
10.200.200.11 4 50001 9 4 1 0 0 00:00: 14 6  
10.200.200.12 4 50001 9 4 1 0 0 00:00: 14 6
```

```
10.200.200.14 4 50001 4 4 1 0 0 00:00: 14 0
PIR#show ip bgp
BGP table version is 1, local router: ID is 10.200.200.13
Status Codes: s suppressed, d damped, h history, * valid, > best, I - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
Network Next Hop Metric LocPrf Weight Path
* i10.0.0.0 10.200.200.12 0 100 0 i
* i 10.200.200.11 0 100 0 i
* i192.168.11.0 10.200.200.12 0 100 0 50998 50222 50223 i
* i 10.200.200.11 0 100 0 50998 50222 50223 i
* i192.168.12.0 10.200.200.12 0 100 0 50998 50222 50223 i
* i 10.200.200.11 0 100 0 50998 50222 50223 i
* i192.168.13.0 10.200.200.12 0 100 0 50998 50222 50223 i
* i 10.200.200.11 0 100 0 50998 50222 50223 i
* i192.168.14.0 10.200.200.11 0 100 0 50998 50222 50223 i
* i 10.200.200.11 0 100 0 50998 50222 50223 i
```

<output omitted>

```
PIR#show ip route
Codes: C - connected, s - static, I IGRP, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, 0 - OSPF< IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
I - is-is, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level - 2
Ia - IS-IS inter area, * - candidate default, U - per-user static route
O - ODR, P - periodic downloaded static route
Gateway of last resort is not set
```

```
10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
o 10.200.200.11/32 [110/11] via 10.1.1.1, 00:06:38, Ethernet0/0
o 10.200.200.14/32 [110/65] via 10.1.3.4, 00:06:38, Serial0/0
o 10.200.200.12/32 [110/75] via 10.1.1.1, 00:06:38, Ethernet0/0
c 10.200.200.13/32 is directly connected, Loopback0
c 10.1.3.0/24 is directly connected, Serial0/0
o 10.1.2.0/72 [110/74] via 10.1.3.4, 00:06:38, Serial0/0
c 10.1.1.0/24 is directly connected, Ethernet0/0
c 10.1.0.0/24 [110/74] via 10.1.1.1, 00:06:38, Ethernet 0/0
```

Router PIR3 is running an IBGP full-mesh with its IBGP neighbors (10.200.200.11, 10.200.200.12, and 10.200.200.14). Based on the BGP configuration and the show command outputs above, why are BGP routes not being selected in the BGP table and placed into the IP routing table?

- A. Because the 10.200.200.11 and 10.200.200.12 neighbors are setting the Weight to 0
- B. Because the 10.200.200.11 and 10.200.200.12 neighbors are setting the MED to 0
- C. Because the 10.200.200.11 and 10.200.200.12 neighbors are not using next-hop-self
- D. Because synchronization is enabled on PIR3
- E. Because there are no routes to reach the next-hops

Answer: D

Explanation:

A BGP router with synchronization enabled will not advertise iBGP-learned routes to other eBGP peers if it is not able to validate those routes in its IGP. Assuming that IGP has a route to iBGP-learned routes, the router will announce the iBGP routes to eBGP peers. Otherwise the router treats the route as not being synchronized with IGP and does not advertise it. Disabling synchronization using the no synchronization command under router BGP prevents BGP from validating iBGP routes in IGP. By default, synchronization is enabled on all BGP routers.

QUESTION 179

With regards to BGP and the administrative distance in a routed environment, which statement is correct?

- A. The administrative distance of all BGP routes is 20, which explains why BGP routes are preferred over any IGP (such as OSPF).
- B. BGP is a path vector protocol, and thus does not employ the concept of administrative distance.
- C. BGP dynamically adjusts its administrative distance to match that of the IGP within the AS to eliminate routing confusion.
- D. BGP actually employs two different administrative distance values: IBGP is 20, while EBGP is 200.
- E. BGP actually employs two different administrative distance values: IBGP is 200, while EBGP is 20.

Answer: E

Explanation:

BGP employs the use of two separate administrative distances, based on the type of BGP route. (Internal or External)

The table below lists the administrative distance default values of the protocols that Cisco supports:

Route Source	Default Distance Values
Connected interface	0
Static route*	1
Enhanced Interior Gateway Routing Protocol (EIGRP) summary route	5
External Border Gateway Protocol (BGP)	20
Internal EIGRP	90
IGRP	100
OSPF	110
Intermediate System-to-Intermediate System (IS-IS)	115
Routing Information Protocol (RIP)	120
Exterior Gateway Protocol (EGP)	140
On Demand Routing (ODR)	160
External EIGRP	170
Internal BGP	200
Unknown**	255

Incorrect Answers:

- A. Only external BGP routes have an AD of 20. Internal BGP routes are given a high AD to prevent these routes from overriding the routes from the IGP routing protocols, such as OSPF, EIGRP, RIP, etc.
- B. BGP is indeed considered a path vector routing protocol, but it does also use the concept of AD, as shown in the table above.
- C. The AD of BGP routes is static, with the default values shown in the table. These values can be configured to use different values, but they will still be considered static and will not change dynamically.
- D. BGP does indeed use two different values, but the values used are the reverse. EBGp is 20 while IBGP is 200.

QUESTION 180

You are configuring the Certkiller Internet router as a BGP peer to your ISP's router. After doing this, which BGP attributes will be carried in every BGP update (both IBGP and EBGp)?

- A. Origin, AS-Path, Next Hop
- B. Origin, local preference, AS-Path
- C. Router-ID, Origin, AS-Path
- D. Router-ID, Local-Preference, Next-Hop
- E. AS-Path, Local Preference, Next-Hop

Answer: A

Explanation:

Origin, AS-PATH, and Next-Hop are all well known, mandatory BGP attributes, which is defined below:

Well known mandatory attributes: These attributes must be recognized by all BGP speakers, and must be included in all update messages. Almost all of the attributes impacting the path decision process, described in the next section, are well known mandatory attributes.

Origin Code

The ORIGIN is a well known mandatory attribute that indicates the origin of the prefix, or rather, the way in which the prefix was injected into BGP. There are three origin codes, listed in order of preference:

1. IGP, meaning the prefix was originated from information learned from an interior gateway protocol
2. EGP, meaning the prefix originated from the EGP protocol, which BGP replaced
3. INCOMPLETE, meaning the prefix originated from some unknown source

AS Path

The AS_PATH is a well-known mandatory attribute and is the list of all autonomous systems the prefixes contained in this update have passed through. The local autonomous system number is added by a BGP speaker when advertising a prefix to an eBGP peer.

Next Hop

The BGP NEXT_HOP is a well-known mandatory attribute. The Next Hop attribute is set when a BGP speaker advertises a prefix to a BGP speaker outside its local autonomous system (it may also be set when advertising routes within an AS, this will be discussed in later sections). The Next Hop attribute may also serve as a way to direct traffic to another speaker, rather than the speaker advertising the route itself.

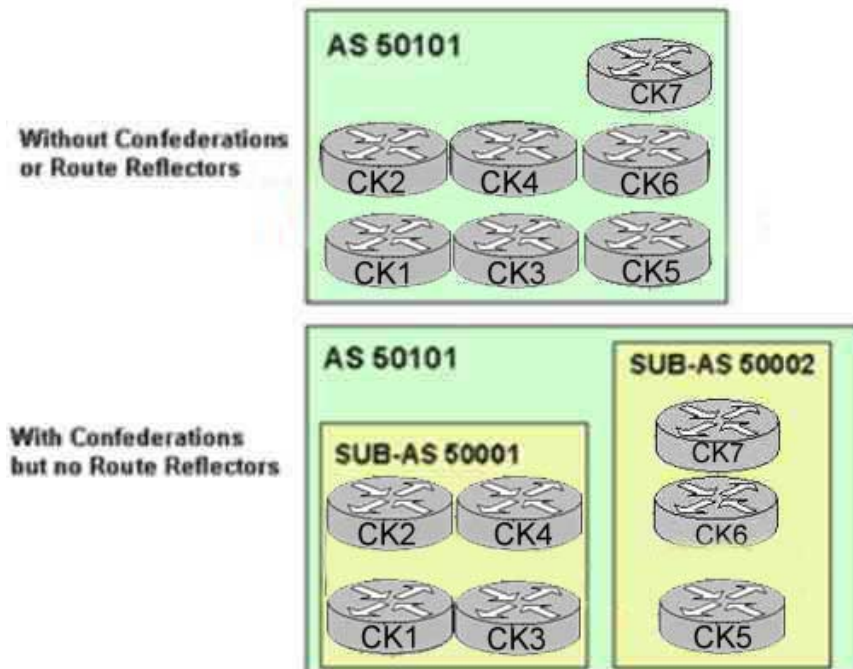
Incorrect Answers:

B, D, E. The LOCAL_PREF attribute is a well-known attribute that represents the network operator's degree of preference for a route within the entire AS. It is not a mandatory attribute that is applied to all BGP updates.

C, D. The router ID is not a well known, mandatory BGP attribute.

QUESTION 181

The Certkiller BGP network has been assigned AS number 50101 as shown below:



The Certkiller AS 50101 network is split into two AS numbers (Sub-AS 50001 and Sub-AS 50002) using Confederations without any route reflectors. Sub-AS 50001 contains 4 routers and sub-AS 50002 contains the other 3 routers. Based on this information, how many IBGP sessions are required?

- A. 9 IBGP sessions using Confederations with two sub-ASs where one of the sub-AS contains 4 routers and the other sub-As contains the other 3 routes.
- B. 11 IBGP sessions using Confederations with two sub-ASs where one of the sub-AS contains 4 routers and the other sub-As contains the other 3 routes.
- C. 18 IBGP sessions using Confederations with two sub-ASs where one of the sub-AS contains 4 routers and the other sub-As contains the other 3 routes.
- D. 21 IBGP sessions using Confederations with two sub-ASs where one of the sub-AS contains 4 routers and the other sub-As contains the other 3 routes.
- E. 25 IBGP sessions using Confederations with two sub-ASs where one of the sub-AS contains 4 routers and the other sub-As contains the other 3 routes.

Answer: A

Explanation:

The advantage of confederations is that they sharply reduce the number of IBGP peering sessions. IBGP is used normally within each member AS, but a special version of EBGP known as confederations. EBGP is run between the autonomous systems.

Confederations are another way of scaling IBGP. Defined in RFC 3065, this feature introduces a divide-and-conquer approach to remove the full mesh requirement.

Using confederations an AS is split into multiple sub-ASs, but the network still appears as one AS to the outside world. Each sub-AS number is stripped from AS path at the confederation border. A full IBGP mesh is only required within each sub-AS, which is usually a manageable number of routers. In very large networks, you can even configure

route reflection within a sub-AS. Typically private ASs are assigned for each sub-AS number.

With IBGP, all routers are to be configured as a fully meshed topology. The number of connections needed for any fully meshed configuration can be found by the formula:

$N(N-1)$

2

There are 4 Sub-AS peers in 5001 so that makes $4*3 / 2 = 6$ peer sessions.

Similarly, there are 3 peers in Sub-As 5002, so we have $3*2 / 2 = 3$ peer sessions

Therefore, the total number of peering sessions is $9(6+3)$.

Reference: Jeff Doyle, "Routing TCP/IP" Vol. II page 287

QUESTION 182

Router CK1 is used as the Certkiller Internet router and is configured for BGP. The Ip BGP information of this router is displayed below:

CK1 # show ip bgp

BGP table version is 12, local router ID is 172.16.1.2

Status code: s supported, d damped, h history, * valid, > best,

i - internal Origin codes: i IGP, e - EGP ? - incomplete

Network Next Hop Metric LocPrf Weight Path

*> 192.168.0.0/16 172.16.1.1 0 0 50103 {50101, 50102} i

Given above information, why does the 192.168.0.0/16 prefix contain an AS-PATH of 50103 { 50101, 50102}

- A. Because AS 50101 and AS 50102 are Transit AS's
- B. Because AS 50103 is using BGP confederations with two sub-ASs (sub-AS 50101 and sub-AS 50102)
- C. Because it is an aggregate route and the more specific routes have passed through AS 50101 and AS 50102
- D. Because AS 50103 is using AS-Path pre-pending to influence the return traffic
- E. Because AS 50103 is performing route summarization using the network 192.168.0.0. mask 255.255.0.0 command

Answer: C

Explanation:

In this example, the 192.168.0.0/16 route includes the SET {50101, 50102}. This indicates that aggregate route of 192.168.0.0 actually summarizes routes that have passed through AS 50101 and AS 50102. The AS-SET information is preserved because it becomes important in avoiding loops as it maintains an indication of where the route has been.

Incorrect Answers:

- A. Transit AS numbers are displayed normally in the IP BGP table.
- B. Confederations are seen as only one single AS to the rest of the Internet, so they will not appear as an AS-SET to EBGP peers.
- D. AS Path prepending is displayed normally, and if this were the case then you would

see multiple entries in a row for the same AS number.

E. Summarized routes only appear in an AS SET when the more specific routes have passed through multiple different AS numbers.

Reference: Bassam Halabi, "Internet Routing Architectures" Cisco Press, page 359.

QUESTION 183

The IP BGP information for a specific network on router CK1 is displayed below:

CK1 #show ip bgp 10.254.0.0

BGP routing table entry for 10.254.0.0/24, version 8

Paths: (2 available, best #1, table

Default-IP-Routing-Table, not advertised

Advertised to non peer-group peers:

10.1.0.2 10.200.200.13 10.200.200.14

50998

172.31.1.3 from 172.31.1.3 (172.31.1.3)

Origin IGP, metric 0, localpref 100, valid, external, best

Community: 50998:1 no-export

50998

172.31.1.3 from 10.1.0.2 (10.200.200.12)

Origin IGP, metric 0, localpref 100, valid, internal

Router CK1, which is in Transit AS 50001, is not propagating the 10.254.0.0/24 prefix to its neighboring ASs. Based on the "show IP BGP 10.254.0.0" output shown, determine a possible cause of this problem.

- A. Because the 10.254.0.0/24 prefix is tagged with the no-export community
- B. Because the best path chosen by BGP is the IBGP learned path
- C. Because the best path chosen by BGP is the EBGP learned path
- D. Because the 10.254.0.0/24 prefix has a MED of 0
- E. Because of the EBGP split horizon rule

Answer: A

Explanation:

From the output shown above, the 10.254.0.0 route is indeed tagged with the BGP community of no-export. The Well Known BGP community of NO EXPORT means that the route can be advertised to other IBGP peers, but it is not to be passed to EBGP peers. If the BGP community of NO ADVERTISE was used instead, then this route would not be forwarded to both IBGP as well as EBGP peers.

Incorrect Answers:

B, C. Regardless of the path that the BGP route was learned, the default behavior is to forward the route to EBGP peers.

D. The metric 0 shown in the example above is the normal behavior for IBGP learned routes.

E. BGP does not use the split horizon rule. This rule applies to distance vector interior routing protocols. BGP is considered to be a path vector external routing protocol.

QUESTION 184

The Certkiller network is using BGP for Internet routing, and part of the router configuration is shown below:

```
router bgp 50101
neighbor 10.1.1.1 remote-as 50102
neighbor 10.2.2.2 remote-as 50103
neighbor 10.2.2.2 route-map test2 out
neighbor 10.1.1.1 route-map test out
!
ip as-path access-list 1 permit _50104$
ip as-path access-list 2 permit .*
!
route-map test permit 10
match as-path 1
set metric 140
!
route-map test permit 20
match as-path 2
!
route-map test2 permit 10
set metric 100
```

Based on the configuration above, which statements is correct? Select all that apply.

- A. All prefixes originating in AS 50104 will be advertised to the 10.1.1.1 neighbor with a MED of 150.
- B. All prefixes not originating in AS 50104 will be advertised to the 10.1.1.1 neighbor with a MED of 0.
- C. All prefixes not originating in AS 50104 will be advertised to the 10.1.1.1 neighbor.
- D. All prefixes will be advertised to the neighbor with a MED of 100.
- E. All prefixes originating in AS 50104 will be advertised to the 10.2.2.2 and the 10.1.1.1 neighbor with a MED of 100.

Answer: A, B

Explanation:

For the 10.1.1.1 BGP peer, route-map "test" is being applied. This route map has two statement entries. The first states that all traffic originating from AS 50104 (as shown by the "ip as-path access-list 1 permit _50104\$" command statement) should have the MED set to 140. The regular expression ".*" matches everything else, so all other traffic is to be routed normally. Since the default MED value is 0, all other traffic not originating in AS will be advertised to the 10.1.1.1 peer with a MED of 0.

Incorrect Answers:

- A. The MED value advertised to the 10.1.1.1 peer that originated from AS 50104 will have the MED value set to 140, not 150.
- C. All prefixes, even the one originating from AS 50104 will be advertised to the 10.1.1.1

neighbor. The only difference with traffic originating from AS 50104 is that the MED values will be changed.

D, E. The default MED value is 0, not 100. The default local preference value is 100.

QUESTION 185

Assume that a BGP router has learned prefix 63.0.0.0/8 from two different BGP neighbors. Which statement regarding the BGP route selection process and how this route will be installed is correct?

- A. The update from the neighbor that has the highest weight and the highest local preference becomes the preferred path.
- B. The update from the neighbor that has the shortest AS path becomes the preferred path.
- C. The update from the neighbor that has the highest local preference and the highest MED becomes the preferred path.
- D. The update from the neighbor that has the lowest local preference becomes the preferred path.
- E. The update from the neighbor that has the highest MED becomes the preferred path.

Answer: A

Explanation:

BGP selects only one path as the best path. When the path is selected, BGP puts the selected path in its routing table and propagates the path to its neighbors. BGP uses the following criteria, in the order presented, to select a path for a destination:

1. If the path specifies a next hop that is inaccessible, drop the update.
2. Prefer the path with the largest weight.
3. If the weights are the same, prefer the path with the largest local preference.
4. If the local preferences are the same, prefer the path that was originated by BGP running on this router.
5. If no route was originated, prefer the route that has the shortest AS_path.
6. If all paths have the same AS_path length, prefer the path with the lowest origin type (where IGP is lower than EGP, and EGP is lower than Incomplete).
7. If the origin codes are the same, prefer the path with the lowest MED attribute.
8. If the paths have the same MED, prefer the external path over the internal path.
9. If the paths are still the same, prefer the path through the closest IGP neighbor.
10. Prefer the path with the lowest IP address, as specified by the BGP router ID.

Incorrect Answers:

B: Although this statement is correct, the weight and local preference values have a higher precedence than the AS path length.

C, E: The lowest MED is preferred, not the highest.

D: A higher local preference is preferred over a lower one.

QUESTION 186

The Certkiller network is using BGP for external routing. If a BGP router has more than one route to the same IP prefix, in what order are BGP attributes examined in

making a best path route selection?

- A. LOCAL_PREF, MED, AS_PATH, WEIGHT, ORIGIN
- B. WEIGHT, LOCAL_PREF, ORIGIN, AS_PATH; MED
- C. WEIGHT, LOCAL_PREF, AS_PATH, ORIGIN, MED
- D. WEIGHT; LOCAL_PREF, AS_PATH, MED, ORIGIN
- E. MED, LOCAL_PREF, WEIGHT, ORIGIN, AS_PATH

Answer: C

Explanation:

BGP assigns the first valid path as the current best path. It then compares the best path with the next path in list, until it reaches the end of the list of valid paths. The following is a list of rules used to determine the best path.

1. Prefer the path with the highest WEIGHT.

Note: WEIGHT is a Cisco-specific parameter, local to the router on which it's configured.

2. Prefer the path with the highest LOCAL_PREF. Note the following:

3. Prefer the path that was locally originated via a network or aggregate BGP subcommand, or through redistribution from an IGP.

4. Prefer the path with the shortest AS_PATH.

5. Prefer the path with the lowest ORIGIN type: IGP is lower than EGP, and EGP is lower than INCOMPLETE.

6. Prefer the path with the lowest multi-exit discriminator (MED). Note the following:

7. Prefer external (eBGP) over internal (iBGP) paths. If bestpath is selected, go to Step 9 (multipath).

8. Prefer the path with the lowest IGP metric to the BGP next hop. Continue, even if bestpath is already selected.

9. Check if multiple paths need to be installed in the routing table for BGP Multipath. Continue, if bestpath is not selected yet.

10. 1. When both paths are external, prefer the path that was received first (the oldest one).

2. Prefer the route coming from the BGP router with the lowest router ID. The router ID is the highest IP address on the router, with preference given to loopback addresses. It can also be set manually using the bgp router-id command.

3. If the originator or router ID is the same for multiple paths, prefer the path with the minimum cluster list length. This will only be present in BGP route-reflector environments. It allows clients to peer with RRs or clients in other clusters. In this scenario, the client must be aware of the RR-specific BGP attribute.

4. Prefer the path coming from the lowest neighbor address. This is the IP address used in the BGP neighbor configuration, and corresponds to the remote peer used in the TCP connection with the local router

Reference:

www.cisco.com/en/US/tech/CK365/technologies_tech_note09186a0080094431.shtml

QUESTION 187

The router CK1 is being configured for BGP, and the configuration will contain both IBGP and EBGP peers. Which statements regarding IBGP and EBGP neighbors are correct? (Select three)

- A. BGP updates from an IBGP peer are propagated to other IBGP and EBGP peers.
- B. BGP updates from an EBGP peer are propagated to other IBGP and EBGP peers.
- C. IBGP peers must be directly connected. If not, the IBGP-multihop option must be configured.
- D. EBGP peers must be directly connected; otherwise, the EBGP-multihop option must be configured.
- E. IBGP neighbors peering can be established using the loopback interface.
- F. EBGP neighbor peering must use the physical interface address to establish peering

Answer: B, D, E

Explanation:


When a BGP router receives a BGP routing update from an EBGP neighbor, the update is propagated to all IBGP neighbors. It is important to note that the same is not true for routing updates received via an IBGP neighbor, as these updates are not passed on to all IBGP peers. This is why IBGP speakers must be configured in a full mesh.

For EBGP peers, the best method is to use the directly connected interfaces as the peering IP addresses. If not, then EBGP multihop must be used. Multihop is used only in EBGP, not in IBGP.

It is recommended to use the loopback interface when configuring IBGP peers, since this interface is always up. For IBGP, the peering IP address needs to only be reachable via the IGP, so they do not need to be directly connected.

QUESTION 188

The Certkiller network is running BGP as displayed in the diagram below:



What path will routers Certkiller 1 and Certkiller 2 take to reach the 156.36.32.0/24 network in AS 5?

- A. Both will use the path through AS 1 due to Certkiller 1 having the shortest AS_PATH attribute.
- B. Certkiller 1 will use the path through AS 1 and Certkiller 2 will use the path through AS 3.
- C. Both will use the path through AS 1 due to Certkiller 1 having a lower LOCAL_PREF value.
- D. Both Certkiller 1 and Certkiller 2 will use the path through AS 3 due to Certkiller 2 having a higher LOCAL_PREF value.

Answer: D

Explanation:

BGP uses the following criteria, in the order presented, to select a path for a destination:

1. If the path specifies a next hop that is inaccessible, drop the update.
2. Prefer the path with the largest weight.
3. If the weights are the same, prefer the path with the largest local preference.
4. If the local preferences are the same, prefer the path that was originated by BGP running on this router.
5. If no route was originated, prefer the route that has the shortest AS_path.
6. If all paths have the same AS_path length, prefer the path with the lowest origin type (where IGP is lower than EGP, and EGP is lower than Incomplete).
7. If the origin codes are the same, prefer the path with the lowest MED attribute.
8. If the paths have the same MED, prefer the external path over the internal path.
9. If the paths are still the same, prefer the path through the closest IGP neighbor.
10. Prefer the path with the lowest IP address, as specified by the BGP router ID.

Based on the information above, the value of the Local Preference is considered before the length of the AS Path. When comparing the Local Preference value, the higher one is preferred.

QUESTION 189

Router CK 1 and CK2 are IBGP peers. Which BGP attributes are carried in all IBGP routing updates? (Select 3)

- A. MED
- B. Local Preference
- C. Weight
- D. Community
- E. AS-path
- F. Cost
- G. Origin

Answer: B, E, G


Explanation:

There are three well-known mandatory attributes. These must be included in updates propagated to all peers (both INGP and EBGp) and includes AS-PATH, NEXT-HOP and ORIGIN.

In addition to these three, all IBGP speakers must also carry the Local Preference information. The Local Preference is relevant when there is more than one path to a network outside of the current AS for instance if your network is connected to more than one ISP. Each of the routers that link to outside the AS can set a preference value for routes advertised into the AS, and this value indicates the router's preference for these routes. Only IBGP routers share the local preference values it does not leave the AS. The higher the value the more preferable the route is so if there are multiple paths to this network the route with the highest Local Preference is chosen and all traffic destined for the network is sent this way.

QUESTION 190

The Certkiller network resides in AS 200 as shown in the diagram below:



A BGP router receives updates for prefix 10.0.0.0/8 sourced from AS 100 from four different BGP neighbors. Certkiller 1 has the Local Preference of the prefix set to 50, while the other three neighbors do nothing with Loc Pref. Neighbor Certkiller 2 advertises the prefix with an AS path length of 3, while all other neighbors have an AS Path length of 2. The advertisement from neighbor Certkiller 3 has the origin code set to E, while the other have it set to I. And Neighbor Certkiller 4 does nothing to any of the attributes. What statement is true?

- A. Neighbor Certkiller 1 is the preferred path to prefix 10.0.0.0/8, since a higher local preference is better, and local preference is compared before the others.
- B. Neighbor Certkiller 2 is the preferred path to prefix 10.0.0.0/8, since a longer AS Path is better, and AS path is compared before the others.
- C. Neighbor Certkiller 3 is the preferred path to prefix 10.0.0.0/8, since an origin code of E is better than I, and origin code is compared before the others.
- D. Neighbor Certkiller 4 is the preferred path to prefix 10.0.0.0/8 only after neighbor Certkiller 2 dies.
- E. Neighbor Certkiller 3 is the preferred path to prefix 10.0.0.0/8 only after neighbor Certkiller 4 dies.

Answer: E

Explanation:

Based on the information provided, the route for 10.0.0.0/8 will be preferred from the following routers, in order:

1. Certkiller 4
2. Certkiller 3
3. Certkiller 2
4. Certkiller 1

BGP uses the following criteria, in the order presented, to select a path for a destination:

1. If the path specifies a next hop that is inaccessible, drop the update.
2. Prefer the path with the largest weight.
3. If the weights are the same, prefer the path with the largest local preference.
4. If the local preferences are the same, prefer the path that was originated by BGP

running on this router.

5. If no route was originated, prefer the route that has the shortest AS_path.

6. If all paths have the same AS_path length, prefer the path with the lowest origin type (where IGP is lower than EGP; and EGP is lower than Incomplete).

7. If the origin codes are the same, prefer the path with the lowest MED attribute.

Incorrect Answers:

A. The default local preference value is 100, so router Certkiller 1 will be the last router used because its local preference was set to 50.


B. A shorter AS path is preferred over a longer one.

C. An origin code of I (Internal) is preferred over an origin code of E (External).

D. Using the information given here, router Certkiller 4 will be preferred over all the others, and router Certkiller 2 will be used only if routers Certkiller 4 and Certkiller 3 both fail.

QUESTION 191

The Certkiller network is using AS 200 in the following BGP network:



Router Certkiller1 Configuration:

```
Certkiller1 (config)# router bgp 200
Certkiller1 (config-bgp)# neighbor 192.168.1.2 remote-as 200
Certkiller1 (config-bgp)# neighbor 192.168.1.2 next-hop-self
Certkiller1 (config-bgp)# neighbor 192.168.2.2 remote-as 200
```

Router Certkiller 1 receives an EBGp update containing 10.0.0.0/8 sourced from AS 300. Router Certkiller 1 then advertises 10.0.0.0/8 to routers Certkiller 2 and Certkiller 3 via IBGP. What does router Certkiller 3 use as a BGP next hop to reach network 10.0.0.0/8?

- A. 172.16.1.1
- B. 172.16.1.2
- C. 192.168.1.1
- D. 192.168.1.2
- E. 192.168.2.1

Answer: A

Explanation:

The EBGP next-hop attribute is the IP address that is used to reach the advertising router. For EBGP peers, the next-hop address is the IP address of the connection between the peers. For IBGP, the EBGP next-hop address is carried into the local AS, as illustrated below:

Figure39-5 BGP AS-path Attribute





Figure39-6 BGP Next-Hop Attribute



Router C advertises network 172.16.1.0 with a next hop of 10.1.1.1. When Router A propagates this route within its own AS, the EBGP next-hop information is preserved. If Router B does not have routing information regarding the next hop, the route will be discarded. Therefore, it is important to have an IGP running in the AS to propagate next-hop routing information.

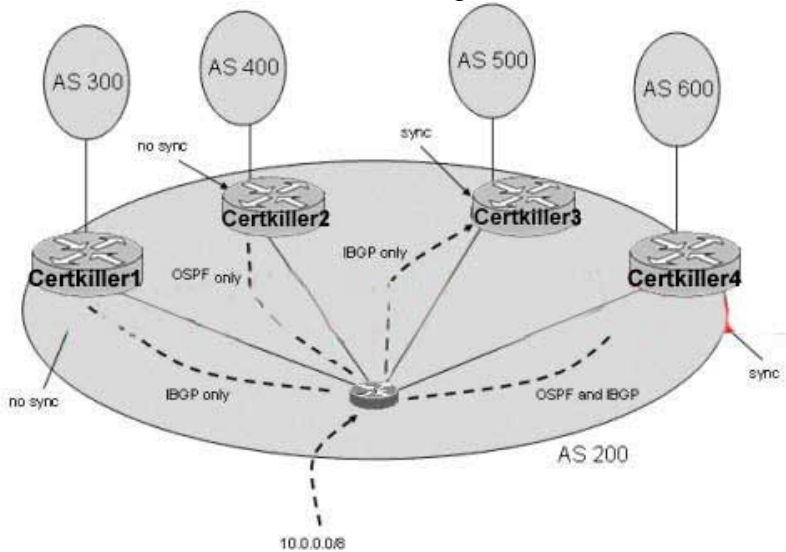
Incorrect Answers:

C: This would be the correct answer from router Certkiller 2's perspective, as the BGP next-hop-self configuration command was used for this peer. However, the next-hop-self command was not used for the Certkiller 3 peer, making the regular next hop rules apply. In order for Certkiller 1 to advertise itself as the next hop to all IBGP peers, it would need the "next-hop-self" command configured for all peers.

Reference: http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.htm

QUESTION 192

The Certkiller BGP network is using AS 200 as shown in the diagram below:



A router receives an EBGP update with prefix 10.0.0.0/8. This update is then forwarded to all BGP neighbors within its AS.

Which neighbors advertise 10.0.0.0/8 with EBGP updates of their own?

- A. Only router Certkiller 1 advertises 10.0.0.0/8 into AS 300.
- B. Both router Certkiller 1 and router Certkiller 2 advertise 10.0.0.0/8 into their respective neighbor ASs.
- C. Both router Certkiller 1 and router Certkiller 4 advertise 10.0.0.0/8 into their respective neighbor ASs.
- D. Both router Certkiller 2 and router Certkiller 4 advertise 10.0.0.0/8 into their respective neighbor ASs.
- E. Routers Certkiller 1, Certkiller 2, and Certkiller 4 advertise 10.0.0.0/8 into their respective neighbor ASs.

Answer: C

Explanation:

A BGP router with synchronization enabled will not advertise iBGP-learned routes to other eBGP peers if it is not able to validate those routes in its IGP. Assuming that IGP has a route to iBGP-learned routes, the router will announce the iBGP routes to eBGP peers. Otherwise the router treats the route as not being synchronized with IGP and does not advertise it. Disabling synchronization using the no synchronization command under router BGP prevents BGP from validating iBGP routes in IGP. By default, synchronization is on for all BGP routers.

In this example, Certkiller 1 will advertise this route to its EBGP peer due to the fact that synchronization is disabled. Although synchronization is enabled on router Certkiller 4, it will advertise the route because it is running both OSPF and BGP, so this route will match the corresponding route within the OSPF table and be advertised.

Incorrect Answers:

A. Both Certkiller 1 and Certkiller 4 will advertise this route.
B, C, D. Certkiller 2 will not advertise this route. Since it is not an IBGP peer, it will not receive the routing update in the first place so it will not be able to forward this route on to the other AS.

QUESTION 193

The Certkiller 1 BGP routing routes are displayed below:

```
CertKiller1#show ip route bgp
B    192.168.12.0/24 [200/0] via 2.2.2.2, 00:53:00
B    192.168.13.0/24 [200/0] via 2.2.2.2, 00:53:00
B    192.168.14.0/24 [200/0] via 2.2.2.2, 00:53:00
B    192.168.15.0/24 [200/0] via 2.2.2.2, 00:53:00
B    192.168.16.0/24 [200/0] via 2.2.2.2, 00:53:00
B    192.168.20.0/24 [200/0] via 2.2.2.2, 00:52:57
B    192.168.21.0/24 [200/0] via 2.2.2.2, 00:52:57
B    192.168.22.0/24 [200/0] via 2.2.2.2, 00:52:57
B    192.168.23.0/24 [200/0] via 2.2.2.2, 00:52:57
B    192.168.24.0/24 [200/0] via 2.2.2.2, 00:52:57

CertKiller1#show run
!
! Partial show run output
!
router ppp 65101
 aggregate-address 192.168.12.0 255.255.252.0 summary-only
 aggregate-address 192.168.20.0 255.255.252.0 as-set
 neighbor 2.2.2.2 remote-as 65101
 neighbor 2.2.2.2 update-source loopback0
 neighbor 2.2.2.2 next-hop-self
 neighbor 10.1.1.1 remote-as 65104
```

Based on the show ip route bgp output and the partial show run output shown, which BGP prefixes will be advertised by Certkiller 1 to the 10.1.1.1 neighbor?

- A. 192.168.12.0/22, 192.168.16.0/24, 192.168.20.0/22, 192.168.20.0/24, 192.168.21.0/24, 192.168.22.0/22, 192.168.23.0/24, 192.168.24.0/24
- B. 192.168.12.0/22, 192.168.20.0/22, 192.168.20.0/24, 192.168.21.0/24, 192.168.22.0/22, 192.168.23.0/24, 192.168.24.0/24
- C. 192.168.12.0/22, 192.168.20.0/22, 192.168.20.0/24, 192.168.21.0/24, 192.168.22.0/22, 192.168.23.0/24
- D. 192.168.12.0/22, 192.168.20.0/22, 192.168.20.0/24
- E. 192.168.12.0/22, 192.168.20.0/22
- F. All routes will be advertised, since there are no route filters in place.

Answer: A

Explanation:

When the aggregate-address command is used within BGP routing, the aggregated address is advertised, along with the more specific routes. The exception to this rule is through the use of the summary-only command. The "summary-only" keyword

suppresses the more specific routes and announces only the summarized route. Using the as-set argument creates an aggregate address with a mathematical set of autonomous systems (AS). This as-set summarizes the AS_PATH attributes of the all of the individual routes. This can be useful to avoid routing loops while aggregating routes. Again, unless the "summary-only" keyword is used with the as-set command the summary route is advertised along with the more specific routes. In the example above, the 192.168.12.0, 192.168.13.0, 192.168.14.0, and 192.168.15.0 networks will be summarized into the only 192.168.12/22 route, which will be advertised. Along with this one route, the others will also be advertised, as well as one additional 192.168.20.0/22 route. In total, 8 different routes will be advertised.

QUESTION 194

Many of the Certkiller BGP routers are configured using peer groups. Which of the following correctly display the common properties of BGP peer groups?

- A. Community values
- B. Inbound policies
- C. Outbound policies
- D. MED inbound policies
- E. Transitive AS policies
- F. None of the above

Answer: C

Explanation:

BGP neighbors who share the same outbound policies can be grouped together in what is called a BGP peer group. Instead of configuring each neighbor with the same policy individually, Peer group allows to group the policies which can be applied to individual peer thus making efficient update calculation along with simplified configuration.

Reference:

www.cisco.com/en/US/tech/CK365/technologies_tech_note09186a0080093fb7.shtml

QUESTION 195

You issue the "show ip bgp regexp [regexp]" command on router CK1 . This command is most useful when performing which type of BGP troubleshooting?

- A. To verify and troubleshoot BGP AS-Path filtering configurations
- B. To verify and troubleshoot BGP route-maps configurations
- C. To verify and troubleshoot BGP synchronization problems
- D. To verify and troubleshoot BGP Prefix-List filtering configurations
- E. To verify and troubleshoot BGP AS-Pathprepending configurations
- F. None of the above

Answer: A

Explanation:

A regular expression is a pattern to match against an input string. You specify the pattern that a string must match when you compose a regular expression. Matching a string to the specified pattern is called "pattern matching." Pattern matching either succeeds or fails. You can use regular expressions in the ip as-path access-list command with Border Gateway Protocol (BGP) to filter AS path information. To display routes matching the autonomous system path regular expression, use the "show ip bgp regexp" command in EXEC mode.

QUESTION 196

Router CK1 is configured for BGP on the Certkiller dual-homed network. Which BGP attributes are carried in every BGP update on this router (both IBGP and EBGP)?

- A. Origin, AS-Path, Next-Hop
- B. Router-ID, Origin, AS-Path
- C. Origin, Local-Preference, AS-Path
- D. Router-ID, Local-Preference, Next-Hop
- E. AS-Path, Local Preference, Next-Hop
- F. None of the above

Answer: A

Explanation:

Explanation:

AS-PATH, NEXT-HOP, and ORIGIN are all well known, mandatory BGP attributes, which is defined below:

Well known mandatory attributes: These attributes must be recognized by all BGP speakers, and must be included in all update messages. Almost all of the attributes impacting the path decision process, described in the next section, are well known mandatory attributes.

Origin Code

The ORIGIN is a well known mandatory attribute that indicates the origin of the prefix, or rather, the way in which the prefix was injected into BGP. There are three origin codes, listed in order of preference:

IGP, meaning the prefix was originated from information learned from an interior gateway protocol

EGP, meaning the prefix originated from the EGP protocol, which BGP replaced

INCOMPLETE, meaning the prefix originated from some unknown source

AS Path

The AS_PATH is a well-known mandatory attribute and is the list of all autonomous systems the prefixes contained in this update have passed through. The local autonomous system number is added by a BGP speaker when advertising a prefix to an eBGP peer.

Next Hop

The BGP NEXT_HOP is a well-known mandatory attribute. The Next Hop attribute is set when a BGP speaker advertises a prefix to a BGP speaker outside its local

autonomous system (it may also be set when advertising routes within an AS, this will be discussed in later sections). The Next Hop attribute may also serve as a way to direct traffic to another speaker, rather than the speaker advertising the route itself.

QUESTION 197

Part of a Certkiller router configuration is shown below:

```
Certkiller1#show route-map setweight
route-map test, permit, sequence 10
  Match clauses:
    ip address prefix-lists: filter
    as-path (as-path filter|: 1 2
  set clauses:
    weight 200
  Policy routing matches: 0 packets, 0 bytes
```

Based upon the "show route-map setweight" output shown above on this Certkiller BGP router, which matching routes will be set to a weight of 200?

- A. Routes that match the prefix-list named filter
AND
also match either the as-path filter 1 OR 2
- B. Routes that match the prefix-list named filter
OR
also match either the as-path filter 1 AND 2
- C. Routes that match the prefix-list named filter
AND
also match either the as-path filter 1 AND 2
- D. Routes that match the prefix-list named filter
OR
also match either the as-path filter 1 OR 2
- E. None of the above

Answer: A

Explanation:

When the match clauses are shown on different lines, then all of the match conditions must be met. In this example, both the IP prefix list named "filter" and the AS path filter must match in order to set the weight to 200 as shown. However, in this configuration, there are two AS path filters configured, numbered 1 and 2. In this case, only one of the two filters needs to be matched. If all three of the criteria had needed to be met, then there would be three distinct lines listed under the match clauses.

QUESTION 198

The Certkiller network is using BGP for Internet routing, and part of the router CK1 configuration is shown below:

```
router bgp 50101
  neighbor 10.1.1.1 remote-as 50102
  neighbor 10.2.2.2 remote-as 50103
  neighbor 10.2.2.2 route-map test2 out
  neighbor 10.1.1.1 route-map test out
!
ip as-path access-list 1 permit _50104$
ip as-path access-list 2 permit .*
!
route-map test permit 10
  match as-path 1
  set metric 150
!
route-map test permit 20
  match as-path 2
!
route-map test2 permit 10
  set metric 100
```

Based on the configuration of CK1 shown above, which statements are correct?
(Select all that apply)

- A. All prefixes originating in AS 50104 will be advertised to the 10.1.1.1 neighbor with a MED of 150.
- B. All prefixes not originating in AS 50104 will not be advertised to the 10.1.1.1 neighbor.
- C. All prefixes not originating in AS 50104 will be advertised to the 10.1.1.1 neighbor with a MED of 0.
- D. All prefixes will be advertised to the 10.2.2.2 neighbor with a MED of 100.
- E. All prefixes not originating in AS 50104 will be advertised to the 10.2.2.2 and the 10.1.1.1 neighbor with a MED of 100.
- F. None of the above

Answer: A, C, D

Explanation:

For the 10.1.1.1 BGP peer, route-map "test" is being applied. This route map has two statement entries. The first states that all traffic originating from AS 50104 (as shown by the "ip as-path access-list 1 permit _50104\$" command statement) should have the MED set to 150. The regular expression ".*" matches everything else, so all other traffic is to be routed normally. For the 10.2.2.2 neighbor, the route-map test2 is being applied, which sets the metric (MED) to 100. Since the default MED value is 0, all other traffic not originating in AS will be advertised to the 10.1.1.1 peer with a MED of 0.

Incorrect Answers:

B: All prefixes that do not match the route-map will be advertised normally with default

settings.

D: The default MED value is 0, not 100. The default local preference value is 100.

QUESTION 199

In a BGP peering relationship with a Certkiller customer where routing information is exchanged by router CK1, which prefix list filter(s) will ensure that only subnets in the class-B address space are accepted by this router?

- A. ip prefix-list list-A permit 191.0.0.0/3 le 16
- B. ip prefix-list list-B permit 0.0.0.0/0 ge 16 le 24
- C. ip prefix-list list-C permit 128.0.0.0/2 ge 17
- D. ip prefix-list list-D permit 0.0.0.0/0 ge 16
ip prefix-list list-D permit 0.0.0.0/0 le 23
- E. ip prefix-list list-E permit 128.0.0.0/1 ge 16
ip prefix-list list-E permit 191.0.0.0/3 le 23

Answer: E

Explanation:

In a prefix list configuration, the "ge" keyword means greater than or equal to, while the "le" keyword means less than or equal to. Choice E correctly describes the two statements that are needed. The first line specifies that any route larger than 128.0.0.0/1 with a prefix range greater than or equal to 16 will match the filter. The second line specifies that any route less than 191.0.0.0/3 with a network mask of less than or equal to 23 will also match. Therefore, only addresses that fall in the class B range will pass through the filter.

Incorrect Answers:

- A. This will allow all class A and B networks to pass through.
 - B. This will permit address space from 16 to 24 bits in length from all network class ranges from passing through the filter.
 - C. This will allow all 128.0.0.2 prefixes with network masks greater than or equal to 17 bits in length. It is not restrictive enough to allow only class B networks.
 - D. This will allow all routes (from every network class) with network masks of between 16 and 23 bits in length.
-

QUESTION 200

Certkiller uses BGP in their multi-homed network. Which of the following are considered to be attributes of BGP routes? (Choose all that apply)

- A. Origin
- B. Weight
- C. Local Preference
- D. Community
- E. Cluster List

Answer: A, C, D, and E

Explanation:

Origin, Local Preference, Community, and Cluster List are all BGP attributes.

ORIGIN Well-known mandatory, Type code 1 RFC 1771

LOCAL_PREF Well-Known discretionary, Type code 5 RFC 1771

COMMUNITY Optional transitive, Type 8 RFC 1997


CLUSTER_LIST Optional nontransitive, Type code 10 RFC 1966

Incorrect Answers:

B. Cisco routers do indeed use weight during the BGP route decision making process. In fact, it is the first parameter that is looked at. However, weight is a Cisco-only parameter, and is therefore not considered a BGP attribute.

QUESTION 201

The Certkiller EIGRP network is displayed below:



You work as a network engineer at Certkiller .com. Please study the exhibit carefully. In the illustrated network, all routers are configured to run EIGRP on all links. Router Certkiller 2 is configured to only send a summary route to Certkiller 4, Certkiller 5 and Certkiller 6. If the link between Certkiller 1 and Certkiller 2 fails, what is the maximum number of queries Certkiller 3 will receive for 192.168.1.0/24, given that all the packets transmitted during convergence are transmitted once (there are no dropped or retransmitted packets)? Will Certkiller 4, Certkiller 5 or Certkiller 6 receive any queries?

- A. Certkiller 3, Certkiller 4, Certkiller 5, and Certkiller 6 won't receive any queries for 192.168.1.0/24, since there's no alternate path to this destination within the network
- B. Certkiller 3 will receive one query for 192.168.1.0/24. Certkiller 4, Certkiller 5 and Certkiller 6 will not receive any queries for this destination because Certkiller 2 isn't advertising this network towards them
- C. Certkiller 3 will receive one query for 192.168.1.0/24, from Certkiller 2. Certkiller 4, Certkiller 5 and Certkiller 6 will each receive and reply to, one query each
- D. Certkiller 3 will receive 4 queries for 192.168.1.0/24, one each from Certkiller 2, Certkiller 4, Certkiller 5 and Certkiller 6. Certkiller 4, Certkiller 5 and Certkiller 6 will each

receive and reply to one query each
E. None of the above

Answer: C

Explanation:

Queries and replies are sent when destinations go into Active state. Queries are always multicast unless they are sent in response to a received query. In this case, it is unicast back to the successor that originated the query. Replies are always sent in response to queries to indicate to the originator that it does not need to go into Active state because it has feasible successors. Replies are unicast to the originator of the query. Both queries and replies are transmitted reliably.

Reference:

http://www.cisco.com/en/US/tech/CK365/technologies_tech_note09186a0080093f07.shtml

QUESTION 202

An EIGRP multicast flow timer is defined as which of the following?

- A. The timeout timer after which EIGRP retransmits to the neighbor in non CR mode, through unicasts.
- B. The time interval that EIGRP hello packets are sent.
- C. The timer after which EIGRP will not forward multicast data traffic.
- D. The timer interval between consecutive transmitted EIGRP hello intervals.
- E. The timeout timer after which EIGRP retransmits to the neighbor in CR mode, through unicasts.
- F. None of the above.

Answer: E

Explanation:

After pair of routers become neighbors, they will send routing updates (and other packets) to one another using a reliable multicast scheme. For example, if router one has a series of packets which must be transmitted to routers two, three, and four such as a routing table update, it will send the first packet to the EIGRP multicast address, 224.0.0.10, and wait for an acknowledgment from each of its neighbors on its Ethernet interface (in this case, routers two, three and four). Let's assume that routers two and four do answer, but router three does not.

Router one will wait until the multicast flow timer expires on the Ethernet interface, then send out a special packet, a sequence TLV, telling router three not to listen to any further multicast packets from router one, then it will continue transmitting the remainder of the update packets as multicast to all other routers on the network. The sequence TLV indicates an out-of-sequence multicast packet. Those routers not listed in the packet enter Conditional Receive (CR) mode, and continue listening to multicast. While there are some routers in this mode, the Conditional Receive bit will be set in multicast packets. In this case, router one will send out a sequence TLV with router three listed, so routers two and four will continue listening to further multicast updates.

QUESTION 203

Which components are factored in by default when an EIGRP metric is calculated?
(Choose all that apply)

- A. MTU
- B. Delay
- C. Load
- D. Bandwidth
- E. Reliability

Answer: B, D

Explanation:

By default, EIGRP uses only bandwidth and Delay when calculating the metric. EIGRP uses these scaled values to determine the total metric to the network:

1.
$$\text{metric} = [K1 * \text{bandwidth} + (K2 * \text{bandwidth}) / (256 - \text{load}) + K3 * \text{delay}] * [K5 / (\text{reliability} + K4)]$$

The default values for K are:

1. K1 = 1
2. K2 = 0
3. K3 = 1
4. K4 = 0
5. K5 = 0

For default behavior, you can simplify the formula as: Metric = Bandwidth + Delay

Incorrect Answers:

A. The MTU is tracked but never used in calculating the metric for IGRP or EIGRP at any time.

C, E. Although Load and Reliability are K values that can indeed be factored into the metric, by default their K value is 0 so they are not used.

Reference:

<http://www.cisco.com/warp/public/103/eigrp-toc.html#eigrpmetrics>

QUESTION 204

The topology of a network changes causing an EIGRP router to go into the active state. The DUAL process shows a new route that meets the EIGRP Feasibility Condition. In regards to this specific route, which of the following is true?

- A. The Feasible Distance of the new route must be equal to one.
- B. The Feasible Distance of the new route must be higher than one.
- C. The Reported Distance of the new route must be equal to Feasible Distance.
- D. The Reported Distance of the new route must be higher than Feasible Distance.
- E. The Reported Distance of the new route must be lower than Feasible Distance.

Answer: E

Explanation:

The following are some terms relating to EIGRP:

1. Feasible Distance: The lowest calculated metric to each destination
2. Feasibility Condition: A condition that is met if a neighbor's advertised distance to a destination is lower than the router's Feasible Distance to that same destination.
3. Successor: The neighbor that has been selected as the next hop for a given destination based on the Feasibility Condition.

Reference:

Jeff Doyle, Routing TCP/IP, Volume I, Chapter 8: Enhanced Interior Gateway Routing Protocol (EIGRP), p.336-337, Cisco Press, (ISBN 1-57870-041-8)

Incorrect Answers:

A: The metric of the new route needs only to be less than the current metric to the destination (feasible distance), and does not necessarily need to equal one.

B: It is feasible that the new metric to the destination could equal one, and also be lower than the current metric.

C, D: The reported distance must be lower than the feasible distance.

Additional info:

The Feasible Condition is met when the receiving router has a Feasible Distance (FD) to a particular network and it receives an update from a neighbor with a lower advertised or Reported Distance (RD) to that network. The neighbor then becomes a Feasible Successor (FS) for that route because it is one hop closer to the destination network.

There may be a number of Feasible Successors in a meshed network environment.

The RD for a neighbor to reach a particular network must always be less than the FD for the local router to reach that same network. In this way EIGRP avoids routing loops. This is why routes that have RD larger than the FD are not entered into the Topology table.

Reference:

Ravi Malhotra, IP Routing, Chapter 4: Enhanced Interior Gateway Routing Protocol (EIGRP), O'Reilly Press, January 2002 (ISBN 0-596-00275-0)

QUESTION 205

Which of the following EIGRP packets require an acknowledgement? (Choose all that apply)

- A. Hello
- B. Query
- C. Reply
- D. Update
- E. Ack
- F. None of the above

Answer: B, C, D

Explanation:

Updates are used to convey reachability of destinations. When a new neighbor is

discovered, update packets are sent so the neighbor can build up its topology table. In this case, update packets are unicast. In other cases, such as a link cost change, updates are multicast. Updates are always transmitted reliably.

Queries and replies are sent when destinations go into Active state. Queries are always multicast unless they are sent in response to a received query. In this case, it is unicast back to the successor that originated the query. Replies are always sent in response to queries to indicate to the originator that it does not need to go into Active state because it has feasible successors. Replies are unicast to the originator of the query. Both queries and replies are transmitted reliably.

EIGRP reliable packets are: Update, Query and Reply.

EIGRP unreliable packets are: Hello and Ack.

Incorrect Answers:

A, E. Hellos are multicast for neighbor discovery/recovery. They do not require acknowledgment. A hello with no data is also used as an acknowledgment (ack). Acks are always sent using a unicast address and contain a non-zero acknowledgment number.

Reference: Cisco BSCN version 1.0 study guide, pages 6-18.

QUESTION 206

Which of the following types of EIGRP packets contain the Init flag?

- A. Hello/Ack
- B. Query
- C. Reply
- D. Update
- E. None of the above

Answer: D

Explanation:

In EIGRP header there is an 8-bit flag value. The rightmost bit is init.

Which when set to 0x00000001 indicates that the enclosed route entries are the first in a new neighbor relationship.

Also the route entries are carried in update packet not hello packet.

Additional Info:

The following debug output displays the Init Sequence increasing only with the update packet.

Router# debug eigrp packetEIGRP: Sending HELLO on Ethernet0/1 AS 109, Flags 0x0,

Seq 0, Ack 0EIGRP: Sending HELLO on Ethernet0/1 AS 109, Flags 0x0, Seq 0, Ack

0EIGRP: Sending HELLO on Ethernet0/1 AS 109, Flags 0x0, Seq 0, Ack 0

EIGRP: Received UPDATE on Ethernet0/1 from 192.195.78.24,

AS 109, Flags 0x1, Seq 1, Ack 0EIGRP: Sending HELLO/ACK on

Ethernet0/1 to 192.195.78.24, AS 109, Flags 0x0, Seq 0, Ack

1EIGRP: Sending HELLO/ACK on Ethernet0/1 to 192.195.78.24,

AS 109, Flags 0x0, Seq 0, Ack 1EIGRP: Received UPDATE on

Ethernet0/1 from 192.195.78.24, AS 109, Flags 0x0, Seq 2,

Ack 0Incorrect Answers:

A. Hellos are multicast for neighbor discovery/recovery. They do not require acknowledgment. A hello with no data is also used as an acknowledgment (ack). Acks are always sent using a unicast address and contain a non-zero acknowledgment number.
B, C. Queries and replies are sent when destinations go into Active state. Replies are always sent in response to queries to indicate to the originator that it does not need to go into Active state because it has feasible successors. Replies are unicast to the originator of the query. Both queries and replies are transmitted reliably.
Reference: "Routing TCP/IP" Jeff Doyle Pg364

QUESTION 207

In your EIGRP network you notice that the neighbor relationship between two of your routers was recently restarted. Which of the following could have occurred to have caused this? (Choose all that apply)

- A. The clear ip route command was issued.
- B. The ARP cache was cleared.
- C. The IP cache was cleared.
- D. An update packet with Init flag set from a known, already established neighbor relationship was received by one of the routers.
- E. The IP EIGRP neighbor relationship was cleared manually.

Answer: D, E

Explanation:

D as well as E will result in EIGRP relationship to be restarted.

The reason for D: If a router receives an update packet with the init flag set it clearly implies that this packet is the first after a new neighbor relationship has been established.

The reason for E: If we clear the IP EIGRP neighbor relationship it will automatically result in EIGRP neighbor relationship to be restarted.

Incorrect Answers:

- A. This will clear the IP routing table, but will not have any affect on the EIGRP neighbor relationship.
 - B. This will only clear the MAC address learned ARP cache.
 - C. This also will not have any affect on the EIGRP neighbor relationship.
-

QUESTION 208

The Certkiller EIGRP network has a router named Router CK2 . Router CK2 is connected to an EIGRP neighbor, CK1 . CK1 is defined as a stub. With regard to this network, which of the following are true?

- A. Router CK1 will not advertise any network routes to CK2 .
- B. Router CK2 will send only summary routes to CK1 .
- C. Router CK2 will not query CK1 about any internal route.
- D. Router CK2 will not query CK1 about any external route.

- E. Router CK2 will not query CK1 about any route.
- F. None of the above.

Answer: E

Explanation:

E is the best choice, as an EIGRP router will not query a stub neighbor about any route.

Incorrect Answers:

- A. CK1 will still be required to advertise its network routes to the neighbor, even though it is configured as a stub.
- B. CK2 still sends all routes to CK1 .
- C, D. Although both of these are true, since CK2 will not query CK1 about any route, E is a better choice.

Reference:

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120s/120s15/eigrpstb.htm>

QUESTION 209

How is the metric for a summarized route derived when the interface summary command for EIGRP is used?

- A. It is derived from the route that has the biggest metric.
- B. It is derived from the route that has the smallest metric.
- C. It is derived from the interface that has the summary command configured on it.
- D. It is derived from the route that has the shortest matching mask.
- E. It is derived from the default-metric.

Answer: B

Explanation:

According to Cisco's EIGRP design guide, "The metric is the best metric from among the summarized routes."


Reference:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fipr_c/ipcp2/1cfeigrp.htm#1001078

"...EIGRP will advertise the summary address out the interface with a metric equal to the minimum of all more specific routes."

QUESTION 210

The Certkiller EIGRP network topology is displayed below, along with the EIGRP metric values for each link:



From the perspective of router A shown above, which of the following routers would be considered the successor and the feasible successors to the 172.16.9.0/24 network? (Select two choices below)

- A. B is the successor
- B. C is the successor
- C. D is the successor
- D. B is a feasible successor
- E. C is a feasible successor
- F. D is a feasible successor.
- G. E is a feasible successor

Answer: A, E

Explanation:

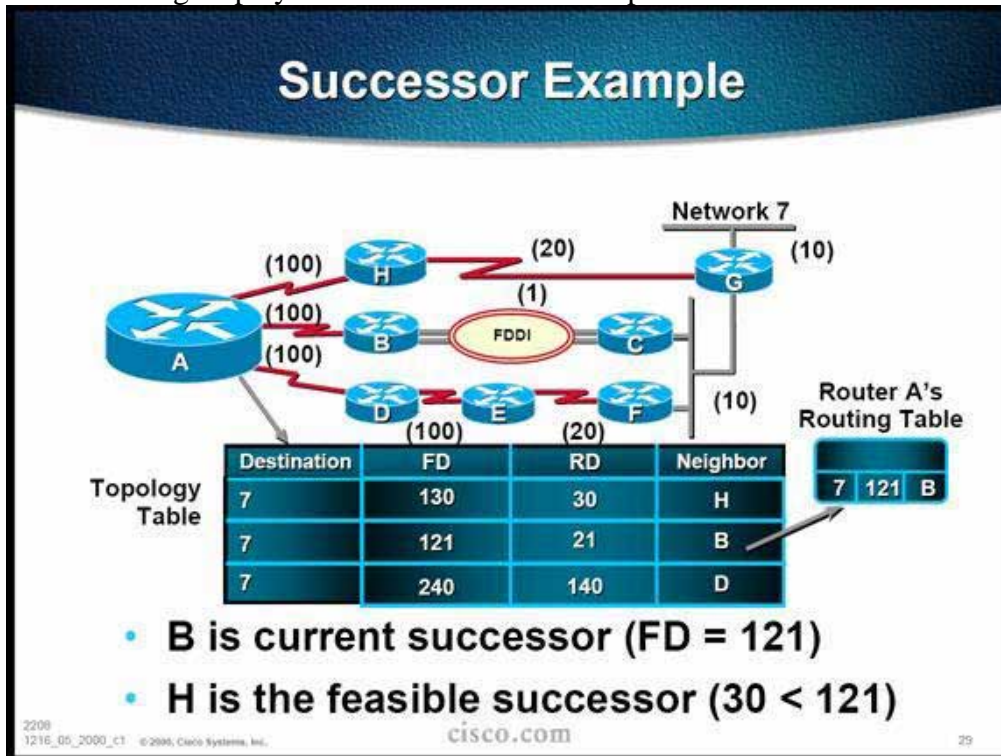
The following are some terms relating to EIGRP:

1. Feasible Distance: The lowest calculated metric to each destination
2. Feasibility Condition: A condition that is met if a neighbor's advertised distance to a destination is lower than the router's Feasible Distance to that same destination.
3. Successor: The neighbor that has been selected as the next hop for a given destination based on the Feasibility Condition.
4. Feasible Successor: A neighbor whose Reported Distance (RD) is less than the Feasible Distance (FD).

The Feasible Condition is met when the receiving router has a Feasible Distance (FD) to a particular network and it receives an update from a neighbor with a lower advertised or Reported Distance (RD) to that network. The neighbor then becomes a Feasible

Successor (FS) for that route because it is one hop closer to the destination network. There may be a number of Feasible Successors in a meshed network environment. The RD for a neighbor to reach a particular network must always be less than the FD for the local router to reach that same network. In this way EIGRP avoids routing loops. This is why routes that have RD larger than the FD are not entered into the Topology table. In this example, Router B would be the successor, with a feasible distance of 7680 (2560+2560+2560). Therefore, only routers with an AD of less than 7680 will become successors. In this case, router C will have an Advertised Distance of 6400 so it is a FS. Router D has a RD of 10240, and since it must be less than the current FD, it will not become a FS.

The following display further describes an example of a Feasible Successor:




Reference:

Jeff Doyle, Routing TCP/IP, Volume I, Chapter 8: Enhanced Interior Gateway Routing Protocol (EIGRP), p.336-337, Cisco Press, (ISBN 1-57870-041-8)

QUESTION 211

The Certkiller EIGRP network is displayed in the following diagram:



The associated EIGRP metric is listed as shown above for each of the links. Based on this information, what is the reported distance (advertised distance) to network 172.16.9.0/24 from router C to router A?

- A. 5120
- B. 6400
- C. 17,920
- D. 10,240
- E. 11,520
- F. 2560
- G. None of the above

Answer: B

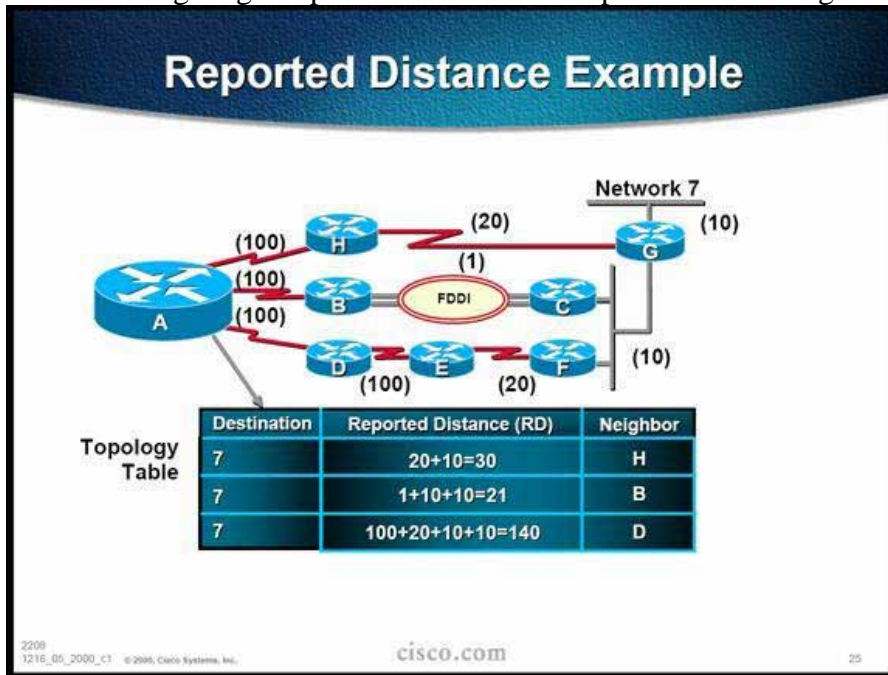
Explanation:

The following are some terms relating to EIGRP:

1. Feasible Distance: The lowest calculated metric to each destination
2. Feasibility Condition: A condition that is met if a neighbor's advertised distance to a destination is lower than the router's Feasible Distance to that same destination.
3. Successor: The neighbor that has been selected as the next hop for a given destination based on the Feasibility Condition.
4. Feasible Successor: A neighbor whose Reported Distance (RD) is less than the Feasible Distance (FD).


In the example above, the Feasible Distance to network 172.16.9.0/24 from C to A would be the distance that this network is from router C. In this case, the distance is $2560 + 3840 = 6400$, so Choice B is correct.

The following diagram provides another example for calculating the RD:



QUESTION 212

The Certkiller EIGRP network is displayed in the exhibit below:



The "Show IP EIGRP neighbor" command is issued on the router

A. Router A is

configured with the default EIGRP settings. After issuing this command, which of the following answer choices correctly describe the expected output?

A. routerA#show ip eigrp neighbor

IP-EIGRP neighbors for process 1

H Address Interface Hold Uptime SRTT RTO Q Seq

(Sec) (ms) Cnt Num

1 10.1.1.2 Et1 13 12:00:53 12 300 0 620

0 10.1.2.2 S0 174 12:00:56 17 200 0 645

B. routerA#show ip eigrp neighbor

IP-EIGRP neighbors for process 1

H Address Interface Hold Uptime SRTT RTO Q Seq

(Sec) (ms) Cnt Num

1 10.1.1.2 Et1 20 12:00:53 12 300 0 620


```
0 10.1.2.2 S0 190 12:00:56 17 200 0 645
C. routerA#show ip eigrp neighbor
IP-EIGRP neighbors for process 1
H Address Interface Hold Uptime SRTT RTO Q Seq
(Sec) (ms) Cnt Num
1 10.1.1.2 Et1 174 12:00:53 12 300 0 620
0 10.1.2.2 S0 13 12:00:56 17 200 0 645
D. routerA#show ip eigrp neighbor
IP-EIGRP neighbors for process 1
H Address Interface Hold Uptime SRTT RTO Q Seq
(Sec) (ms) Cnt Num
1 10.1.1.2 Et1 185 12:00:53 12 300 0 620
0 10.1.2.2 S0 19 12:00:56 17 200 0 645
```

Answer: A

Explanation:

The value in the Hold column of the command output should never exceed the hold time, and should never be less than the hold time minus the hello interval (unless, of course, you are losing hello packets). If the Hold column usually ranges between 10 and 15 seconds, the hello interval is 5 seconds and the hold time is 15 seconds. If the Hold column usually has a wider range - between 120 and 180 seconds - the hello interval is 60 seconds and the hold time is 180 seconds.

The EIGRP default timer settings are:

Hello Interval: 5 seconds for all high speed links

60 seconds for low speed links (T1 or less)

The default hold timer is less 3 times the hello interval. Since this question tells us that the default values are used, the Router A would have a value of not more than 15 seconds for the Ethernet peer and 180 seconds for the serial peer, so choice A is correct.

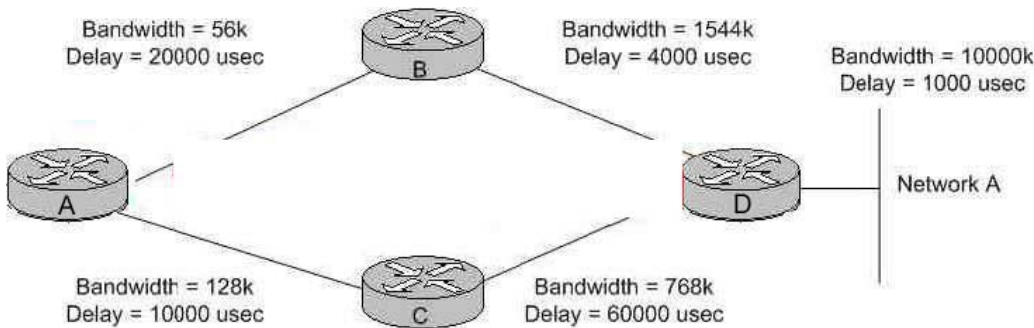
Incorrect Answers:

B, D. The timers for both the Ethernet and serial peers are above the maximum theoretical values for a working EIGRP network, assuming that the default timers are being used.

C. The timer values in this choice is wrong. High speed links such as ethernet use a shorter hello interval than low speed T1 links.

QUESTION 213

The Certkiller EIGRP network, along with the configured bandwidth statements of the routers and the delay of each link, is shown below:



Assuming that all EIGRP routers are all using default configurations, what path would router A choose to route packets to network A?

- A. Router A takes the path through router B.
- B. Router A takes the path through router C.
- C. Router A would load balance to both router B and router C.
- D. Neither path would be chosen as there is a loop in the network.
- E. The metrics shown are too large, and the route to network A would be considered unreachable.

Answer: B

Explanation:

When all 5 of the K values are set to the default values, the EIGRP metric calculation for each link is found by the default formula of $(\text{Bandwidth} + \text{Delay}) \times 256$. The metric calculation is the same as IGRP, but the result is multiplied by 256 for finer granularity. In this case, the bandwidth component is found in the same way as the OSPF metric, which is $10,000,000/\text{bandwidth}$. It is important to note that only the minimum outgoing bandwidth is used, so along any path from A to Z, the slowest link among all the hops is used as the chosen metric for the bandwidth portion. This is true for both IGRP and EIGRP (See Routing TCP/IP by Jeff Doyle, page 243-244). The delay metric is found by adding the total delay of the path (in microseconds) and dividing by 10.

For this question, the shortest path can be found by comparing the two different choices that we really have (through router B or router C). For the path through router B the bandwidth metric is:

$$(10 \text{ million}/56) + (24000/10) \times 256 = (178571 + 2400) \times 256 = 46328683.$$

For the path going through router C:

$$(10 \text{ million}/128) + (70000/10) \times 256 = (78125 + 7000) \times 256 = 21792000.$$

Note that only the lowest bandwidth metric was used along the entire path, where as the delay was added at each hop. Also note that the calculation for the path from router D to network A was omitted, since this value would be simply added to the metrics above would not change the answer.

Incorrect Answers:

- A. The path through this router has a higher metric and so would not be used.
- C. By default, EIGRP would load balance over equal cost paths. Although these paths are not equally valued, load balancing could occur despite this if the "variance" EIGRP feature was used. However, the variance command is not enabled by default.

D. Although a loop does exist, EIGRP routers maintain loop avoidance techniques, including keeping track of hop counts used. For IGRP and EIGRP, there is a maximum hop count of 100 hops.

E. Both of the metric listed above are well within the maximum limits set by EIGRP.

QUESTION 214

The Certkiller network is using EIGRP as the routing protocol, and the EIGRP topology information for router R4 is displayed below:

```
R4#sh ip eigrp top
IP-EIGRP Topology Table for AS(10)/ID(140.140.3.1)

Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia status

P 140.140.0.0/16, 1 successors, FD is 128256
   via Summary (128256/0), Null0
P 170.170.0.0/16, 1 successors, FD is 2809856
   via 116.16.34.3 (2809856/2297856), Serial1/0
P 190.190.0.0/16, 1 successors, FD is 2297856
   via 116.16.34.9 (2297856/128256), Serial1/0
P 130.130.0.0/16, 1 successors, FD is 2297856
   via 116.16.34.3 (2297856/128256), Serial1/0
P 140.140.1.0/24, 1 successors, FD is 128256
   via Connected, Loopback1
P 116.16.37.0/30, 1 successors, FD is 2681856
   via 116.16.34.3 (2681856/2169856), Serial1/0
P 116.16.34.0/23, 1 successors, FD is 2169856
   via Connected, Serial1/0
P 116.0.0.0/8, 1 successors, FD is 2169856
   via Summary (2169856/0), Null0
```

Based on the information above, which of the following statements is true?

- A. The routers 116.16.34.3 and 116.16.34.9 are EIGRP neighbors to CK4 .
- B. The 116.16.37.0 network is reachable via the 116.16.34.9 interface.
- C. A static route has been configured to summarize the 140.0.0.0 network and route it to the NULL 0 interface.
- D. Interface serial 1/0 is most likely a frame relay interface with four DLCIs: one to the 170.170.0.0 network, one to the 130.130.0.0 network and one to the 116.16.37.0 network.
- E. All of the above

Answer: A

Explanation:

The IP address following the "via" entry is the peer that told the software about this destination. When issuing this command, the first n of these entries, where N is the number of successors, are the current successors. The remaining entries on the list are feasible successors. In the example above, the router CK4 is learning routes from both of these two peers, so they are EIGRP neighbors to CK4 .

Incorrect Answers:

- B. This network is reachable via the 116.16.34.3 neighbor, not 116.16.34.2.

- C. The routing entry for the 140.0.0.0/8 network is known via the summary, not a static route. EIGRP uses auto-summarization by default, which has produced this route.
- D. All three of these networks are known via the same IP peer. Although it is possible that 4 separate PVC's are built to the same IP address peer, there is no reason to assume that this is the case in this example. It actually looks like there may be 3 total DLCIs on this serial interface, not 4.

QUESTION 215

The Certkiller WAN is displayed in the diagram below, along with the partial configuration files of routers CK1 and CK2 :



```
hostname CK1
!
interface Ethernet0/0
 ip address 172.16.1.10 255.255.255.0
!
interface Serial6/0
 ip address 192.168.1.5 255.255.255.252
!
router eigrp 10
 network 172.16.0.0
 network 192.168.1.0
```

```
hostname CK2
!
interface Ethernet0/0
 ip address 172.17.1.10 255.255.255.0
!
interface Serial6/0
 ip address 192.168.1.16 255.255.255.252
!
router eigrp 11
 network 172.17.0.0
 network 192.168.1.0
```

Based on the above information, what would be the most likely reason that the routing tables do not contain routes for each of the remote networks?

- A. The routers interfaces are not functioning properly.
- B. IP routing is not enabled on the routers.
- C. The routers are not members of the same autonomous system.
- D. The routers only pass locally significant routing information.
- E. The routers are using different routing protocols.
- F. Auto-summarization is not disabled on the routers.

Answer: C


Explanation:

The number following the "router eigrp" command is known as the process ID, and is used to denote the Autonomous System of the network that the router is in. The process ID can be any number between 1 and 65535 (0 is not allowed) and it can be randomly chosen, as long as it is the same for all EIGRP processes in routers that are to share the routing information. In the example above, router CK1 is configured to use EIGRP process 10, while router CK2 is using EIGRP process 11

Reference: Jeff Doyle, "Routing TCP/IP volume 1" page 377.

QUESTION 216

Two Certkiller routers are configured for EIGRP as shown in the following exhibit:



Hostname Certkiller1

```
!
interface Ethernet0/0
 bandwidth 10000
 ip address 192.168.1.17 255.255.255.252
 ip address 209.224.111.41 255.255.255.252 secondary
!
router eigrp 200
 network 192.168.1.0
 network 209.224.111.0
```

Hostname Certkiller2

```
!
interface Ethernet0/0
 bandwidth 1000
 ip address 209.224.111.42 255.255.255.252
!
router eigrp 100
 network 209.224.111.0
 no auto summary
```

Router Certkiller 1 and router Certkiller 2 are unable to form an EIGRP neighbor relationship. What are two reasons for this problem? (Select two).

- A. The bandwidth settings on the interfaces do not match.
- B. The routers belong to different autonomous systems.
- C. EIGRP can not form a neighbor relationship using secondary addresses.
- D. The network statement under router EIGRP does not match the subnetted network configured on the Ethernet interface.
- E. Auto summarization has not been correctly configured on router Certkiller 1.

Answer: B, C

Explanation:

EIGRP, unlike OSPF, checks for the Autonomous System number on neighboring routers before becoming neighbors. EIGRP will only form a neighbor relationship with other routers in the same AS.

Since EIGRP always sources data packets from the primary address, Cisco recommends that you configure all routers on a particular subnet with primary addresses that belong to the same subnet. Routers do not form EIGRP neighbors over secondary networks.

Therefore, if all routers' primary IP addresses do not agree, problems can arise with neighbor adjacencies.

Incorrect Answers:

- A. Manually setting the bandwidth will affect the overall metric of the individual EIGRP routes, but will not affect the state of the neighbor relationship.
- D. Although they do not match, EIGRP will still work as long as the EIGRP interface is included within the subnet mask used for the EIGRP process.
- E. In EIGRP, automatic summarization is on by default. Whether this is enabled or disabled will have no effect on the neighbor relationship.

QUESTION 217

Certkiller .com is designing a large network with core, distribution, and access layers. EIGRP is the routing protocol that will be used throughout the network. Each distribution router has WAN connectivity to at least 20 access routers. Every router in the network

has an explicit route to every possible subnet. All hosts in the network should be able to reach any other host, anywhere within the network. What should be done to optimize the routing configuration?

- A. Ensure IP address space is allocated so that routes can be summarized at the core routers.
- B. Filter routes in the distribution layer so that every access router doesn't have an explicit route to every subnet.
- C. Filter routes in the access layer so that every access router doesn't have an explicit route to every subnet.
- D. Ensure IP address space is allocated so that routes can be summarized at each distribution router.

Answer: D

Explanation:


The best way to reduce the number of routes within the routing table is via route summarization. In order to optimize the network using this approach, summarization should take place on the distribution layer of a three tiered network design. When doing this, it is important to ensure that proper planning takes place to ensure that enough IP address space is allocated at each distribution router, in order to summarize all of the remote locations into one single network route.

Incorrect Answers:

- A. Core routers should focus solely on routing data packets as quickly as possible. The use of any ancillary technologies such as access lists, packet classification, and route filtering. These technologies are best suited to be placed on the distribution layer network devices.
- B, C: Either of these choices could result in some hosts becoming unreachable from other hosts within the Certkiller network.

QUESTION 218

The Certkiller frame relay network is displayed below:



Assume that no subinterfaces are used on router Certkiller 1 and EIGRP is the routing protocol in use. What is the effect on routing updates if router Certkiller 1 learns about network A from router Certkiller 2, assuming default configurations are being used?

- A. Router Certkiller 1 will advertise the route to network A to router Certkiller 3.

- B. Router Certkiller 1 will advertise the route to network A to router Certkiller 2 and Certkiller 3.
- C. Router Certkiller 1 will load balance between router Certkiller 2 and router Certkiller 3.
- D. Router Certkiller 1 will not advertise the route to network A to router Certkiller 3.

Answer: D

Explanation:

Split horizon controls the sending of IP Enhanced IGRP update and query packets. When split horizon is enabled on an interface, these packets are not sent for destinations for which this interface is the next hop. This reduces the possibility of routing loops.

By default, split horizon is enabled on all interfaces.

Split horizon blocks route information from being advertised by a router out of any interface from which that information originated. This behavior usually optimizes communications among multiple routing devices, particularly when links are broken. However, with nonbroadcast networks (such as Frame Relay and SMDS), situations can arise for which this behavior is less than ideal. In the example above, when router Certkiller 1 sees the route advertisement from Certkiller 2, it will not advertise this route out the same interface, which is also the connection to Certkiller 3. If we were to use sub-interfaces or disable split horizons on Certkiller 1, then it would indeed advertise the route.

Incorrect Answers:


A, B: The Split Horizon rule, which is enabled by default, prevents this.

C: Certkiller 1 will only load balance over equal cost routes by default. In this case, Certkiller 2 will always be used to reach network A since the route via Certkiller 3 will be seen with a higher metric.

Note: We could configure Certkiller 1 to load balance in this situation, using the "variance" router process configuration command.

QUESTION 219

Five Certkiller routers are part of the following EIGRP network:



Based on the diagram shown above, what is the feasible distance (FD) on router Certkiller 1 to network 172.16.9.0/24 through router Certkiller 3?

- A. 10,240
- B. 5120
- C. 6400
- D. 17,920
- E. 8960
- F. 11,520
- G. None of the above


Answer: F

Explanation:

The Feasible Distance to a network is simply the lowest calculated metric to each destination. In this case, the lowest cost metric to the given network through Certkiller 3 is $5120 + 3840 + 2560 = 11520$. It is important to note that even though the best path to the 172.16.9.0/24 network is through Certkiller 2, not Certkiller 3 but this question asked for the FD via Certkiller 3.

QUESTION 220

Routers Certkiller 1, Certkiller 2, and Certkiller 3 are all configured for EIGRP as shown below:



Certkiller 1 has the following configuration:

```
Router eigrp 1
Network 192.168.10.0
Redistribute connected
```

Which routes would show up in the routing table of Certkiller 3 as EIGRP routes? (Choose all that apply)

- A. 10.1.0.0/16
- B. 10.0.0.0/24
- C. 10.0.0.0/8
- D. 10.1.1.0/24
- E. 192.168.10.0/24

Answer: C, E

Explanation:

EIGRP will perform auto-summarization of External Routes. Since the 10.1.1.0 network was redistributed into EIGRP via a connected network, this will automatically make this route external to EIGRP. The 192.168.10.0 network will also show up in the routing table as an EIGRP route through the normal EIGRP process.

Additional Info:

Auto-Summarization

EIGRP performs an auto-summarization

each time it crosses a border between two different major networks.

For example, in Figure 13, Router Two advertises only the 10.0.0.0/8 network to Router One, because the interface Router Two uses to reach Router One is in a different major network.

On Router One, this looks like the following:




Figure 13

```
one#show ip eigrp topology 10.0.0.0
```

IP-EIGRP topology entry for 10.0.0.0/8

State is Passive, Query origin flag is 1, 1 Successor(s), FD is 11023872

Routing Descriptor Blocks:

172.16.1.1 (Serial0), from 172.16.1.2, Send flag is 0x0

Composite metric is (11023872/10511872), Route is Internal

Vector metric:

Minimum bandwidth is 256 Kbit

Total delay is 40000 microseconds

Reliability is 255/255

Load is 1/255

Minimum MTU is 1500

Hop count is 1

This route is not marked as a summary route in any way; it looks like an internal route. The metric is the best metric from among the summarized routes. Note that the minimum bandwidth on this route is 256k; although there are links in the 10.0.0.0 network that have a bandwidth of 56k.

On the router doing the summarization, a route is built to null0 for the summarized address:

```
two#show ip route 10.0.0.0
```

Routing entry for 10.0.0.0/8, 4 known subnets

Attached (2 connections)

Variably subnetted with 2 masks

Redistributing via eigrp 2000

C 10.1.3.0/24 is directly connected, Serial2

D 10.1.2.0/24 [90/10537472] via 10.1.1.2, 00:23:24, Serial1

D 10.0.0.0/8 is a summary, 00:23:20, Null0

C 10.1.1.0/24 is directly connected, Serial1

The route to 10.0.0.0/8 is marked as a summary through Null0. The topology table entry for this summary route looks like the following:

```
two#show ip eigrp topology 10.0.0.0
```

IP-EIGRP topology entry for 10.0.0.0/8

State is Passive, Query origin flag is 1, 1 Successor(s), FD is 10511872

Routing Descriptor Blocks:

0.0.0.0 (Null0), from 0.0.0.0, Send flag is 0x0

(note: the 0.0.0.0 here means this route is originated by this router)

Composite metric is (10511872/0), Route is Internal

Vector metric:

Minimum bandwidth is 256 Kbit

Total delay is 20000 microseconds

Reliability is 255/255

Load is 1/255

Minimum MTU is 1500

Hop count is 0

Incorrect Answers:

A, B, D. Any more specific routes in the 10.0.0.0 network will be summarized into one 10.0.0.0/8 network. Again, since the 10.0.0.0 network was learned by EIGRP only via redistribution, it is external as far as EIGRP is concerned.

QUESTION 221

You need to configure a Certkiller router to policy route specific traffic across an interface. What are two types of information that can be used to direct traffic along this route when using policy based routing?

- A. The packet time to live (TTL) and the source IP Address
- B. The Source IP address and the layer 2 source address
- C. The source IP Address and the protocol (such as FTP, HTTP and others)
- D. Type of service header and packet length
- E. None of the above

Answer: C

Explanation:

To enable PBR, you must create a route map that specifies the match criteria and the resulting action if all of the match clauses are met. Then, you must enable PBR for that route map on a particular interface. All packets arriving on the specified interface matching the match clauses will be subject to PBR.

To enable PBR on an interface, use the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	Router(config)# route-map map-tag [permit deny] [sequence-number]	Defines a route map to control where packets are output. This command puts the router into route-map configuration mode.
Step 2	Router(config-route-map)# match length min max Router(config-route-map)# match ip address {access-list-number name} [...access-list-number name]	Specifies the match criteria. Although there are many route-map matching options, here you can specify only length and/or ip address. <ul style="list-style-type: none"> ▪ length matches the Level 3 length of the packet. ▪ ip address matches the source or destination IP address that is permitted by one or more standard or extended access lists. If you do not specify a match command, the route map applies to <i>all</i> packets.

With PBR, you can use any standard or extended IP access list which allows you to base the match criteria on any source or destination IP address, or layer 4 port information

Reference:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fqos_c/fqcprt1/qcfpbr.htm

QUESTION 222

A router is being configured to override the normal routed behavior of certain traffic types. To do this, Policy Based Routing is used. Which of the following statements is FALSE with regards to the application of policy based routing (PBR)?

- A. PBR can not be used to set the IP precedence.
- B. PBR can not set the DSCP in one statement.
- C. PBR can be used to set the next hop IP address.
- D. PBR can be used to match on the length of a packet.
- E. All of the above are true

Answer: A

Explanation:

PBR gives you a flexible means of routing packets by allowing you to configure a defined policy for traffic flows, lessening reliance on routes derived from routing protocols. To this end, PBR gives you more control over routing by extending and complementing the existing mechanisms provided by routing protocols. PBR allows you to set the IP precedence. It also allows you to specify a path for certain traffic, such as priority traffic over a high-cost link.

You can set up PBR as a way to route packets based on configured policies. For example, you can implement routing policies to allow or deny paths based on the identity of a particular end system, an application protocol, or the size of packets.

PBR allows you to perform the following tasks:

1. Classify traffic based on extended access list criteria. Access lists, then, establish the match criteria.


2. Set IP Precedence bits, giving the network the ability to enable differentiated classes of service.

3. Route packets to specific traffic-engineered paths; you might need to route them to allow a specific QoS through the network.

Policies can be based on IP address, port numbers, protocols, or size of packets. For a simple policy, you can use any one of these descriptors; for a complicated policy, you can use all of them.

QUESTION 223

The Certkiller WAN and Internet connectivity is displayed below:



Router CK2 is configured as follows:

```
hostname CK2
```

```
!
```

```
interface Ethernet0/0
```

```
ip address 192.168.1.13 255.255.255.252
```

```
!
```

```
interface Ethernet0/1
```

```
ip address 192.168.1.9 255.255.255.252
```

```
!
```

```
interface Ethernet0/2
```

```
ip address 192.168.1.6 255.255.255.252
```

```
ip policy route-map net-200
```

```
!
```

```
router eigrp 1
```

```
network 192.168.1.0
```

```
!
```

```
access-list 101 permit ip 200.156.16.32 0.0.0.15 any
```

```
!  
route-map net-200 permit 10  
match ip address 101  
set interface Ethernet0/1  
!  
route-map net-10 permit 20  
!  
end
```

It is desired that all traffic from network 200.155.16.32/28 be sent to the internal through the firewall-enabled router CK4 . Router CK2 has been configured for policy-based routing as shown on the exhibit above. The policy-based configuration is not working. Debug and show commands indicate that Router CK2 has an "Incomplete" ARP entry for network 192.168.1.20. What is the best method to resolve this issue?

- A. Configure a static route to the 192.168.1.20 network in router CK2
- B. Configure ip proxy-arp on the router's Ethernet 0/1 and 0/2 interface
- C. Configure a static ARP entry for the 192.168.1.20 network on router CK2
- D. Reconfigure the "set interface" command to "set ip next-hop" with the IP address of the firewall
- E. Open the TCP ports on the firewall that are currently blocking ARP requests form router CK2

Answer: D

Explanation:

When configuring policy based routing on a multi-access network such as an Ethernet LAN, issues can arise when the interface is used as the next hop, rather than specifying the IP address. In this specific example, if we issue the "show arp" command we will see something similar to the following:

```
Cisco_Wan_Router# show arp  
Protocol Address Age (min) Hardware Addr Type Interface  
Internet 192.168.1.9 - 00b0.64cb.eab1 ARPA Ethernet0/1  
Internet 192.168.1.10 3 0010.7b81.0b19 ARPA Ethernet0/1  
Internet 192.168.1.20 0 Incomplete ARPA
```

Router CK2 attempts to do what it was instructed and tries to put the packets directly onto the Ethernet 0/1 interface. This requires that the router send an Address Resolution Protocol (ARP) request for the destination address of 192.1.1.1, which the router realizes is not on this interface, and hence the ARP entry for this address is "Incomplete," as seen by the show arp command. An encapsulation failure then occurs as the router is unable to put the packet on the wire with no ARP entry.

By specifying the IP address of the firewall as the next-hop, we can prevent this problem and make the route-map work as intended.

Configuration change should be:

```
!  
route-map net-200 permit 10
```

```
match ip address 101
set ip next-hop 192.168.1.10
!
```

Reference:

http://www.cisco.com/en/US/partner/tech/CK365/technologies_tech_note09186a008009481d.shtml#configforfire

QUESTION 224

Part of the configuration for router CK1 is displayed in the diagram below:

```
Hostname Certkiller1
!
interface Ethernet0/0
 ip address 192.168.1.1 255.255.255.0
 ip local policy route-map
!
ip local policy route-map reroute
ip classless
!
ip access-list extended reroute-acl
 permit tcp any 172.16.1.0 0.0.0.255 eq telnet
!
route-map mark-em-up permit 10
 match ip address reroute-acl
 set ip precedence flash
 set ip next-hop 192.168.1.20
!
route-map reroute permit 10
 match ip address reroute-acl
 set ip next-hop 192.168.1.25
```

Policy-Based routing has been configured on CK1 to sort traffic according to an administrative policy.

Which is the result from applying this configuration to Certkiller 1? (Select all that apply)

- A. All Telnet traffic destined to hosts on the 172.16.1.0/24 network will be forwarded to 192.168.1.20.
- B. All telnet traffic will be marked with IP Precedence Flash.
- C. Telnet traffic to destinations on the 172.16.1.0/24 network initiated from console connections on the router will be policy-routed to 192.168.1.25.
- D. Any telnet traffic transiting this router and exiting interface Ethernet 0/0 will be policy-routed to 192.168.1.20.
- E. If an administrator Telnets to Certkiller 1 and then subsequently telnets to 172.16.1.55, the session will be directed to 192.168.1.25

Answer: A, E

Explanation:

Choice A correctly describes the function of the normal policy based routing part of the configuration. In addition to this, a local policy route map has been configured. By default, packets that are originated from the router are not policy routed, unless a local

policy route map is configured as shown in this example. Because this has been applied to router Certkiller 1, telnet traffic originated from the router as described in choice E will be policy routed to the next hop IP address of 192.168.1.25.

Incorrect Answers:

B. Only telnet traffic destined to the 172.16.1.0/24 subnet will be marked with the flash IP precedence value.

C. Only packets that originate from the router are policy routed according to the local policy. This does not apply to connections that originate from the console interface.

D. Again, only telnet traffic that matches the reroute-acl access list will be policy routed, not all telnet traffic.

QUESTION 225

You need to configure a Certkiller router for Policy Based Routing (PBR). PBR allows network administrators to implement routing policies to allow or deny paths based on which of the following? (Choose four)

- A. Size of packets
- B. Identification of a particular end system
- C. Protocol
- D. Application
- E. Throughput

Answer: A, B, C, D

Explanation:

You can set up PBR as a way to route packets based on configured policies. For example, you can implement routing policies to allow or deny paths based on the identity of a particular end system, an application protocol, or the size of packets.

Reference:

http://www.cisco.com/en/US/docs/ios/12_0/qos/configuration/guide/qcclass.html

QUESTION 226

The Certkiller network administrator has enabled Weighted Random Early Discard on a router. Implementing WRED is most effective under what circumstance?

- A. There is an equal distribution of TCP and UDP traffic
- B. There is a mix of TCP, UDP and Non-IP traffic
- C. Most traffic is TCP based
- D. There are very high bandwidth interfaces such as Gigabit Ethernet
- E. None of the above

Answer: C

Explanation:

WRED avoids the globalization problems that occur when tail drop is used as the congestion avoidance mechanism. Global synchronization manifests when multiple TCP

hosts reduce their transmission rates in response to packet dropping, then increase their transmission rates once again when the congestion is reduced.

WRED is only useful when the bulk of the traffic is TCP/IP traffic. With TCP, dropped packets indicate congestion, so the packet source will reduce its transmission rate. With other protocols, packet sources may not respond or may resend dropped packets at the same rate. Thus, dropping packets does not decrease congestion.

Reference:

http://www.cisco.com/en/US/docs/ios/12_0/qos/configuration/guide/qcconavd.html

QUESTION 227

A Certkiller router was configured as shown below:

```
match ip precedence 7
class-map match-all Bearer
match ip precedence 5
!
policy-map ProviderOut
class Bearer
priority 48
class Signal
bandwidth 15
class System
bandwidth 15
class class-default
fair-queue
random-detect
shape average 512000
!
interface Ethernet0/1
description Provider Interface
ip address dhcp client-id Ethernet0/1
ip access-group 111 in
ip nat outside
full-duplex
no cdp enable
service-policy output ProviderOut
!
```

Based on the information shown above, what is the overall type of queuing that is being used on the outgoing data of the interface Ethernet0/1?

- A. Priority Queuing
- B. FIFO
- C. IP RTP Priority Queuing
- D. CBWFQ
- E. LLQ
- F. Weighted Fair Queuing
- G. None of the above

Answer: E

Explanation:

LLQ provides strict priority queueing for CBWFQ, reducing jitter in voice conversations. Configured by the priority command, LLQ enables use of a single, strict priority queue within CBWFQ at the class level, allowing you to direct traffic belonging to a class to the

CBWFQ strict priority queue. To enqueue class traffic to the strict priority queue, you specify the named class within a policy map and then configure the priority command for the class. (Classes to which the priority command is applied are considered priority classes.) Within a policy map, you can give one or more classes priority status. When multiple classes within a single policy map are configured as priority classes, all traffic from these classes is enqueued to the same, single, strict priority queue.

Example:

```
policy-map llqpolicy
class voice
priority 50
```

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800b

QUESTION 228

A new Certkiller router has been configured with a policy map as shown below:

```
class-map match-any tcp
match protocol http
match protocol ftp
class-map match-all acl100
match access-group 180

policy-map police
class tcp
police 150000 1000 conform-action transmit exceed-action drop
class acl100
police 150000 conform-action set-prec-transmit 2 exceed-action set-prec-transmit 1 violate-action set-prec-transmit 0

Interface Tunnel1:
ip address 20.1.1.2 255.255.255.252
service-policy input police
load-interval 30
tunnel source 4.4.4.2
tunnel destination 4.4.4.1
```

When applying this policy-map on the tunnel 1 interface, you see packet loss for class tcp starting at around 100000 bps instead of the configured 150000bps. What is the most likely cause for this?

- A. Policing on tunnel interfaces is not supported
- B. The load-interval on the tunnel interface prevents proper policing calculations
- C. The burst size is too low
- D. There should not be a violate-action configured
- E. The CIR keyword is missing in the policer
- F. None of the above

Answer: C

Explanation:

The policy map dictates what we want done to the traffic class previously defined. The police configuration command sets our rate limit in this example to 1.5Mbps with a burst size of 1000. The burst size is the trickiest part of this command. If the burst is set too low, your traffic will not be able to approach the maximum allowed throughput do to packet drops.

Because TCP window scaling halves the window size for each dropped packet, it's

important to set the burst size at a level that doesn't impact performance. The rule of thumb is that the burst size should be double the amount of traffic sent at the maximum rate at a given round-trip time.

QUESTION 229

A Certkiller router was configured as shown below:

```
class-map match-all af31
match dscp af31
class-map match-all af32
match dscp af32
class-map match-all af33
match dscp af33

policy-map Hairing
class af31
set dscp af31
class af32
set dscp af32
class af33
set dscp af33

policy-map limit
class af33
police cir 150000 bc 50000 pir 200000 be 50000
conform-action set-dscp-transmit af31
exceed-action set-dscp-transmit af32
violate-action set-dscp-transmit default
class class-default
bandwidth 300

interface Ethernet0/1
ip address 3.3.3.1 255.255.255.0
no ip proxy-arp
load-interval 30
half-duplex
no keepalive
no cdp enable
service-policy input marking

interface Serial0/0
ip address 4.4.4.1 255.255.255.252
ip load-sharing per-packet
encapsulation ppp
load-interval 30
no dce-terminal-timing-enable
service-policy output limit
```

As a network administrator at Certkiller , you have configured a dual-rate, dual bucket policer as per FRC 2698 on the serial interface of your router, connecting to you provider. The SLA with your provider states that you should send only AF31 (limited to 150kbps), AF32 (limited to 50kbps) and AF33 (best effort. Your service provider claims you are not conforming to the SL

- A. What is wrong? (Choose 2)
- A. This policer configuration is not implementing RFC 2698 dual bucket, dual-rate
 - B. The policer is configured in the wrong class
 - C. Configuring a service-policy on half-duplex Ethernet interfaces is not supported
 - D. The violate action is wrong
 - E. Class-default in policy-map limit should be marking traffic to dscp defaults

Answer: B, D

Explanation:

In this example we can see that there are two problems, one is that the policer is configured in the wrong class, and in order to only send AF31, AF32, and AF33 it's violate action should be set to drop. It is currently configured to set DSCP markings and transmit all other traffic, not drop it.

QUESTION 230

Weighted Random Early Discard (WRED) has been configured on a Certkiller WAN link. What are two benefits of applying WRED to this connection? (Choose 2)

- A. Completely eliminates congestion
- B. Avoids TCP synchronization
- C. Provides minimal bandwidth guarantees
- D. Provides bounded low latency
- E. A different drop profile can be manually enabled per IP Precedence or DSCP

Answer: B, E

Explanation:

WRED and distributed WRED (DWRED)-both of which are the Cisco implementations of RED-combine the capabilities of the RED algorithm with the IP Precedence feature.

Within the section on WRED, the following related features are discussed:

Flow-based WRED. Flow-based WRED extends WRED to provide greater fairness to all flows on an interface in regard to how packets are dropped.

DiffServ Compliant WRED. DiffServ Compliant WRED extends WRED to support Differentiated Services (DiffServ) and Assured Forwarding (AF) Per Hop Behavior (PHB). This feature enables customers to implement AF PHB by coloring packets according to differentiated services code point (DSCP) values and then assigning preferential drop probabilities to those packets.

WRED avoids the globalization problems that occur when tail drop is used as the congestion avoidance mechanism on the router. Global TCP synchronization occurs as waves of congestion crest only to be followed by troughs during which the transmission link is not fully utilized. Global synchronization of TCP hosts, for example, can occur because packets are dropped all at once. Global synchronization manifests when multiple TCP hosts reduce their transmission rates in response to packet dropping, then increase their transmission rates once again when the congestion is reduced.

Reference:

<http://www.google.com/search?hl=en&q=WRED%2C+A+different+drop+profile+can+be+manually+enabled+p>

QUESTION 231

QPPB is being utilized within the Certkiller network. QPPB allows which of the following marking behaviors?

- A. Provides no marking or classification
- B. Assigns only a specific BGP community based on the ingress packet DSCP marking
- C. Assigns a specific BGP attribute based on the IP precedence/DSCP of the inbound packet
- D. Uses NBAR to associate an IP Precedence to a packet
- E. None of the above

Answer: C

Explanation:

The Policy Propagation via BGP feature allows you to classify packets by IP Precedence based on BGP community lists, BGP autonomous system paths, and access lists. After a packet has been classified, you can use other quality of service features such as committed access rate (CAR) and Weighted Random Early Detection (WRED) to specify and enforce policies to fit your business model.

Note that this allows you to set up a policy at one BGP speaking router, and propagate that to other routers via BGP. Hence the name. This means that at the service provider router connecting to a site, a policy can be set up so that inbound traffic elsewhere is classified into the right class of service (IP Precedence bits). This can then interact with Tag Switching, or MPLS.

Reference:

http://www.cisco.com/en/US/docs/ios/12_0/qos/configuration/guide/qcbgprop.html

QUESTION 232

You are the administrator of the Certkiller network which has a main site and multiple remote sites. Your network carries both VOIP and data traffic and you need to implement QoS. You agree with your service provider to classify VOIP and data traffic according to the differentiated services RFCs. Which of the following are valid combinations for this?

- A. Data marked with DSCP AF21, VOIP marked as DSCP EF
- B. VOIP marked with DSCP AF31, data marked with DSCP EF
- C. Data marked as DSCP AF51 and VOIP marked DSCP EF
- D. Data marked with DE-bit, VOIP marked with CLP-bit
- E. Data marked with ip precedence 5, VOIP marked with DSCP EF
- F. None of the above

Answer: A

Explanation:

When addressing the QoS needs of Transactional Data and Interactive Data traffic, the following guidelines are recommended:

1. Transactional Data traffic should be marked to DSCP AF21; excess Transactional Data traffic can be marked down by a policer to AF22 or AF23.

RFC 2598 defines the Expedited Forwarding (EF) PHB: "The EF PHB can be used to build a low loss, low latency, low jitter, assured bandwidth, end-to-end service through DS (Diffserv) domains. Such a service appears to the endpoints like a point-to-point connection or a "virtual leased line." This service has also been described as Premium service." This is the recommended value for VOIP traffic.

QUESTION 233

Which of the following is FALSE regarding differences between Generic Traffic Shaping (GTS) and Frame Relay Traffic Shaping (FRTS)?

- A. GTS supports the traffic group command while FRTS does not.
- B. For GTS, the shaping queue is weighted fair queue (WFQ). FRTS does not support WFQ. With FRTS, the queue can be a CQ, PQ or FIFO.
- C. FRTS supports shaping on a per-DLCI basis, while GTS is configurable per interface or subinterface.
- D. GTS works with a variety of Layer 2 technologies, including Frame Relay, ATM, Switched Multimegabit Data Service, and Ethernet. FRTS is supported only on Frame Relay interfaces.

Answer: B

Explanation:

B. For FRTS, the queue can indeed be a weighted fair queue (configured by the frame-relay fair-queue command), a strict priority queue with WFQ (configured by the frame-relay ip rtp priority command in addition to the frame-relay fair-queue command), custom queuing (CQ), priority queuing (PQ), or first-in, first-out (FIFO).

Differences Between Traffic-Shaping Mechanisms

Generic traffic shaping (GTS), class-based shaping, distributed traffic shaping (DTS), and Frame Relay traffic shaping (FRTS) are similar in implementation, share the same code and data structures, but differ in regard to their CLIs and queue types used.

Following are some examples in which these mechanisms differ:

1. For GTS, the shaping queue is a weighted fair queue. For FRTS, the queue can be a weighted fair queue (configured by the frame-relay fair-queue command), a strict priority queue with WFQ (configured by the frame-relay ip rtp priority command in addition to the frame-relay fair-queue command), custom queuing (CQ), priority queuing (PQ), or first-in, first-out (FIFO).
2. For class-based shaping, GTS can be configured on a class, rather than only on an access control list (ACL). To do so, you must first define traffic classes based on match criteria including protocols, access control lists (ACLs), and input interfaces. Traffic shaping can be applied to each defined class.
3. FRTS supports shaping on a per-DLCI basis; GTS is configurable per interface or subinterface.

Incorrect Answers:

A, C, D. These statements are all true. For more on FRTS and GTS see the following URL (towards the bottom):

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800b

QUESTION 234

In the Certkiller Frame Relay network, Class Based Shaping is being used to increase network performance. Which of the following is a true statement regarding Class Based Shaping?

- A. CB shaping allows to rate-limit traffic in both incoming and outgoing directions.

- B. CB shaping provides a rate-limiting functionality with an associated amount of buffers, to store temporary out of profile traffic.
- C. CB shaping can only be configured in a child policy in a hierarchical policy map.
- D. CB shaping is a versatile feature which allows to both queue and remark traffic in input.
- E. None of the above
- F. All of the above

Answer: B

Explanation:

Traffic shaping allows you to control the traffic going out an interface in order to match its transmission to the speed of the remote, target interface and to ensure that the traffic conforms to policies contracted for it. Traffic adhering to a particular profile can be shaped to meet downstream requirements, thereby eliminating bottlenecks in topologies with data-rate mismatches. This is done with the use of buffers, which are used to temporarily store traffic that is queued. An optional Class Based Shaping command allows for the maximum number of buffers to be adjusted.

Incorrect Answers:

- A, D. Class Based Shaping is used for rate limiting outgoing traffic only. It does not provide for any mechanism to shape of mark incoming traffic.
- C. Class Based Shaping uses class-map statements. A set of hierarchical policy maps are not required for configuring CBS.

QUESTION 235

The Certkiller network is using FRTS to optimize the data flows within the network. In frame-relay traffic shaping (FRTS), what is the committed burst (Bc) parameter?

- A. The Bc is optional, and can be 0. It tells IOS how much extra bandwidth can be used on top of the CIR.
- B. The Bc is a parameter which needs to be negotiated with the provider of the frame-relay circuit. It defines the percentage of the frame-relay circuit IOS will use to send bursty traffic.
- C. Bc is a mandatory parameter when configuring FRTS. It defines a traffic rate up to which IOS will send traffic.
- D. Bc defines the amount of tokens added to the token bucket at each interval. The token bucket algorithm is used in FRTS. If not configured, it defaults to 56000 bits.
- E. Bc is total size of the token bucket. This includes the excess burst and conform burst.
- F. None of the above.

Answer: D

Explanation:

Bc (Committed Burst) is defined as the Maximum number of bits the frame relay network commits to transfer over a Committed Rate Measurement Interval (Tc). $Tc = Bc / CIR$. It is an optional parameter that defaults to 56000 bits, but it can indeed be set to 0, which

means that no traffic will be able to burst above the CIR.

To configure FRTS:

CK1 (config)# map-class frame-relay ?

WORD Static map class name

CK1 (config)# map-class frame-relay priority

CK1 (config-map-class)# frame-relay ?

adaptive-shaping Adaptive traffic rate adjustment, Default = none

bc Committed burst size (Bc), Default = 56000 bits

be Excess burst size (Be), Default = 0 bits

cir Committed Information Rate (CIR), Default = 56000 bps

CK1 (config-map-class)# frame-relay priority-group ?

<1-16> Priority group number

Incorrect Answers:

A. This would be true for Be, not Bc

B. The Bc is a value specified in bits per interval, not in a percentage.

C. Bc is optional, not mandatory. The only mandatory configuration guidelines for FRTS is to specify the interface with frame-relay encapsulation, and to enable FRTS with the "frame-relay traffic-shaping" interface command.

D. The default committed burst size is 7000 bits, when no value is specified. The default

E. Bc is the committed burst rate, not the total burst.

Reference:

http://www.cisco.com/en/US/tech/CK543/CK545/technologies_tech_note09186a0080104819.shtml

QUESTION 236

What is true about Class based Weighted Fair Queuing (CBWFQ)?

A. CBWFQ provides delay, jitter and bandwidth guarantees to traffic.

B. CBWFQ can be configured on any interface in either input or output.

C.

CBWFQ has to be configured with the Modular QoS CLI. The resulting service-policy has to be applied on output.

D. CBWFQ can only be configured in a hierarchical policy-map. The parent policy-map does policing and the child policy-map does CBWFQ.

E. All of the above

F. None of the above

Answer: C

Explanation:

To configure CBWFQ, there are 3 required steps: Defining class maps, configuring class policy in the policy map, and attaching the service policy and enabling CBWFQ. This is done using the new IOS syntax called Modular QoS. You must use the Modular QoS CLI to configure class based marking.

Incorrect Answers:

- A. There is no way to specify traffic guarantees for jitter and delay, as the underlying network that is used for transport will have the greatest impact on these values.
- B. The CBWFQ service policy can only be applied to outbound interfaces.
- D. Although CBWFQ is typically configured using this method, it not not required for all implementations.

Reference: Distributed QoS, Odom/Cavanaugh, Cisco Press, page 176.

QUESTION 237

CAR has been configured on router CK1 . What best defined Committed Access Rate (CAR)?

- A. CAR allows metering of traffic for traffic shaping.
- B. CAR is a feature that allows the rate limiting of traffic in either the incoming or outgoing direction.
- C. CAR is part of a set of features to be used in conjunction with queuing to form a hierarchical policy. CAR must always be applied in a parent policy-map, whereas CBWFQ should be applied in a child policy-map.
- D. CAR is a queuing feature.
- E. CAR matches only on UDP port range {16384 - 32767}.

Answer: B

Explanation:

The Committed Access Rate (CAR) and Distributed CAR (DCAR) services limit the input or output transmission rate on an interface or subinterface based on a flexible set of criteria.

The rate-limiting feature of CAR provides the network operator with the means to define Layer 3 aggregate or granular access, or egress bandwidth rate limits, and to specify traffic handling policies when the traffic either conforms to or exceeds the specified rate limits. Aggregate access or egress matches all packets on an interface or subinterface. Granular access or egress matches a particular type of traffic based on precedence. You can designate CAR rate-limiting policies based on physical port, packet classification, IP address, MAC address, application flow, and other criteria specifiable by access lists or extended access lists. CAR rate limits may be implemented either on input or output interfaces or subinterfaces including Frame Relay and ATM subinterfaces.

An example of use of CAR's rate-limiting capability is application-based rates limiting HTTP World Wide Web traffic to 50 percent of link bandwidth, which ensures capacity for non-Web traffic including mission-critical applications.

QUESTION 238

QoS mechanisms have been put in place within the Certkiller network using IP precedence. This IP precedence can be defined as:

- A. The ToS byte is the IP precedence
- B. The middle 4 bits of the ToS byte
- C. The 3 left most bits of the ToS byte

- D. The 3 right most bits of the ToS byte
- E. The right most bit of the ToS byte

Answer: C

Explanation:

Within the Type of Service (TOS) byte, the three most significant bits are the IP precedence bits. The TOS byte is displayed below:

P2	P1	P0	T2	T1	T0	CU1	CU0
----	----	----	----	----	----	-----	-----

1. IP precedence-three bits (P2 to P0)
2. Delay, Throughput and Reliability-three bits (T2 to T0)
3. CU (Currently Unused)-two bits(CU1-CU0)

Reference: <http://www.cisco.com/warp/public/105/dscpvalues.html>

QUESTION 239

Due to intermittent congestion issues on a link, Committed Access Rate (CAR) has been configured on an interface. During a period of congestion, a packet arrives that causes the compounded debt to be greater than the value set for the extended burst. Which of the following will occur due to this? (Choose all that apply).

- A. CAR's exceed action takes effect, dropping the packet.
- B. A token is removed from the bucket.
- C. The packet will be queued and eventually serviced.
- D. The compounded debt value is effectively set to zero (0).
- E. The packet is buffered by the CAR process.

Answer: A, D

Explanation:

Here is how the extended burst capability works. If a packet arrives and needs to borrow n number of tokens because the token bucket contains fewer tokens than its packet size requires, then CAR compares the following two values:

1. Extended burst parameter value
2. Compounded debt. Compounded debt is computed as the sum over all ai.
3. 1. i indicates the ith packet that attempts to borrow tokens since the last time a packet was dropped.
2. a indicates the actual debt value of the flow after packet i is sent. Actual debt is simply a count of how many tokens the flow has currently borrowed.

If the compounded debt is greater than the extended burst value, CAR's exceed action takes effect. After a packet is dropped, the compounded debt is effectively set to 0. CAR will compute a new compounded debt value equal to the actual debt for the next packet that needs to borrow tokens.

If the actual debt is greater than the extended limit, all packets will be dropped until the actual debt is reduced through accumulation of tokens in the token bucket.

Incorrect Answers:

- B. Dropped packets do not count against any rate or burst limit. That is, when a packet is

dropped, no tokens are removed from the token bucket.

C, E. After the exceed action takes place, the packet is dropped immediately and is not buffered.

Reference:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/qos_c/qcpart4/qcpolts.htm

QUESTION 240

In an effort to minimize the risks associated from DOS and ICMP flooding attacks, the following is configured on the serial interface of a router:

```
interface serial 0
```

```
rate-limit input access-group 199 128000 4000 4000 conform-action
```

```
transmit exceed-action drop
```

```
access-list 199 permit icmp any any
```

What QoS feature is this an example of?

- A. CBWFQ
- B. LLQ
- C. RSVP
- D. CAR
- E. WFQ
- F. FRTS

Answer: D

Explanation:

Committed Access Rate (CAR) is used to rate limit traffic. In this example, all ICMP traffic that exceeds the defined level will be dropped. This will prevent an ICMP flood attack from saturating the link.

CAR definition:

Rate limiting is one mechanism to use to allow a network to run in a degraded manner, but remain up when it is receiving a stream of Denial of Service (DoS) attack packets as well actual network traffic. Rate limiting can be achieved in a number of methods using Cisco IOS(r) software. Namely, through Committed Access Rate (CAR), Traffic Shaping, and both Shaping and Policing through Modular Quality of Service Command Line Interface (QoS CLI).

Incorrect Answers:

A. Class-based weighted fair queuing (CBWFQ) extends the standard WFQ functionality to provide support for user-defined traffic classes. For CBWFQ, you define traffic classes based on match criteria including protocols, access control lists (ACLs), and input interfaces. Packets satisfying the match criteria for a class constitute the traffic for that class. A queue is reserved for each class, and traffic belonging to a class is directed to the queue for that class.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1830/products_feature_guide09186a0080087a84.html

B, C. RSVP and LLQ (low latency queuing) are often implemented in voice and video

data networks, but are not typically used for preventing DOS attacks.

F. FRTS is frame relay traffic shaping. It is not clear from this example that the link is even using frame relay as the transport link.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_tech_note09186a00800fb50a.shtml

QUESTION 241

Which of the following are functions of Random Early Discard (RED)? (Choose all that apply)

- A. To avoid global synchronization for TCP traffic.
- B. To provide unbiased support for bursty traffic.
- C. To minimize packet delay jitter.
- D. To ensure that high priority traffic gets sent first.
- E. To prevent the starvation of the lower priority queues.

Answer: A, B, C

Explanation:

When it comes to Quality of Service, there are 2 separate approaches. The first is congestion management, which is setting up queues to ensure that the higher priority traffic gets serviced in times of congestion. The other is congestion avoidance, which works by dropping packets before congestion on the link occurs. Random Early Detection (RED) is a congestion avoidance mechanism that takes advantage of TCP's congestion control mechanism.

RED takes a proactive approach to congestion. Instead of waiting until the queue is completely filled up, RED starts dropping packets with a non-zero drop probability after the average queue size exceeds a certain minimum threshold. A drop probability ensures that RED randomly drops packets from only a few flows, avoiding global synchronization. A packet drop is meant to signal the TCP source to slow down. Responsive TCP flows slow down after packet loss by going into slow start mode.

Incorrect Answers:

- D. This would be a function of priority queuing, not RED. Weighted RED (WRED) is used to assign priorities to traffic and works to not drop the higher priority traffic types, but RED does not.
- E. This is a function of custom queuing, which is a congestion management mechanism, not a congestion avoidance mechanism such as RED.

Reference:

'IP Quality of Service' page 130, Cisco Press.

QUESTION 242

Rate Limiting is configured on the Ethernet interface of a router as follows:

```
interface Ethernet 0
```

```
rate-limit input access-group rate limit 1 1000000 10000 10000
```

```
conform-action
```

```
access-list rate-limit 1 mask 07
```

What effect will this configuration have?

- A. The command access rate policing limits all TCP traffic to 10Mbps.
- B. Traffic matching access-list 7 is rate limited.
- C. Voice traffic with DiffServ code point 43 is guaranteed.
- D. Traffic with IP Precedence values of 0, 1, and 2 will be policed.

Answer: D

Explanation:

Use the mask keyword to assign multiple IP precedence's to the same rate-limit list. To determine the mask value, perform the following steps:

Step 1 Decide which precedence's you want to assign to this rate-limit access list.

Step 2 Convert the precedence's into an 8-bit numbers with each bit corresponding to one precedence. For example, an IP precedence of 0 corresponds to 00000001, 1 corresponds to 00000010, 6 corresponds to 01000000, and 7 corresponds to 10000000.

Step 3 Add the 8-bit numbers for the selected precedence's together. For example, the mask for precedence's 1 and 6 is 01000010.

Step 4 Convert the binary mark into the corresponding hexadecimal number. For example, 01000010 becomes 0x42. This value is used in the access-list rate-limit command. Any packets that have an IP precedence of 1 or 6 will match this access list. A mask of FF matches any precedence, and 00 does not match any precedence.

In this example, a mask of 07 translates to 00000111, so IP precedence 0, 1, and 2 will be policed.

QUESTION 243

When configuring Low Latency Queuing (LLQ), a bandwidth parameter is needed. What does this parameter specify?

- A. It provides a built in policer to limit the priority traffic in the LLQ during congestion.
- B. This parameter is optional, since the LLQ will always have precedence over other queues.
- C. This parameter should be as low as possible. It represents bandwidth which will always be reserved. It reduces the amount of bandwidth on the interface, even if it is not used by any LLQ traffic.
- D. It represents the reference CIR to calculate the burst size of the token bucket of the built-in policer.
- E. None of the above.

Answer: A

Explanation:

The bandwidth argument is used to specify the maximum amount of bandwidth allocated for packets belonging to a class configured with the priority command. The bandwidth parameter both guarantees bandwidth to the priority class and restrains the flow of packets from the priority class.

When the device is not congested, the priority class traffic is allowed to exceed its allocated bandwidth. When the device is congested, the priority class traffic above the allocated bandwidth is discarded.

Reference

:http://www.cisco.com/en/US/products/sw/iosswrel/ps1834/products_feature_guide09186a0080080232.html#47

QUESTION 244

What statement is FALSE with regards to Weighted RED (WRED)?

- A. WRED is a congestion avoidance mechanism, based on the adaptive nature of TCP traffic for congestion.
- B. WRED is a queuing feature.
- C. WRED allows for differentiated dropping behavior based on either IP precedence or DSCP.
- D. WRED is configurable in a CBWFQ policy-map.
- E. All of the above are false statements.

Answer: B

Explanation

The WRED algorithm provides congestion avoidance on network interfaces by providing buffer management, and by allowing Transmission Control Protocol (TCP) traffic to throttle back before buffers are exhausted. This helps avoid tail drops and global synchronization issues, maximizing network usage and TCP-based application performance. WRED works by selectively dropping packets before congestion occurs, so it is considered to be a congestion avoidance feature, not a queuing feature.

Incorrect Answers:

- A. WRED is only useful when the bulk of the traffic is TCP/IP traffic. With TCP, dropped packets indicate congestion, so the packet source will reduce its transmission rate. With other protocols, packet sources may not respond or may resend dropped packets at the same rate. Thus, dropping packets does not decrease congestion.
- C. WRED works with the IP precedence or DSCP values to determine which packets get dropped first. You can configure WRED to ignore IP Precedence when making drop decisions so that nonweighted RED behavior is achieved.
- D. WRED can indeed be configured in a policy map that is applied to class based weighted fair queuing as specified in the following:
http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a00801b2406.html

QUESTION 245

A newly deployed Cisco router is configured for Weighted Random Early Discard (WRED) on the Certkiller network. WRED has which of the following two characteristics? (Choose two.)

- A. Non-IP traffic is given the lowest priority and is more likely to be dropped
- B. When the minimum threshold is crossed, WRED begins dropping all incoming packets

(tail-drop)

C. Global Synchronization is avoided by selectively dropping packets from multiple TCP flows

D. Low bandwidth flows are experiencing packet drop at a higher rate than higher bandwidth flows

Answer: A, C

Explanation:

By randomly dropping packets prior to periods of high congestion, WRED tells the packet source to decrease its transmission rate. If the packet source is using TCP, it will decrease its transmission rate until all the packets reach their destination, which indicates that the congestion is cleared.

WRED generally drops packets selectively based on IP precedence. Packets with a higher IP precedence are less likely to be dropped than packets with a lower precedence. Thus, the higher the priority of a packet, the higher the probability that the packet will be delivered.

WRED reduces the chances of tail drop by selectively dropping packets when the output interface begins to show signs of congestion. By dropping some packets early rather than waiting until the queue is full, WRED avoids dropping large numbers of packets at once and minimizes the chances of global synchronization. Thus, WRED allows the transmission line to be used fully at all times.

In addition, WRED statistically drops more packets from large users than small.

Therefore, traffic sources that generate the most traffic are more likely to be slowed down than traffic sources that generate little traffic.

WRED avoids the globalization problems that occur when tail drop is used as the congestion avoidance mechanism. Global synchronization manifests when multiple TCP hosts reduce their transmission rates in response to packet dropping, then increase their transmission rates once again when the congestion is reduced.

WRED is only useful when the bulk of the traffic is TCP/IP traffic. With TCP, dropped packets indicate congestion, so the packet source will reduce its transmission rate. With other protocols, packet sources may not respond or may resend dropped packets at the same rate. Thus, dropping packets does not decrease congestion.

WRED treats non-IP traffic as precedence 0, the lowest precedence. Therefore, non-IP traffic, in general, is more likely to be dropped than IP traffic.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800c

QUESTION 246

Traffic Classification using NBAR is configured using which IOS command?

A. Router CK1 (Config-if)#ip nbar protocol-discovery

B. Router CK1 (Config)#ip nbar port-map {protocol} [tcp|udp] {port-number} {port-number}

C. Router CK1 (Config-cmap)#match protocol {protocol-name}

- D. Router CK1 (Config)#ip nbar {pdlm-file}
- E. Router CK1 (Config-cmap)#match access-group {number}
- F. Router CK1 (Config-pmap)#class bar

Answer: C

Explanation:

To configure a traffic class and the match criteria that will be used to identify traffic as belonging to that class, use the class-map global configuration command.

	Command	Purpose
Step 1	Router(config)# class-map [match-all match-any] <i>class-name</i>	Specifies the user-defined name of the class map. The match-all option specifies that all match criteria in the class map must be matched. The match-any option specifies that one or more match criteria must match.
Step 2	Router(config-cmap)# match protocol <i>protocol-name</i>	Specifies a protocol supported by NBAR as a matching criterion.

Example:

In the following example, the class-map class1 command uses the NBAR classification of SQL*Net as its matching criterion:

```
Router(config)# class-map class1
Router(config-cmap)# match protocol sqlnet
```

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800c

QUESTION 247

A Certkiller router has been configured with traffic policing. What functionality does this traffic policing provide?

- A. Traffic Policing provides a means of providing low latency on a congested interface
- B. Traffic Policing is mandatory when enabling NBAR and CBWQ
- C. Traffic Policing allows you to meter and limit bandwidth utilization
- D. Traffic Policing is also referred to as 'buffer tuning'. It is an optimal way to manage the routers I/O memory
- E. Traffic policing should be configured on every router to avoid memory corruptions
- F. None of the above

Answer: C

Explanation:

Traffic policing allows you to control the maximum rate of traffic sent or received on an interface, and to partition a network into multiple priority levels or class of service (CoS). The Traffic Policing feature manages the maximum rate of traffic through a token bucket algorithm. The token bucket algorithm can use the user-configured values to determine

the maximum rate of traffic allowed on an interface at a given moment in time. The token bucket algorithm is affected by all traffic entering or leaving (depending on where the traffic policy with Traffic Policing configured) and is useful in managing network bandwidth in cases where several large packets are sent in the same traffic stream.

The token bucket algorithm provides users with three actions for each packet: a conform action, an exceed action, and an optional violate action. Traffic entering the interface with Traffic Policing configured is placed in to one of these categories. Within these three categories, users can decide packet treatments. For instance, packets that conform can be configured to be transmitted, packets that exceed can be configured to be sent with a decreased priority, and packets that violate can be configured to be dropped.

Traffic Policing is often configured on interfaces at the edge of a network to limit the rate of traffic entering or leaving the network. In the most common Traffic Policing configurations, traffic that conforms is transmitted and traffic that exceeds is sent with a decreased priority or is dropped. Users can change these configuration options to suit their network needs.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800b

QUESTION 248

Network-Based Application Recognition (NBAR) has been configured on numerous devices within the Certkiller network. Network-Based Application Recognition is used to provide which Quality of Service function?

- A. Classification
- B. Policing
- C. CBWFQ Bandwidth guarantees
- D. Shaping
- E. None of the above

Answer: A

Explanation:

The Network-Based Application Recognition (NBAR) feature adds intelligent network classification to network infrastructures. NBAR is a classification engine that recognizes a wide variety of applications, including web-based and other difficult-to-classify protocols that utilize dynamic TCP/User Datagram Ports (UDP) port assignments. When an application is recognized and classified by NBAR, a network can invoke services for that specific application.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800c

QUESTION 249

You want to mark all VOIP packets in the Certkiller network with the DSCP value of EF. According to the differentiated services RFCs, traffic with Expedited

Forwarding Per Hop Behavior should be marked as:

- A. DSCP decimal 36
- B. IP ToS of 0xEF
- C. IP Experimental ECN
- D. DSCP decimal 5
- E. Binary Value of 101110
- F. None of the above

Answer: E

Explanation:

The table below lists the IP precedence and DSCP values, and their names, for review.

Note that not all DSCP values are listed; only the DSCP values suggested by the DiffServ RFCs are listed in the table. QoS tools that are capable of setting DSCP can set any of the actual 64 values.

Field and Value (Decimal)	Binary Value	Name	Defined by This RFC
Precedence 0	000	routine	791
Precedence 1	001	priority	791
Precedence 2	010	immediate	791
Precedence 3	011	flash	791
Precedence 4	100	flash override	791
Precedence 5	101	critic	791
Precedence 6	110	internetwork control	791
Precedence 7	111	network control	791
DSCP 0	000000	best effort	2475
DSCP 8	001000	CS1	2475
DSCP 16	010000	CS2	2475
DSCP 24	011000	CS3	2475
DSCP 32	100000	CS4	2475
DSCP 40	101000	CS5	2475
DSCP 48	110000	CS6	2475
DSCP 56	111000	CS7	2475
DSCP 10	001010	AF11	2597
DSCP 12	001100	AF12	2597
DSCP 14	001110	AF13	2597
DSCP 18	010010	AF21	2597
DSCP 20	010100	AF22	2597
DSCP 22	010110	AF23	2597
DSCP 26	011010	AF31	2597
DSCP 28	011100	AF32	2597
DSCP 30	011110	AF33	2597
DSCP 34	100010	AF41	2597
DSCP 36	100100	AF42	2597
DSCP 38	100110	AF43	2597
DSCP 46	101110	EF	2598

Reference: <http://www.ciscopress.com/articles/article.asp?p=101170&seqNum=2&rl=1>

QUESTION 250

You need to configure NBAR on a Certkiller router for QoS purposes. What Quality of Service functionality does NBAR (Network Based Application Recognition) provide? (Choose two)

- A. NBAR is an advanced algorithm acting as scheduler in a MQC policy-map
- B. NBAR provides deep packet inspection and is used for advanced packet classification
- C. NBAR provides a per-protocol packet and bytes accounting functionality. This is used to track bandwidth utilization for all protocols described in the loaded PDLs
- D. NBAR is a mandatory component of MQC (Modular QoS CLI). Without NBAR it is impossible to do any MQC configuration
- E. The use of NBAR is configured with an application-policy

Answer: B, C

Explanation:

NBAR addresses IP QoS classification requirements by classifying application-level protocols so that QoS policies can be applied to the classified traffic. NBAR addresses the ongoing need to extend the classification engine for the many existing and emerging application protocols by providing an extensible Packet Description Language (PDL). NBAR can determine which protocols and applications are currently running on a network so that an appropriate QoS policy can be created based upon the current traffic mix and application requirements.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800c

QUESTION 251

A Certkiller router is configured using MQC as shown below:

```
!
class-map match-all Signal
  match ip precedence 3
class-map match-any System
  match access-group name Security
  match ip precedence 6
  match ip precedence 7
class-map match-all Bearer
  match ip precedence 5
!
!
policy-map ProviderOut
  class Bearer
    priority 48
  class Signal
    bandwidth 15
  class System
    bandwidth 15
  class class-default
    fair-queue
    random-detect
    shape average 512000
!
interface Ethernet0/1
  description Provider Interface
  ip address dhcp client-id Ethernet0/1
  ip access-group 111 in
  ip nat outside
  full-duplex
  no cdp enable
  service-policy output ProviderOut
!
```

Based on the information shown above, which of the following is applied to class Bearer?

- A. WRED
- B. Traffic Shaping
- C. Packet Marking
- D. Packet Classification
- E. FIFO queuing within the class

F. None of the above

Answer: E

Explanation:

The "priority" command is used to configure Low Latency queuing (LLQ) where the traffic assigned to the LLQ is given strict priority with FIFO queuing.

The priority command can be configured in multiple classes, but it should only be used for voice-like, constant bit rate (CBR) traffic. If the traffic is not CBR, you must configure a large enough bandwidth parameter to absorb the data bursts.

Configuring the priority command in multiple classes provides the ability to police the priority classes individually. For an example, refer to the following configuration:

```
policy-map policy1
```

```
class voice1
```

```
priority 24
```

```
class voice2
```

```
priority 48
```

```
class data
```

```
bandwidth 20
```

In this example, voice1 and voice2 classes of traffic go into the high priority queue and get strict priority queueing over data traffic. However, voice1 traffic will be rate-limited to 24 kbps and voice2 traffic will be rate-limited to 48 kbps. The classes will be individually rate-limited (and given first-in first-out [FIFO] treatment) even if they go into the same queue.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1830/products_feature_guide09186a0080087b13.html

QUESTION 252

Switch Certkiller 2 was configured as shown below:

```
Certkiller2(config)# wrr-queue bandwidth 10 20 70 1
```

```
Certkiller2(config)# no wrr-queue cos-map
```

```
Certkiller2 (config)# wrr -queue cos-map 1 0 1
```

```
Certkiller2 (config)# wrr-queue cos-map 2 2 4
```

```
Certkiller2 (config)# wrr -queue cos-map 3 3 6 7
```

```
Certkiller2 (config)# wrr -queue cos-map 4 5
```

Based on the information shown above, what does the IOS configuration on the Certkiller 2 Catalyst 2950 accomplish?

- A. It sets up the WRR queuing where frames with a CoS of 3 or 6 or 7 will have the highest priority.
- B. It sets up the CoS-to-DSCP mappings and DSCP-to-CoS mappings.
- C. It enables frames with a CoS 5 marking to be serviced by the expedite queue.
- D. 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.
- E. 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.
F. None of the above

Answer: A

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

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

Defaults

The defaults are as follows:

Receive queue1/drop threshold 1 and transmit queue1/drop threshold 1:CoS 0 and 1.

Receive queue1/drop threshold 2 and transmit queue1/drop threshold 2:CoS 2 and 3.

Receive queue2/drop threshold 3 and transmit queue2/drop threshold 1:CoS 4 and 6.

Receive queue2/drop threshold 4 and transmit queue2/drop threshold 2:CoS 7.

QUESTION 253

You need to ensure that mission critical traffic gets prioritized in the Certkiller network. The IP precedence of a packet can be determined from which of the

following?

- A. All 8 bits of the ToS byte
- B. Bits 4, 5, and 6 of the ToS byte
- C. The three least significant bits of the ToS byte
- D. The three most significant bits of the ToS byte
- E. None of the above

Answer: D

Explanation:

This DSCP field definition allows for up to 64 distinct values (levels of service), 0 through 63, of classification on IP frames. The last two bits represent the Early Congestion Notification (ECN) bits. IP Precedence is only the 3 most significant bits of the ToS field. As a result, IP Precedence maps to DSCP by using IP Precedence as the 3 high-order bits and padding the lower-order bits with 0.

QUESTION 254

You need to configure priority queuing on router CK1 . What statement is true with regards to priority queuing?

- A. The high and medium queues have precedence over the default queue.
- B. There are 4 priority queues: high, medium, normal, low.
- C. The classification is configurable via the "priority-list" command.
- D. The default queue is the normal queue.
- E. All of the above
- F. None of the above

Answer: E

Explanation:

There are four priority queues: high, medium, normal, and low- listed in order from highest to lowest priority.

The default queue is the normal queue. Traffic that is not explicitly defined in the priority list will be assigned this priority by default.

Priority queuing is configured using the "priority-list" command.

Example:

In the following example, queuing priority for all telnet and SMTP traffic is assigned the high priority.

```
priority-list 1 protocol ip high tcp 23
```

```
priority-list 1 protocol ip high tcp 25
```

QUESTION 255

The relevant part of a Certkiller router's configuration is displayed below:

```
ip cef
|
class-map match-all VoIP-Remark
  match ip dscp ef
  match ip dscp cs3
  match ip dscp af31
|
class-map match-any VoIP-RTP
  match protocol rtp audio
  match access-group name VoIP-RTP
|
policy-map Policy-NoTrust
  class VoIP-RTP
    set ip dscp ef
  class VoIP-Remark
    set ip dscp default
|
interface Ethernet0
  description Outside Interface
  ip address 192.168.1.1 255.255.255.0
  service-policy input Policy-NoTrust
|
interface FastEthernet0
  description Inside Interface
  ip address 192.168.2.1 255.255.255.0
|
|
ip access-list extended VoIP-RTP
  permit udp any any range 16384 32767
```

Classed-Based marking has been configured on this Certkiller router as shown above to sort traffic into classes for appropriate treatment by upstream routers. Unfortunately, traffic received by an upstream router on the 192.168.1.0/24 network is not appropriately marked; VoIP packets are not marked *EF* and packets previously set by end-users as CS3 and EF are not remarked to DSCP 00000. Which of the following issues could be the cause of the problem? (Select all that apply).

- A. Using NBAR to classify RTP traffic requires that IP CEF be disabled.
- B. The VoIP-RTP access list includes both even and odd-numbered ports starting from 16384; while RTP only uses even-numbered ports.
- C. The router has been improperly configured to only mark traffic flowing in the wrong direction.
- D. The set ip dscp default command does not mark the packet with DSCP 0000, but instead resets the command use back to Cisco IOS default settings.
- E. The Class-map VoIP-Remark has been improperly configured to simultaneously match more than one traffic type.
- F. None of the above

Answer: C, E

Explanation:

The first problem is that the service policy named Policy-NoTrust is being applied to the input direction of the ethernet interface, when it should be applied in the outbound direction for the upstream routers to see the correct DSCP markings of the packets. The second problem is the fact that the keyword "match-all" is being applied to the

VOIP-Remark class map. This keyword instructs the IOS that all of the criteria in the entire map must match in order to be applied. In the configuration above, only a packet that matches all three of the criteria (DSCP EF, DSCP CS3, and DSCP AF21) will be marked, instead of packets that match any one of those. In this example, the correct syntax should have been "class-map match-any VoIP-Remark"

QUESTION 256

The Certkiller network plans to implement some method of quality of service using DSCP information. In comparing the different options which of the following statements is correct?

- A. The IP precedence and DSCP have no overlapping fields.
- B. DSCP is only for TCP; IP precedence is for UDP
- C. The DSCP is exactly the same as IP precedence; the name change is merely a marketing naming convention.
- D. The last 2 bits of the DSCP overlap with the IP precedence
- E. The DSCP contains class selectors for backward compatibility with the IP precedence.
- F. None of the above

Answer: E

Explanation:

The Per-Hop Behavior is indicated by encoding a 6-bit value-called the Differentiated Services Code Point (DSCP)-into the 8-bit Differentiated Services (DS) field of the IP packet header.

In theory, a network could have up to 64 (26) different traffic classes using different markings in the DSCP. The DiffServ RFCs recommend, but do not require, certain encodings, which gives a network operator great flexibility in defining traffic classes. In practice, however, most networks use the following commonly-defined Per-Hop Behaviors:

Default PHB-which is typically best-effort traffic

Expedited Forwarding (EF) PHB-for low-loss, low-latency traffic

Assured Forwarding (AF)-behavior group

Class Selector PHBs-which are defined to maintain backward compatibility with the IP Precedence field.

Prior to DiffServ, IP networks could use the Precedence field in the Type of Service (TOS) byte of the IP header to mark priority traffic. The TOS byte and IP precedence was not widely used. The IETF agreed to reuse the TOS byte as the DS field for DiffServ networks. In order to maintain backward compatibility with network devices that still use the Precedence field, DiffServ defines the Class Selector PHB.

The Class Selector codepoints are of the form 'xxx000'. The first three bits are the IP precedence bits. Each IP precedence value can be mapped into a DiffServ class. If a packet is received from a non-DiffServ aware router that used IP precedence markings, the DiffServ router can still understand the encoding as a Class Selector codepoint.

Reference: http://en.wikipedia.org/wiki/Differentiated_services

QUESTION 257

Router CK1 is configured for QoS as shown below:

```
|
ip cef
|
class-map match-all cos3and4
match cos 3 4
|
class-map match-all trans
match protocol http
match protocol telnet
|
class-map match-all scavenger
match protocol napster
match-cos 1

policy-map ccietest
class cos3and4
set dscp af33
class trans
set dscp af21
class scavenger
set dscp cs1
|
interface fastethernet0/0
ip address 10.1.1.1 255.255.255.0
service-policy input ccietest
```

Based on the configuration displayed in the exhibit, what statement is correct about ingress traffic to the fa0/0 interface on CK1 ?

- A. All ingress frames marked as COS 0 will be marked as DSCP 0.
- B. All ingress frames marked as COS 1 will be marked as DSCP cs1.
- C. All ingress HTTP traffic will be marked as DSCP af21.
- D. All ingress Napster traffic will be marked as DSCP cs1.
- E. All ingress frames marked as COS 3 or COS 4 will be marked as DSCP af33.
- F. None of the above.

Answer: E

Explanation:

Since the "match cos 3 4" statement lies within a single configuration line, only one or the other need to match. When the keyword "match-all" is used, all distinct lines must match for the rule to take effect. Since the values shown in choice E are displayed in a single line, all traffic with COS values of 3 or 4 will match, and will subsequently be forwarded after being marked as AF33.

Additional info:

To access the QoS class map configuration mode to configure QoS class maps, use the class-map command. Use the no form of this command to delete a class map.

class-map name [match-all | match-any]

no class-map name [match-all | match-any]

Syntax Description

name	Class map name.
match-all	(Optional) Matches all match criteria in the class map.
match-any	(Optional) Matches one or more match criteria.

Defaults:

When you do not specify the match-all or match-any keyword, the default is match-all.

Incorrect Answers:

A. CoS values of 0 are not automatically marked with a DSCP value of 0

B, D. Here, only frames marked as COS and using the Napster protocol will be marked with a DSCP value of 1.

C. Only traffic that is both HTTP and Telnet will be marked as such. This is obviously not possible since HTTP uses port 80 while telnet uses port 23.

QUESTION 258

Part of the configuration file of router CK1 is shown in the diagram below:

```
!
ip cef
!
class-map match-all bulk
match protocol ftp
match protocol tftp
!
policy-map mark
class bulk
set dscp all 1
!
int fastethernet0/0
ip address 10.1.1.1 255.255.255.0
service-policy input mark
!
```

Based upon the MQC configuration shown above, what statement is correct?

A. ip cef must be disabled (using no ip cef) in order for the NBAR classification (match protocol) commands to function.

B. All non-FTP and non-TFTP incoming traffic to the fa0/0 interface will be classified into the class-default traffic class and marked as DSCP 0.

C. All incoming traffic to the fa0/0 interface will be classified into the class-default traffic class and no DSCP marking will be performed.

D. Either FTP or TFTP incoming traffic to the fa0/0 interface will be marked as af11.

E. None of the above.

Answer: C

Explanation:

Based on the configuration above, the service policy named mark will be applied to all traffic incoming on the fast ethernet 0/0 interface. In the policy map, all matching traffic will be assigned the Differentiated Services Code Point of assured forwarding 11. In this case, only traffic that is both FTP and TFTP will match the class-match due to the "match-all" keyword. Since a packet can not be both TFTP and FTP, no traffic will match and the default action will be taken.

Incorrect Answers:

D: This would be true if the "match-any" keyword was used in the "bulk" policy, but in this example the traffic must be both FTP and TFTP, which is not possible.

QUESTION 259

You need to configure Weighted Fair Queuing on router CK1 . In WFQ, you can configure a "congestive-discard-threshold" (CDT). What is the CDT value used for?

- A. This threshold specifies from which point on IOS should start using WFQ.
- B. The CDT specifies the number of messages allowed in each queue.
- C. The CDT specifies the maximum amount of messages to be used by WFQ for high bandwidth traffic, dropping packets from the most aggressive flow.
- D. The CDT defines a value from when IOS starts to account all messages in the WFQ system in conjunction with Netflow.
- E. CDT refers to the maximum amount of dynamic flows IOS will allow for WFQ.
- F. None of the above

Answer: C

Explanation:

WFQ provides traffic priority management that automatically sorts among individual traffic streams without requiring that you first define access lists. WFQ can also manage duplex data streams such as those between pairs of applications, and simplex data streams such as voice or video. There are two categories of WFQ sessions: high bandwidth and low bandwidth. Low-bandwidth traffic has effective priority over high-bandwidth traffic, and high-bandwidth traffic shares the transmission service proportionally according to assigned weights.

Example: The following example requests a fair queue with a congestive discard threshold of 64 messages, 512---dynamic queues, and 18 RSVP queues:

```
interface Serial 3/0
ip unnumbered Ethernet 0/0
fair-queue 64 512 18
```

When WFQ is enabled for an interface, new messages for high-bandwidth traffic streams are discarded after the configured or default congestive messages threshold has been met. However, low-bandwidth conversations, which include control message conversations, continue to enqueue data. As a result, the fair queue may occasionally contain more messages than its configured threshold number specifies.

QUESTION 260

Multiple Certkiller routers are configured for custom queuing to ensure that mission critical traffic gets prioritized. What statement is FALSE regarding custom queuing on these routers? (Select all that apply)

- A. Custom queuing defines 16 queues, not counting the system queue.
- B. With custom queuing you can specify a minimum bandwidth guarantee
- C. In custom queuing there is a strict priority assigned to each queue which specifies how each queue is treated.
- D. Custom queuing has one preemptive priority queue. This can be extended to multiple priority queues by configuring the 'lowest-custom' queue in the 'queue-list'.
- E. In custom queuing you can classify based on the incoming interface.
- F. None of the above

Answer: B, C

Explanation:

B: This is false because with CQ, bandwidth is allocated proportionally for each different class of traffic. CQ allows you to specify the number of bytes or packets to be drawn from the queue, which is especially useful on slow interfaces. Although you can specify the proportion of bandwidth a traffic class can get, you can not specify a minimum bandwidth guarantee.

C: This is false because with custom queuing there is no pre-emptive queue. Bandwidth is statically serviced based on the configuration, and the queue that is being serviced at any given time will finish before servicing the next queue.

QUESTION 261

Using the 3 layer hierarchical approach to a network, What QoS functions are performed at the access layer in the Certkiller network? (Choose 2)

- A. Packet classification
- B. Congestion management
- C. Classification preservation.
- D. Congestion avoidance
- E. Admission control

Answer: A, E

Explanation:

As per Cisco's Hierarchical network model, There are 3 network layers: The Core layer, The distribution layer, and the Access layer. The Core or backbone of the network should not be involved in Processor intensive tasks. Tasks such as packet classification and access control are limited to the access layer and in some cases to the distribution layer. It is the edge routers that classify the QoS traffic, as well as control the QoS admissions into the network.

Incorrect Answers:

B, D. These are functions of the distribution layer. It could be argued that certain aspects of congestion avoidance and management are handled by edge routers, options A and D are better choices.

C. Differentiated Services markings are marked at the access layer edge routers, but preserved throughout the network at the distribution and core layers.

QUESTION 262

You want traffic on the Certkiller frame relay link to conform to specific policies.

Because of this, you configure traffic shaping on a CK router as follows:

Router configuration:

```
ip cef
class-map match-all gold
match ip dscp 10 12 14
class-map match-all bronze
match ip dscp 26 28 30
class-map match-all silver
match ip dscp 18 20 22
policy-map SHAPE
class gold
shape peak 512000
bandwidth percent 50
class bronze
shape average 384000
bandwidth percent 20
class silver
bandwidth percent 30
shape peak 448000
interface Serial4/0
encapsulation frame-relay
ipaddress 14.34.34.51 255.255.255.0
service-policy output SHAPE
end
```

You verify your configuration using the "show policy-map" command as shown below:

Router CertK #sh policy-map inter s4/0

Serial4/0

Service-policy output: SHAPE (1865)

Class-map: gold (match-all) (1866/2)

0 packets, 0 bytes

1 minute offered rate 0 bps, drop rate 0 bps

Match: ip dscp 10 12 15 (1868)

Traffic Shaping

Target Byte Sustain Excess Interval Increment Adapt

Rate Limit bits/int bits/int (ms) (bytes) (active)

1024000 3200 12800 12800 25 3299 -

Queue Packets Bytes Packets Bytes

Depth Delayed Delayed Active

0 0 0 0 no

Weighted Fair Queueing

Output Queue: Conversation 265

Bandwidth 50% Max Threshold 64 (packets)

(pkts matched/bytes matched) 0/0

(pkts discards/bytes discards/tail drops) 0/0/0

Based on this information, what is the CIR value for all the traffic marked with DSCP values 10?

- A. 128000
- B. 256000
- C. 512000
- D. 1024000
- E. Cannot be determined

Answer: C

Explanation:

The value following the "shape peak" command is the CIR. From the configuration above, DSCP value 10 falls under the class gold, which has a CIR value of 512000. To shape traffic to the indicated bit rate according to the algorithm specified, use the shape policy-map class configuration command.

shape [average | peak] mean-rate [[burst-size] [excess-burst-size]]

Syntax Description

average	(Optional) Committed Burst (Bc) is the maximum number of bits sent out in each interval.
peak	(Optional) Bc + Excess Burst (Be) is the maximum number of bits sent out in each interval.
mean-rate	(Optional) Also called committed information rate (CIR). Indicates the bit rate used to shape the traffic, in bits per second. When this command is used with backward explicit congestion notification (BECN) approximation, the bit rate is the upper bound of the range of bit rates that will be permitted.
burst-size	(Optional) The number of bits in a measurement interval (Bc).
excess-burst-size	(Optional) The acceptable number of bits permitted to go over the Be.

QUESTION 263

You need to use NBAR to classify traffic for QoS use in the Certkiller network.

Which of the following is a required configuration parameter for setting up NBAR on a Cisco device in the Certkiller network?

- A. match protocol IP
- B. match nbar type 1
- C. match ftp session passive
- D. match protocol http
- E. match url www.cisco.com
- F. None of the above

Answer: D

Explanation:

Configuring a Traffic Class

To configure a traffic class and the match criteria that will be used to identify traffic as belonging to that class, use the class-map global configuration command. To define the match criteria, use the following commands beginning in global configuration mode.

In the following procedure, all traffic matching a specified protocol will be classified as belonging to the traffic class. The traffic class will classify traffic while the traffic policy configuration will determine how to treat the traffic.

For instance, if you wanted all FTP traffic to be marked with the QoS group value of 1, you would use the match protocol ftp command in class-map configuration mode, and use the set qos-group 1 command in policy-map class configuration mode (assuming the traffic policy uses the specified class). Therefore, the classification purpose (classifying FTP traffic) would be handled in the traffic class, while the QoS feature (marking the QoS group value to 1) would be handled in the traffic policy.

Configuring a Traffic Class with NBAR Example

In the following example, the class-map class1 command uses the NBAR classification of SQL*Net as its matching criterion:

```
Router(config)# class-map class1
Router(config-cmap)# match protocol sqlnet
```

QUESTION 264

The Certkiller network is using QoS to prioritize the critical traffic over busy links. What command would be used to configure Modular QoS CLI (MQC) to allow for a maximum bandwidth of 64 kb/s during times of network congestion; and when there is no congestion, to allow the use of more bandwidth?

- A. bandwidth 64
- B. priority 64
- C. police 64000 conform-action transmit exceed-action drop
- D. shape average 64000
- E. all of the above

Answer: A

Explanation:

MQC is a framework that provides a clear separation between a classification policy and the specification of other parameters that act on the results of that applied classification policy.

Broadly, MQC is configured and implemented as follows:

1. Define a traffic class with the class-map command.
2. Create a service policy by associating the traffic class with one or more QoS features (using the policy-map command).
3. Attach the service policy to the interface with the service-policy command.

To specify the bandwidth to be applied, configure the bandwidth as follows:

Router(config-pmap-c)# bandwidth { bandwidth-kbps percent percent }	Specifies a minimum bandwidth guarantee to a traffic class. A minimum bandwidth guarantee can be specified in kilobits per second or by a percentage of the overall available bandwidth.
--------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Incorrect Answers:

- B. This is used to specify the priority of the traffic, but not the actual bandwidth to be used.
- C. This command configures policing on the interface, so any traffic exceeding the 64 kbps will be dropped, even when there is no congestion.
- D. This is used to specify the average traffic shaping, as specified by the CIR.

Reference:

[http://www.cisco.com/en/US/products/sw/iosswrel/ps5014/products_feature_guide_chapter09186a008008813a.h](http://www.cisco.com/en/US/products/sw/iosswrel/ps5014/products_feature_guide_chapter09186a008008813a.html)

QUESTION 265

Routers CK1 , CK2 , and CK3 are configured in a hub and spoke frame relay environment, with router CK1 as the hub. You have configured Router CK1 , Router CK2 , and Router CK3 to run IGRP over the frame relay connections. No sub-interfaces are used. You have configured a single IP subnet on all the Frame Relay interfaces. Router CK1 can reach both router CK2 and CK3 , but CK2 and CK3 can not reach each other.

What is the probable cause of this problem?

- A. Router CK1 is missing frame maps.
- B. Router CK2 and Router CK3 are not performing frame map updates.
- C. LMI mismatches between routers CK2 and CK3 .
- D. Split-horizon is enabled on Router CK1 .
- E. Split-horizon is disabled on Router CK1 .

Answer: D

Explanation:

The rule of split horizons is the problem with distance vector protocols such as IGRP. The split horizon rule prohibits a router from advertising a route through an interface that the router itself uses to reach that destination. Without sub-interfaces, split-horizon goes into effect, and all routes learned from the Serial interface will not be advertised out of that interface.

Incorrect Answers:

A, B. If the problem was related to missing frame maps or missing updates, then any given location would have issues reaching any location. In this case, router CK2 and CK3 are both able to reach CK1 with no problems.

QUESTION 266

You are troubleshooting a frame relay problem with the serial0 interface on one of your Certkiller routers. When the interface is brought up, it stays up for a short time before it goes back down. You issue the show interface command, and from this you can see that your interface shows LMI status messages sent, but none received. What could be the problem?

- A. There are too many input errors on the line.
- B. The Frame-Relay lmi-type is set incorrectly.
- C. Too many sub-interfaces are exceeding IDB limits.
- D. The DCD not set correctly for a Frame-Relay circuit.
- E. Keepalives are not set correctly on both ends.

Answer: B

Explanation:

In a frame relay configuration, the router's interface always assumes that the connection is up first. Only after missing three consecutive LMI status messages will the interface go down. This explains why the interface shows an "up" status for a short time before going back down. In this case the counters for LMI sent is increasing while the counters for LMI rcvd is still 0. This clearly indicates a case of misconfigured LMI type.

For a detailed discussion on how to troubleshoot serial lines, refer the link below.

http://www.cisco.com/univercd/cc/td/doc/cisintwk/itg_v1/tr1915.htm#xtocid195571

Incorrect Answers:

C. IDB units are Individual Data Blocks, which are units that consume memory resources for each sub-interface that is created. In order to surpass the IDB limits of most routers, thousands of sub-interfaces will need to be created. In addition, after this threshold is met, no more sub-interfaces can be created. Since this question is referring to an already configured router, this is not the problem.

D, E. Although both of these issues could cause problems with the serial lines staying up, they do not explain the fact that LMI status inquiries are received back. Even with keepalives or DCD information set incorrectly, the LMI messages should still be sent and received.

QUESTION 267

The Certkiller WAN consists of a frame relay network using ANSI LMI. What is the maximum theoretical number of DLCI's that can be advertised on a Frame-Relay interface with an MTU of 1500 bytes when using ANSI LMI?

- A. 1024
- B. 1023
- C. 992
- D. 297
- E. 186
- F. 796

Answer: D

Explanation:

The formula for finding the maximum number of DLCI's for ANSI is $(1500-13)/5 = \text{max DLCIs} = 297.4$. See below for the specifics for how this formula is generated:

Analysis

In a PVC information packet, the Report Type (RT) portion is one byte long and the KeepAlive (KA) portion is two bytes long. For the ANSI and Q933a LMIs, the PVC information is 3 bytes long, whereas for the Cisco LMI it is 6 bytes long due to the additional "bw" (for BandWidth) value. The "bw" value represents the Committed Information Rate (CIR); the actual bw value will only be seen if the frame relay switch is configured to forward this information.


The static overhead in each case is 13 bytes [Entire LMI packet minus IEs (10 bytes) + RT (1 byte) + KA (2 bytes)]. We can subtract this number from the Maximum Transmission Unit (MTU) to get the total available bytes for DLCI information. We then divide that number by the length of the PVC IE (5 bytes for ANSI and Q933a, 8 bytes for Cisco) to get the maximum theoretical number of DLCIs for the interface:

For ANSI or Q933a, the formula is: $(\text{MTU} - 13) / 5 = \text{max DLCIs}$.

For Cisco, the formula is $(\text{MTU} - 13) / 8 = \text{max DLCIs}$.

QUESTION 268

The Certkiller frame relay network is displayed in the diagram below, along with the partial configuration of router CK1 :



```
!Hostname CK1
interface Serial0/0
 ip address 192.168.1.1 255.255.255.0
 encapsulation frame-relay
```

What command must be added to interface serial 0/0 of CK1 to allow it to ping CK3 ?

- A. frame-relay inverse-arp ip
- B. frame-relay interface-dlci 302
- C. encapsulation frame-relay ietf
- D. frame-relay map ip 192.168.1.3 102 broadcast
- E. None of the above

Answer: D

Explanation:

The frame relay map command is used to map layer 3 addresses to layer 2 DLCI information. In this case, the router CK1 is configured to statically map IP address 192.168.1.3 to DLCI 102.

Incorrect Answers:

- A. Inverse ARP creates dynamic address mappings, as contrasted with the frame-relay map command, which defines static mappings between a specific protocol address and a specific DLCI. In this case, there is not a PVC that directly connects CK1 and CK3 , so Inverse ARP alone will not be sufficient.
- B. This command should be used on point to point subinterfaces, not on the physical serial interface because it will map all IP addresses to the 302 DLCI, which is incorrect.
- C. This command should be used to connect a Cisco router to a non-Cisco frame relay router.

QUESTION 269

The CK1 frame relay router is configured for frame relay traffic shaping as shown in the diagram below:



hostname CK1

!

interface Serial0/0

bandwidth 384

encapsulation frame-relay

!

interface Serial0/0.101

bandwidth 128

ip address 192.168.1.1 255.255.0

frame-relay interface-dlci 101

class ccie

!

map-class frame-relay ccie

frame-relay cir 128000

frame-relay bc 16000

frame-relay be 0

frame-relay adaptive-shaping becn

Router CK1 is receiving BECNs. What is the lowest rate CK1 will shape its output traffic to?

A. 0 kbps

B. 16 kbps

C. 64 kbps

D. 128 kbps

E. 384 kbps

F. None of the above

Answer: C

Explanation:

The minimum CIR is value is specified by the "frame-relay mincir" command. This command is optional, and if it is omitted from the configuration, the default value is found by dividing the CIR value that is specified by two. In this specific example, it is 128000/2 for a minimum value of 64 kbps, so choice C is correct.

Some additional information on Frame Relay Adaptive Traffic Shaping can be found below:

The Adaptive Frame Relay Traffic Shaping for Interface Congestion feature enhances Frame Relay traffic shaping functionality by adjusting permanent virtual circuit (PVC) sending rates based on interface congestion. When this new feature is enabled, the traffic-shaping mechanism monitors interface congestion. When the congestion level exceeds a configured value called queue depth, the sending rate of all PVCs is reduced to the minimum committed information rate (minCIR). As soon as interface congestion drops below the queue depth, the traffic-shaping mechanism changes the sending rate of

the PVCs back to the committed information rate (CIR). This process guarantees the minCIR for PVCs when there is interface congestion.

Note The sum of the minCIR values for all PVCs on the interface must be less than the usable interface bandwidth.


This new feature works in conjunction with backward explicit congestion notification (BECN) and Foresight functionality. If interface congestion exceeds the queue depth when adaptive shaping for interface congestion is enabled along with BECN or ForeSight, then the PVC sending rate is reduced to the minCIR. When interface congestion drops below the queue depth, then the sending rate is adjusted in response to BECN or ForeSight.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide09186a0080087b91.html

QUESTION 270

The Certkiller frame relay network is shown in the display below:



At which interfaces can the DE bit be set for frame relay packets flowing from CK1 to CK2 ? (Select all that apply)

- A. A
- B. B
- C. C
- D. D
- E. E

Answer: A, B, C, D

Explanation

The frame relay provider's backbone is shared by many users and possibly multiple services. To keep you (and everybody else) from sending more data than the network can hold, frames sent above your contracted rate may be marked as Discard Eligible (DE). DE bits are set by the carrier network, not your equipment. They are also an indication of congestion within the frame relay network, so the DE bits are set on the interior of the carrier network, not at the provider edge to customer edge portion. If your equipment receives DE-marked frames, this indicates that data sent at this rate in the future may get dropped. This may be an early indicator of traffic rates that you didn't plan for in the design of your frame relay WAN.

Frame relay equipment notices congestion when it sees frames marked with the Forward Error Correction Notification (FECN) and Backward Error Correction (BECN) bits.

These merely indicate an overload within the carrier network, and are only of value in monitoring the carrier's health.


The Cisco router can be configured to mark packets as DE using the "frame-relay de-list" command, making choice A also correct.

Incorrect Answers:

E: The DE bits are set between the carrier's frame relay switches, not between the frame relay switches and the customer provided routers. The DE frames are also used only within the network provider, so they would not be marked on interface E since the frame is going directly to the customer router.

QUESTION 271

The Certkiller frame relay network is depicted below:



Traffic from Certkiller 1 to Certkiller 2 is experiencing congestion. What device sets the BECN bit?

- A. Certkiller 1 sets the BECN bit on outgoing packets.
- B. Certkiller 2 sets the BECN bit on outgoing packets.
- C. SwB sets the BECN bit on packets from Certkiller 1 to Certkiller 2.
- D. SwB sets the BECN bit on packets from Certkiller 2 to Certkiller 1.
- E. None of the above

Answer: D

Explanation:

If device A is sending data to device B across a Frame Relay infrastructure and one of the intermediate Frame Relay switches encounters congestion, congestion being full buffers, over subscribed port, overloaded resources, etc, it will set the BECN bit on packets being returned to the sending device and the FECN bit on the packets being sent to the receiving device. This has the effect of telling the sending router to Back off and apply flow control like traffic Shaping and informs the receiving device that the flow is congested and that it should inform upper layer protocols, if possible, that it should close down windowing etc to inform the sending application to slow down.

A FECN tells the receiving device that the path is congested so that the upper layer protocols should expect some delay. The BECN tells the transmitting device that the Frame Relay network is congested and that it should "back off" to allow better throughput.


FECN (Forward Error Congestion Notification)

BECN (Backward Error Congestion Notification)

Reference: <http://www.sins.com.au/network/frame-relay-fecn-becn.html>

QUESTION 272

The Certkiller frame relay network is displayed in the following diagram:



What command must be added to interface serial 0/0 of Certkiller 1 to allow it to ping the Certkiller 3 remote site?

- A. frame-relay inverse-arp ip
- B. framy-relay interface-dlci 302
- C. encapsulation frame-relay ietf
- D. frame-relay map ip 192.168.1.3 102 broadcast
- E. frame-relay map ip 192.168.1.3 302 broadcast


Answer: D

Explanation:

Connecting from Spoke to Spoke:

You cannot ping from one spoke to another spoke in a hub and spoke configuration using multipoint interfaces because there is no mapping for the other spokes' IP addresses. Only the hub's address is learned via the Inverse Address Resolution Protocol (IARP). If you configure a static map using the frame-relay map command for the IP address of a remote spoke to use the local data link connection identifier (DLCI), you can ping the addresses of other spokes. The local DLCI should be specified when using the "frame-relay map" command, which is 102 in this example.

Example:



Configuration:Prasit

```
prasit#show running-configinterface Ethernet0ip address 123.123.123.1
255.255.255.0!interface Serialip address 3.1.3.2 255.255.255.0encapsulation
frame-relayframe-relay map ip 3.1.3.3 150frame-relay interface-dlci 150 Reference:
http://www.cisco.com/en/US/tech/ CK7 13/ CK2 37/technologies_tech_note09186a008014f8a7.shtml
```

QUESTION 273

On router CK1 , you want to view the status of a frame relay connection. Which "show" commands will display the status of a Frame-Relay PVC? (Select all that apply)

- A. show frame relay pvc
- B. show frame-relay pvc
- C. show frame-relay interface
- D. show frame-relay lmi
- E. show frame-relay map
- F. show frame relay interface

Answer: B, E

Explanation:

The following is the example output from the show frame-relay pvc command, which explicitly displays the PVC status:

```
CK1 #show frame-relay pvc
PVC Statistics for interface Serial0 (Frame Relay DCE)
Active Inactive Deleted Static
Local 1 0 0 0
Switched 0 0 0 0
Unused 0 0 0 0
DLCI = 101, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE = Serial0
input pkts 207 output pkts 239 in bytes 15223
out bytes 14062 dropped pkts 0 in FECN pkts 0
in BECN pkts 0 out FECN pkts 0 out BECN pkts 0
in DE pkts 0 out DE pkts 0
out bcast pkts 17 out bcast bytes 3264
PVC create time 00:11:32, last time PVC status changed 00:11:32
Similarly, for show frame-relay map:
CK1 #show frame-relay map
Serial3/1/0.100(D1) (up): point-to-point(D2) dlci, dlci
401(D3)(0x191,0x6410),
broadcast(D4)
status defined, active(D5)
Serial3/1/0.120 (up): point-to-point dlci, dlci 402(0x192,0x6420),
broadcast
status defined, active
```

Incorrect Answers:

A: This is an invalid command, as the Cisco IOS syntax uses frame-relay, not frame relay.

C: The "show frame-relay interface resource" command is a FR to ATM internetworking command that will show PVC stats for Cisco ATM switches. It is not really a valid router IOS command, and it does not show the status of the individual PVC's

D: The following is sample output from the show frame-relay lmi command when the interface is a data terminal equipment (DTE) device:

CK1 # show frame-relay lmi

LMI Statistics for interface Serial1 (Frame Relay DTE) LMI TYPE = ANSI

Invalid Unnumbered info 0 Invalid Prot Disc 0

Invalid dummy Call Ref 0 Invalid Msg Type 0

Invalid Status Message 0 Invalid Lock Shift 0

Invalid Information ID 0 Invalid Report IE Len 0

Invalid Report Request 0 Invalid Keep IE Len 0

Num Status Enq. Sent 9 Num Status msgs Rcvd 0

Num Update Status Rcvd 0 Num Status Timeouts 9

None of the fields above explicitly show the status of the the PVC.

F: This is an invalid command

QUESTION 274

The Certkiller network is utilizing Dynamic Multipoint VPNs. Which of the following is a difference between Phase I and Phase II DMVPN in this network?

- A. Support for multicast
- B. Utilization of Spoke-to-Spoke dynamic tunnels
- C. Utilization of Hub-to-Spoke dynamic tunnels
- D. Utilization of multi-point GRE tunnels at the hub site
- E. None of the above

Answer: B

Explanation:

The 3 DMVPN Phases are:

Phase 1: Hub and spoke functionality

Phase 2: Spoke-to-spoke functionality

Phase 3: Architecture and scaling

Reference:

http://www.cisco.com/application/pdf/en/us/guest/products/ps6658/c1161/cdccont_0900aecd80313c97.pdf

QUESTION 275

Router CK1 is seeing a large number of Discard Eligible (DE) frames. In Frame-relay, as part of the Congestion-Control Mechanism, the DE bit works in conjunction with which of the following mechanism?

- A. Forward Explicit Congestion Notification (FECN) settings
- B. Class-Of-Service (COS) settings

- C. Type-of-service (TOS) settings
- D. Frame-relay Traffic Shaping (FRTS) settings
- E. Differentiated Services Code Point (DSCP) Settings
- F. None of the above

Answer: A

Explanation:

Frame Relay reduces network overhead by implementing simple congestion-notification mechanisms rather than explicit, per-virtual-circuit flow control. Frame Relay typically is implemented on reliable network media, so data integrity is not sacrificed because flow control can be left to higher-layer protocols. Frame Relay implements two congestion-notification mechanisms:

Forward-explicit congestion notification (FECN)

Backward-explicit congestion notification (BECN)

FECN and BECN each is controlled by a single bit contained in the Frame Relay frame header. The Frame Relay frame header also contains a Discard Eligibility (DE) bit, which is used to identify less important traffic that can be dropped during periods of congestion. The FECN bit is part of the Address field in the Frame Relay frame header. The FECN mechanism is initiated when a DTE device sends Frame Relay frames into the network. If the network is congested, DCE devices (switches) set the value of the frames' FECN bit to 1. When the frames reach the destination DTE device, the Address field (with the FECN bit set) indicates that the frame experienced congestion in the path from source to destination. The DTE device can relay this information to a higher-layer protocol for processing. Depending on the implementation, flow control may be initiated, or the indication may be ignored.

Reference:

<http://www.cisco.com/en/US/tech/CK1>

330/tsd_technology_support_technical_reference_chapter09186a00807598

QUESTION 276

Certkiller has chosen to use DMVPN over IPSec tunnels with GRE for secure WAN connectivity. Which of the following are reasons that the Certkiller network administrator would chose to implement DMVPN over IPSEC+GRE for remote site connectivity? (Choose three)

- A. Reduced Configuring at the hub site
- B. Support for dynamically addressed (DHCP) spoke routers
- C. Support for dynamically addressed (DHCP) hub routers
- D. Dynamic Spoke-to-Spoke tunneling
- E. Support for dynamic Spoke Control Protocol (DSCP)
- F. Reduced configuration at the spokes

Answer: A, B, D

Explanation:

No GRE or IPsec information about a spoke is configured on the hub router in the DMVPN network. The spoke router's GRE tunnel is configured (via NHRP commands) with information about the hub router. When the spoke router starts up, it automatically initiates the IPsec tunnel with the hub router as described above. It then uses NHRP to notify the hub router of its current physical interface IP address. This is useful for three reasons:

1. If the spoke router has its physical interface IP address assigned dynamically (such as with ADSL or CableModem via the use of DHCP), then the hub router cannot be configured with this information since each time the spoke router reloads it will get a new physical interface IP address. With DMVPN, spoke routers can still connect dynamically even if they are dynamically addressed through DHCP.
2. Configuration of the hub router is shortened and simplified since it does not need to have any GRE or IPsec information about the peer routers. All of this information is learned dynamically via NHRP.
3. When you add a new spoke router to the DMVPN network, you do not need to change the configuration on the hub or on any of the current spoke routers. The new spoke router is configured with the hub information, and when it starts up, it dynamically registers with the hub router. The dynamic routing protocol propagates the routing information for this spoke to the hub. The hub propagates this new routing information to the other spokes. It also propagates the routing information from the other spokes to this spoke. With the DMVPN solution, one router is the hub, and all the other routers (spokes) are configured with tunnels to the hub. The spoke-to-hub tunnels are up continuously, and spokes do not need configuration for direct tunnels to any of the other spokes. Instead, when a spoke wants to transmit a packet to another spoke (such as the subnet behind another spoke), it uses NHRP to dynamically determine the required destination address of the target spoke. The hub router acts as the NHRP server and handles this request for the source spoke. The two spokes then dynamically create an IPsec tunnel between them (via the single mGRE interface) and data can be directly transferred. This dynamic spoke-to-spoke tunnel will be automatically torn down after a (configurable) period of inactivity.

Reference:

http://www.cisco.com/en/US/tech/CK583/CK372/technologies_white_paper09186a008018983e.shtml#solution

QUESTION 277

Once a dynamic tunnel is built between two spoke sites in a DMVPN scenario utilizing OSPF, what is the neighbor state between the two spoke routers?

- A. 2WAY/DROTHER
- B. FULL/DR
- C. FULL/DROTHER
- D. ESTABLISHED
- E. None of the above

Answer: E

Explanation:

When using the Open Shortest Path First (OSPF) routing protocol, OSPF must be configured to use broadcast network mode. Using broadcast network mode limits the number of hubs in a DMVPN network to just two. Only two hubs may be used because OSPF broadcast networks require a designated router (DR) and a backup designated router (BDR), and each spoke must be connected to both the DR and BDR. This configuration effectively limits the DMVPN network to two hubs when using OSPF. Therefore, there will only be a neighbor relationship between the spoke and the hub, not between the two spokes.

Reference:

http://www.cisco.com/en/US/products/ps6441/products_feature_guide09186a0080641515.html

QUESTION 278

A Certkiller WAN T1 circuit is generating red alarms. What does a red alarm on this T1 indicate?

- A. The CSU cannot synchronize with the framing pattern on the T1 line.
- B. The far end equipment has a problem with the signal it is receiving from the upstream equipment.
- C. There is an alarm on the line upstream from the equipment connected to the port generating the alarm.
- D. There is an alarm from the equipment connected to the port generating the alarm.
- E. The CSU is in a loopback.
- F. None of the above

Answer: A

Explanation:

A RED alarm is known as a Transmit Sending Remote Alarm.

A Red alarm is declared when the channel service unit (CSU) cannot synchronize with the framing pattern on the T1 line.

Reference:

http://www.cisco.com/en/US/tech/CK7_13/CK6_28/technologies_tech_note09186a00801069ff.shtml#topic5

QUESTION 279

On the Certkiller Unified Communications network, what should be used to compress Voice over IP packets on a lower-speed Frame Relay circuit?

- A. Predictor payload compression
- B. RTP header compression
- C. Cisco proprietary payload compression
- D. FRF.9 payload compression
- E. TCP header compression
- F. None of the above

Answer: B

Explanation:

Since VOIP uses the real time protocol (RTP), compressing this type of traffic will be best. RTP is the Internet-standard protocol for the transport of real-time data. It is intended to provide end-to-end network transport functions for applications that support audio, video, or simulation data over multicast or unicast network services.

RTP provides support for real-time conferencing of groups of any size within the Internet. This support includes source identification and support for gateways such as audio and video bridges as well as multicast-to-unicast translators. RTP offers QoS feedback from receivers to the multicast group, as well as support for the synchronization of different media streams.

RTP includes a data portion and a header portion. The data portion of RTP is a thin protocol that provides support for the real-time properties of applications, such as continuous media, including timing reconstruction, loss detection, and content identification.

The header portion of RTP is considerably large. The minimal 12 bytes of the RTP header, combined with 20 bytes of IP header (IPH) and 8 bytes of UDP header, create a 40-byte IP/UDP/RTP header. For compressed-payload audio applications, the RTP packet typically has a 20-byte to 160-byte payload. Given the size of the IP/UDP/RTP header combinations, it is inefficient to transmit the IP/UDP/RTP header without compressing it.

To avoid the unnecessary consumption of available bandwidth, the RTP header compression feature-referred to as CRTP-is used on a link-by-link basis.

RTP can be used over frame relay, HDLC, and PPP links and is meant to be used over slow links (less than 2 Mbps).

QUESTION 280

While troubleshooting an issue with a serial interface on the Certkiller 2 router, you issue the "show interface" command as shown below:

```
certkiller2#sho int ser 0/1
Serial0/1 is up, line protocol is up
Hardware is PowerQUICC Serial
Internet address is 142.16.13.3/24
MTU 1500 bytes, BW 1544 Kbit, DLY 20000 usec,
  reliability 255/255, txload 1/255, rxload 1/255
Encapsulation FRAME-RELAY, loopback not set
Keepalive set (10set)
LMI eng sent 154467, LMI stat recvd 154468, LMI upd recvd 0, DTE LMI up
LMI eng recvd 0, LMI stat sent 0, LMI upd sent 0
LMI DLCI 1023, LMI type is CISCO frame relay DTE
FR SVC disabled, LAPF state down
Broadcast queue 0/64, broadcasts sent/dropped 0/0, interface broadcasts 0
Last input 00:00:02, output 00:00:02, output hang never
Last clearing of "show interface" counters 2w3d
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy hanhy
Output queue: 0/1000/64/0 (size/max total/threshold/drops)
  Conversations 0/1/256 (active/max active/max total)
  Reserved Conversations 0/0 (allocated/max allocated)
  Available Bandwidth 1158 kilobits/sec
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
259810 packets input, 33080103 bytes, 0 no buffer
Received 0 broadcasts, 0 runs, 0 giants, 4 throttles
1 input errors, 0 CRC, 1 frame, 0 overrun, 0 ignored, 0 abort
260370 packets output, 32082918 bytes, 0 underruns
0 output errors, 0 confessions, 1 interface resets
0 output buffer failures, 0 output buffers swapped out
3 carrier transitions
DCD=up, DSR=up, DTR=up, RTS=up, CTS=up
```

Based on the information above, how many times has the interface been reset by the telco service provider?

- A. 0
- B. 1
- C. 3
- D. 4
- E. 1023
- F. None of the above

Answer: C

Explanation:

Carrier transitions appear in the output of the show interfaces serial exec command whenever there is an interruption in the carrier signal (such as an interface reset at the remote end of a link).

Incorrect Answers:

B. The output shows 1 interface reset, but Interface resets that appear in the output of the "show interfaces serial" exec command are the result of missed keepalive packets, and are not indicative of resets sent by the service provider.

Reference: "Troubleshooting Serial Lines"

http://www.cisco.com/univercd/cc/td/doc/cisintwk/itg_v1/tr1915.htm#xtocid7

QUESTION 281

What are the primary reasons to implement traffic shaping on the Certkiller network? (Choose all that apply).

- A. To regulate and thus control the average queue size by indicating when transmission of packets should be halted temporarily.
- B. To control access to available bandwidth on the network.
- C. To define Layer 3 aggregate or granular bandwidth rate limits.
- D. To control the maximum rate of traffic on an interface.
- E. To ensure that traffic conforms to the policies established for it.
- F. To prevent denial of service attacks.
- G. To drop high levels of unwanted traffic.
- H. None of the above

Answer: B, E

Explanation:

According to Cisco, the primary reasons to use traffic shaping are to control access to available bandwidth, to ensure that traffic conforms to specific policies, and to regulate the flow of traffic in order to avoid congestion.

Reference:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4224/sw_config/traffic.htm

QUESTION 282

A Certkiller router's interface is configured for traffic shaping as follows:

```
interface Serial1.1 point-to-point
ip address 10.16.1.1 255.255.255.252
frame-relay class Certkiller
frame-relay interface-dlci 220
!!
map-class frame-relay Certkiller
frame-relay cir 128000
frame-relay bc 8000
frame-relay be 8000
no frame-relay adaptive-shaping
```

In what are the bc and be parameters measured in the above configuration?

- A. Bits per millisecond.
- B. Bits per interval.
- C. Bytes per interval.
- D. Bytes per second.
- E. Bits per second.
- F. Bytes per millisecond.
- G. None of the above

Answer: B

Explanation:

The Sustain (bc) and excess (be) are configured bit per interval.

The following is sample output of the show traffic-shape command:

Target Rate = CIR = 100000 bits/s

Mincir = CIR/2 = 100000/2 = 50000 bits/s

Sustain = Bc = 8000 bits/int

Excess = Be = 8000 bits/int

Interval = Bc/CIR = 8000/100000 = 80 ms


Increment = Bc/8 = 8000/8 = 1000 bytes

Byte Limit = Increment + Be/8 = 1000 + 8000/8 = 2000 bytes

Reference: http://www.cisco.com/warp/public/125/framelay_ts_cmd.html

QUESTION 283

The Certkiller router shown in the following exhibit has a frame relay link with a port speed of 256K and a PVC CIR speed of 64K. The Certkiller router is receiving a notification from the Frame Relay provider that there is congestion in the network.



You want the Certkiller router to react dynamically to this notification from the Frame relay provider. What command should you issue to do this?

- A. traffic-shape adaptive 64000
- B. fair-queue 64000
- C. shape peak 256000 64000
- D. frame-relay class Certkiller
- E. None of the above

Answer: A

Explanation:

When the provider is sending messages to the frame relay customer that there is congestion notifications, they send Backward Explicit Congestion Notification messages (BECNs). As you can see from the definition below, the traffic shape adaptive command

enables the router to react to this:

traffic-shapeadaptive [bit-rate] configures minimum bit rate to which traffic is shaped when backward explicit congestion notifications (BECNs) are received on an interface.

With adaptive GTS, the router uses backward explicit congestion notifications (BECNs) to estimate the available bandwidth and adjust the transmission rate accordingly. The actual maximum transmission rate will be between the rate specified in the traffic-shape adaptive command and the rate specified in the traffic-shape rate command.

As you can see this fulfills the requirement of the question about the Frame Relay network sending information about congestion.

Incorrect Answers:

B, C. These commands will not enable the router to dynamically react to the BECN messages.

D. The Frame-relay class is command is setting up a map class. When a map class is applied to the main interface all the VC gets the traffic shaping from the main interface. This command needs too much assuming while the traffic-rate command does not.

QUESTION 284

In the Certkiller Frame Relay network, what is the maximum theoretical number of DLCI's that can be advertised on a Frame-Relay interface with an MTU of 1500 bytes when using ANSI LMI?

- A. 1024
- B. 1023
- C. 992
- D. 297
- E. 186
- F. 796

Answer: D

Explanation:

The formula for finding the maximum number of DLCI's for ANSI is $(1500-13)/5 = \text{max DLCIs} = 297.4$. See below for the specifics for how this formula is generated:

Analysis

In a PVC information packet, the Report Type (RT) portion is one byte long and the KeepAlive (KA) portion is two bytes long. For the ANSI and Q933a LMIs, the PVC information is 3 bytes long, whereas for the Cisco LMI it is 6 bytes long due to the additional "bw" (for BandWidth) value. The "bw" value represents the Committed Information Rate (CIR); the actual bw value will only be seen if the frame relay switch is configured to forward this information.

The static overhead in each case is 13 bytes [Entire LMI packet minus IEs (10 bytes) + RT (1 byte) + KA (2 bytes)]. We can subtract this number from the Maximum Transmission Unit (MTU) to get the total available bytes for DLCI information. We then divide that number by the length of the PVC IE (5 bytes for ANSI and Q933a, 8 bytes for Cisco) to get the maximum theoretical number of DLCIs for the interface:

For ANSI or Q933a, the formula is: $(MTU - 13) / 5 = \text{max DLCIs}$.

For Cisco, the formula is $(MTU - 13) / 8 = \text{max DLCIs}$.

QUESTION 285

What is the relation between the entities CIR, BC, and TC when considering traffic shaping and frame relay?

- A. $Tc = Bc / CIR$
- B. $CIR = Tc / Bc$
- C. $CIR = Be / TC$
- D. $Tc = CIR / Bc$
- E. None of the above

Answer: A

Explanation:

FRTS Non-Configurable Parameters: interval (Tc) The interval during which you send the Bc bits in order to maintain the average rate of the CIR in seconds.

$Tc = Bc / CIR$ in seconds The range for Tc is between 10 ms and 125 ms. The router internally calculates this value based on the CIR and Bc values in the map class. If Bc/CIR is more than or equal to 125 msec, it uses the internal Tc value. If Bc/CIR is less than 125 ms, it uses the Tc calculated from that equation.

Reference:

http://www.cisco.com/en/US/tech/CK713/CK237/technologies_configuration_example09186a00800942f8.shtml

QUESTION 286

You need to decide if you should configure traffic policing or traffic shaping on router CK1. Which of the following describe differences between these two? (Select two)

- A. Policing uses a token bucket algorithm, shaping use SPD algorithm
- B. Shaping tuned only be applied for ingress traffic policing for egress
- C. with policies you can tune the buffer usage for traffic exceeding the specified CIR
- D. Traffic shaping stores excess traffic in packet buffer until bandwidth is available again
- E. With traffic shaping you can tune the buffer usage for traffic exceeding the specified CIR

Answer: D, E

Explanation:

With traffic shaping, when the excess burst (Be) is configured to a value different than 0, the shaper allows tokens to be stored in the bucket, up to $Bc + Be$. The largest value that the token bucket can ever reach is $Bc + Be$, and overflow tokens are dropped. The only way to have more than Bc tokens in the bucket is to not use all Bc tokens during one or more Tc. Since the token bucket is replenished every Tc with Bc tokens, you can

accumulate unused tokens for later use up to $B_c + B_e$.

In contrast, class-based policing and rate-limiting adds tokens continuously to the bucket. Specifically, the token arrival rate is calculated as follows:

$(\text{time between packets} < \text{which is equal to } t - t_1 > * \text{policer rate}) / 8 \text{ bits per byte.}$

In other words, if the previous arrival of the packet was at t_1 and the current time is t , the bucket is updated with $t - t_1$ worth of bytes based on the token arrival rate. Note that a traffic policer uses burst values specified in bytes, and the above formula converts from bits to bytes.

Reference: "Comparing Traffic Policing and Traffic Shaping for Bandwidth Limiting "
<http://www.cisco.com/warp/public/105/policevsshape.html>

QUESTION 287

Host CK1 wants to join a multicast based video session within the Certkiller network. Hosts that wish to join a multicast group must do which of the following?

- A. Unicast an IGMPv2 membership report to the default router on the local subnet
- B. Multicast an IGMPv2 membership report on the local subnet with a destination IP Address set to the multicast group being joined
- C. Unicast an IGMPv2 membership report to the Rendezvous Point for the group
- D. Multicast an IGMP v2 membership report to the "All-PIM-Routers" multicast group, 224.0.0.2 on the local subnet
- E. None of the above

Answer: B

Explanation:

IGMP is designed to be used by hosts to inform routers that they wish to receive Multicast traffic on specific addresses. In this way, routers can decide whether to forward Multicast traffic based on whether a host on a given subnet has requested this or not. In addition, some vendors such as Cisco, extend this functionality by having routers share this information with switches so that the switches will only forward the Multicast traffic to ports with hosts that have requested it. Without this feature, the traffic would effectively be broadcast traffic.

When a host joins a Multicast session, it sends out an IGMPv2 packet to let any listening routers know that it wants to receive Multicast traffic sent to a particular address. This packet is addressed to the Multicast address that the host wants to join. This is called "Joining a Multicast Group". Similarly, when the session has ended, the host sends out another IGMP packet to "Leave the Multicast Group".

QUESTION 288

The Certkiller network administrator plans to migrate their IP multicast network from IGMP version 2 to IGMP version 3. IGMPv3 made which of the following major functionality changes over IGMPv2? (Choose three)

- A. IGMPv3 added the ability for a host to specify which sources in a multicast group it does not wish to receive

- B. IGMPv3 added the ability for a host to specify which sources in a multicast group it wishes to receive
- C. IGMPv3 added "Request-to-send", "Clear-to-send" Signaling between sources and the local IGMP Querier
- D. IGMPv3 removed the "report-suppression" feature for IGMP Membership Reports
- E. IGMPv3 removed the ability to do a wildcard join of all sources in a multicast group

Answer: A, B, D

Explanation:

Table of IGMP Versions

IGMP Version	Description
IGMPv1	Provides the basic Query-Response mechanism that allows the multicast router to determine which multicast groups are active and other processes that enable hosts to join and leave a multicast group. RFC 1112 defines Host Extensions for IP Multicasting.
IGMPv2	Extends IGMP, allowing such features as the IGMP leave process, group-specific queries, and an explicit maximum query response time. IGMPv2 also adds the capability for routers to elect the IGMP querier without dependence on the multicast protocol to perform this task. (See RFC 2236.)
IGMPv3	Provides for source filtering, which enables a multicast receiver host to signal to a router which groups it wants to receive multicast traffic from, and from which sources this traffic is expected. In addition, IGMPv3 supports a link local address, 224.0.0.22, which allows routing protocols to communicate with each other.

Note:


IGMP report suppression is supported only when the multicast query has IGMPv1 and IGMPv2 reports. This feature is not supported when the query includes IGMPv3 reports.

References:

http://www.cisco.com/en/US/products/ps6350/products_configuration_guide_chapter09186a00805a344c.html
http://www.cisco.com/en/US/docs/switches/lan/catalyst2970/software/release/12.2_25_se/configuration/guide/sw

QUESTION 289

Part of the Certkiller IP multicast network is shown below:



H1, H2, H3, and rtr-a are all IGMP version 2 devices. Host 2 and Host 3 belong to the 224.1.1.1 group. After a while, H2 sends out an IGMPv2 Leave message to leave the 224.1.1.1 group. How will rtr-a react to this leave message?

- A. It will send an IGMPv2 Query to the all multicast hosts address 224.255.255.255.
- B. It will send an IGMPv2 Group Specific Query to 224.1.1.1
- C. It will send an IGMPv2 Leave Acknowledgement to Hosts H1 and H3.
- D. It will send an IGMPv2 General Query to 224.1.1.1
- E. It will send an IGMPv2 Group Specific Query to 224.0.0.1.

Answer: B

Explanation:


In IGMP version 2, a Leave message is responded by a group specific query from the router to check if there are any additional hosts participating in the multicast session. The group specific query is always destined for the multicast address that is being used.

Incorrect Answers:

- A. The address 224.255.255.255 would never be used in this situation. In fact, the notion of a multicast "broadcast" does not exist.
- C. Leave messages are not acknowledged.
- D. General Query messages are not used.
- E. The group specific query is always destined for the multicast address that is being used, which is 224.1.1.1 in this case.

QUESTION 290

The Certkiller network has a mix of IGMP version 1 and version 2 devices in its IP multicast network as shown below:



H1 and H2 are both IGMPv2 speakers and are also members of group 224.1.1.15. H3 is an IGMPv1 speaker and sends an IGMPv1 Membership Report to join group 224.1.1.15. What will happen?

- A. The router rtr-a will do nothing, since there are already members of group 224.1.1.15 on the subnet.
- B. The router rtr-a will ignore all IGMPv2 Leave messages while the IGMPv1 host is a member of group 224.1.1.15.
- C. The router rtr-a will stop sending IGMPv2 Group-Specific queries in response to IGMPv1 Leaves received on this subnet for groups 224.1.1.15, while the IGMPv1 hosts is a member of group 224.1.1.15.
- D. The router rtr-a will ignore the IGMPv1 Membership Report because router rtr-a is an IGMPv2 speaker and IGMPv2 and IGMPv1 are not compatible.

Answer: B

Explanation:

With IGMP version 1 and version 2 on the same network, routers will revert to v1, so the router will ignore the leave requests from all v2 members as long as the v1 member is still active for that multicast session.

Incorrect Answers:

- A. Although there are already members on the same segment, the routers must be aware of the fact that there are a mix of v1 and v2 devices, so that the v2 leave messages can be ignored.
- C. When the v1 device leaves the multicast session, the router must still send the group query out to see if the v2 devices are also still actively participating in the multicast session.
- D. IGMP version 2 was designed to be backward compatible with version 1.

Reference:

"CCIE Professional Development Routing TCP/IP Volume II" by Jeff Doyle and Jennifer De Haven Carroll, Page 414.

QUESTION 291

IGMP is used to dynamically register individual hosts in a multicast group on a particular LAN. Hosts identify group memberships by sending IGMP messages to their local multicast router. Under IGMP, routers listen to IGMP messages and periodically send out queries to discover which groups are active or inactive on a particular subnet.

Hosts need to actively communicate to the local multicast router that they intend to leave a group. If there are no replies, the router times out the group and stops forwarding the traffic. In order for this to work, what needs to be implemented?

- A. IGMPv1
- B. IGMPv2
- C. IGMPv3
- D. IGMPv4
- E. CGMP

Answer: B

Explanation:

The Internet Group Management Protocol (IGMP) is used by IP hosts to report their host group memberships to any immediately neighboring multicast routers. IGMP messages are encapsulated in IP datagrams, with an IP protocol number of 2. IGMP has versions IGMP v1, v2 and v3.

In IGMPv2, leave messages were added to the protocol. This allowed group membership termination to be quickly reported to the routing protocol, which is important for high-bandwidth multicast groups and/or subnets with highly volatile group membership.

Incorrect Answers:

- A. IGMPv1: Hosts can join multicast groups. There were no leave messages. Routers were using a time-out based mechanism to discover the groups that are of no interest to the members.
- C. IGMPv3: Major revision of the protocol. It allows hosts to specify the list of hosts from which they want to receive traffic from. Traffic from other hosts is blocked inside the network. It also allows hosts to block inside the network packets that come from sources that sent unwanted traffic.
- D. IGMPv4 is not yet in use.
- E. CGMP is the Cisco Group Management Protocol (CGMP) which is a multicast protocol used by Cisco LAN switches, and not routers.

QUESTION 292

In the Certkiller network, hosts need to actively communicate to the local multicast router that they intend to leave a group. The router then sends out a group-specific query and determines if any remaining hosts are interested in receiving the traffic. If there are no replies, the router times out the group and stops forwarding the traffic. In order for this to work, what needs to be implemented?

- A. IGMPv1
- B. IGMPv2
- C. IGMP snooping
- D. DVMRO
- E. CGMP
- F. RGMP

Answer: B

Explanation:

IGMP version 2 is the Industry-standard protocol for managing multicast group membership, including support for IGMP-leave messages and group-specific queries. Leave Group message is a new type different from IGMP version 1. Membership Report is issued by host that want to join a specific multicast group (GDA). When IGMP router receive the Membership Report, it will add the GDA to the multicast routing table and start forwarding the IGMP traffic to this group. Membership Queries are issued by router at regular intervals to check whether there is still a host interested in the GDA in that segment. Host Membership Reports are sent either when the host wants to receive GDA traffic or response for a membership query from IGMP router.

If a host does not want to receive the IGMP traffic any more, it sends a Leave Group message. When the multicast router receives this Leave Group message, it removes the GDA from the multicast routing table. In addition, IGMP multicast routers periodically send Host Membership Query messages (hereinafter called Queries) to discover which host groups have members on their attached local networks. If no Reports are received for a particular group after some number of Queries, the routers assume that that group has no local members and that they need not forward remotely-originated multicasts for that group onto the local network. In addition, IGMP version 2 has leave mechanisms.

Incorrect Answers:

A. In IGMP version 1, hosts can join multicast groups. There were no leave messages. Routers were using a time-out based mechanism to discover the groups that are of no interest to the members.

D, F. These are incorrect terms for this IP multicasting functionality.

E. CGMP is used between Cisco switches and routers to provide for IP multicast information to be passed between the two.

QUESTION 293

You want to optimize the Certkiller IP Multicast network on the LAN. Choose the correct protocols to handle IP multicast efficiently in this layer 2 switched network.

- A. Use IGMP Snooping on subnets that include end users or receiver clients and routed segments that contain only routers, such as in a collapsed backbone.
- B. Use Router-Port Group Management Protocol (RGMP) on subnets that include end users or receiver clients. Use Cisco Group Management Protocol (CGMP), IGMP Snooping on routed segments that contain only routers, such as in a collapsed backbone.
- C. Use Cisco Group Management Protocol (CGMP) on subnets that include end users or

receiver clients and routed segments that contain only routers, such as in a collapsed backbone.

D. Use Router-Port Group Management Protocol (RGMP) on subnets that include end users or receiver clients and routed segments that contain only routers, such as in a collapsed backbone.

E. Use Cisco Group Management Protocol (CGMP) and IGMP Snooping on subnets that include end users or receiver clients. Use Router-Port Group Management Protocol (RGMP) on routed segments that contain only routers, such as in a collapsed backbone.

F. None of the above

Answer: E

Explanation:

Explanation:

The purpose of Cisco Group Management Protocol (CGMP) and Internet Group Management Protocol (IGMP) snooping is to restrain multicast traffic in a switched network. By default, a LAN switch floods multicast traffic within the broadcast domain. This can consume a lot of bandwidth if many multicast servers are sending streams to the segment.


IGMP snooping is a feature that allows the switch to "listen in" on the IGMP conversations between hosts and routers. When a switch hears an IGMP report from a host for a given multicast group, the switch adds the host's port number to the GDA list for that group. And, when the switch hears an IGMP Leave, it removes the host's port from the CAM table entry.

RGMP constrains multicast traffic that exits the Cisco Router through ports to which only disinterested multicast routers are connected. RGMP reduces network congestion by forwarding multicast traffic to only those routers that are configured to receive it.

Note: To use RGMP, you must enable IGMP snooping on the Cisco router. IGMP snooping constrains multicast traffic that exits through LAN ports to which hosts are connected. IGMP snooping does not constrain traffic that exits through LAN ports to which one or more multicast routers are connected.

QUESTION 294

Exhibit:



Switch A is a catalyst switch running IGMP snooping. In the Certkiller IP Multicast network shown, which routers will multicast stream (239.1.2.3)? (Choose two).

- A. 1
- B. 2
- C. 3
- D. IGMP snooping does not know receiver in vlan 33, so switch A will drop the multicast stream.
- E. Router 2 only after a PC joins group 239.1.2.3

Answer: A, C

Explanation:

IGMP Snooping is the process of listening to IGMP traffic. IGMP snooping, as implied by the name, is a feature that allows the switch to "listen in" on the IGMP conversation between hosts and routers by processing the layer 3 IGMP packets sent in a multicast network.

When IGMP snooping is enabled in a switch it analyses all the IGMP packets between hosts connected to the switch and multicast routers in the network. When a switch hears an IGMP report from a host for a given multicast group, the switch adds the host's port number to the multicast list for that group. And, when the switch hears an IGMP Leave, it removes the host's port from the table entry.

IGMP snooping can very effectively reduce multicast traffic from streaming and other bandwidth intensive IP applications. While a switch that does not understand multicast will broadcast multicast traffic to all the ports in a broadcast domain (a LAN), a switch using IGMP snooping will only forward multicast traffic to the hosts interested in that traffic. This reduction of multicast traffic reduces the packet processing at the switch (at the cost of needing additional memory to handle the multicast tables) and also reduces the workload at the end hosts since their network cards (or operating system) will not have to receive and filter all the multicast traffic generated in the network.

In this case, since the router interfaces on 1 and 3 are configured to use sparse-dense mode, they will both participate in streaming multicast traffic from 239.1.2.3.

QUESTION 295

Devices on the Certkiller Multicast network need to send source specific multicast (SSM) messages. What range of IP Multicast Addresses are reserved by the IANA for source-specific multicast?

- A. 232.0.0.0-232.0.0.255
- B. 233.0.0.0-233.255.255.255
- C. 232.0.0.0-232.255.255.255
- D. 239.0.0.0-239.255.255.255
- E. None of the above

Answer: C

Explanation:

IANA Assigned Multicast Address Blocks

The IETF has provided the IANA with guidance on how IP Multicast address space should be allocated in RFC 3171bis, "IANA Guidelines for IPv4 Multicast Address Assignments." Table 1 below lists the current assignments blocks documented in RFC 3171bis.

Table 1 IANA Multicast Address Assignments		
Range	Mask	Description
224.0.0.0-224.0.0.255	224.0.0/24	Local Network Control Block
224.0.1.0-224.0.1.255	224.0.1/24	Internetwork Control Block
224.0.2.0-224.0.255.255	-	Ad hoc Block
224.1.0.0-224.1.255.255	-	Unassigned
224.2.0.0-224.2.255.255	224.2/16	SDP/SAP Block
224.3.0.0-231.255.255.255	-	Unassigned
232.0.0.0-232.255.255.255	232/8	Source Specific Multicast Block
233.0.0.0-233.255.255.255	233/8	GLOP Block
234.0.0.0-238.255.255.255	-	Unassigned
239.0.0.0-239.255.255.255	239/8	Administratively Scoped Block

Reference:

http://www.cisco.com/en/US/tech/CK828/technologies_white_paper09186a00802d4643.shtml

QUESTION 296

Routers CK1 and CK2 are IP Multicast routers. These routers use (S, G) entries for multicast packets forwarding. Which of the following address types are used in the "S" entry?

- A. Source Specific multicast addresses
- B. The block of administratively scoped multicast addresses
- C. SDP / SAP addresses
- D. Any class A, Class B or Class C host address
- E. GLOP addresses

Answer: D

Explanation:

State entries for a source tree use the notation (S, G) pronounced S comma G. The letter S represents the IP address of the source (any valid class A, B, or C host IP address), and G represents the group address.

Reference: <http://www.ciscopress.com/articles/article.asp?p=32100>

QUESTION 297

A multicast Application is being deployed within the Certkiller network. The scope of the application's multicast traffic is to remain entirely within the Certkiller network. From what address range should the multicast address be assigned for this application?

- A. The network administrator should pick ANY IP multicast address for use by the application since the application scope is entirely within the Enterprise network and will conflict with Global Internet multicast
- B. The network administrator should assign an address from the administratively scoped address range (239.0.0.0-239.255.255.255 to the application)
- C. The network administrator should assign an address in the 232.0.0.0 - 232.255.255.255 address range to the application
- D. The Enterprise should apply to IANA to have an address permanently assigned to this application
- E. None of the above

Answer: B

Explanation:

RFC 2365-Administratively Scoped Addresses

RFC 2365 provides limited guidelines on how the multicast address space can be divided and used privately by enterprises. The terminology "Administratively Scoped IPv4 multicast space" relates to the group address range of 239.0.0.0 to 239.255.255.255. The key properties of Administratively Scoped IP Multicast are that:

Packets addressed to Administratively Scoped multicast addresses do not cross configured administrative boundaries. The limits of these scope boundaries often are called "Zones" or "Scoped Zones."

Administratively Scoped multicast addresses are locally assigned and are not required to be unique across administrative boundaries.

Reference:

http://www.cisco.com/en/US/tech/ CK8 28/technologies_white_paper09186a00802d4643.shtml#wp1002679

QUESTION 298

IP multicast addresses in the range of 224.0.0.0 through 224.0.0.255 are reserved for what purpose?

- A. It is reserved for Administratively Scoped multicast traffic intended to remain inside a

private network.

B. It is reserved for Administratively Scoped multicast traffic that is not supposed to be transmitted onto the Internet.

C. It is reserved for link-local multicast traffic consisting of network control messages that is not supposed to leave the local subnet.

D. Any valid multicast data stream used by multicast applications.

E. Global Internet multicast traffic intended to travel throughout the Internet.

Answer: C

Explanation:

As found in RFC1112. These addresses are used by many routing protocols such as OSPF and RIPv2, in order to sent updates to all neighbors on the same segment.

Incorrect Answers:

Administratively Scoped IP multicast addresses are contained in the 239.0.0.0-239.255.255.255 range. (Not A, Not B)

The 224.0.0.0/8 network range is not intended to be used outside of the local subnet link. (Not D, Not E)

QUESTION 299

What is the class D IP address range 239.0.0.0-239.255.255.255 used for?

A. Administratively Scoped multicast traffic meant for internal use.

B. Link-local multicast traffic made up of network control messages meant to stay in the local subnet.

C. Global Internet multicast traffic meant to travel throughout the Internet.

D. Any valid multicast data stream for use with multicast applications.

E. Routing protocol use.

Answer: A

Explanation:

The 239 address range is reserved for IP multicast traffic that is to be used for internal use only. It is similar to RFC 1918 private IP address space, except instead of specifying unicast address ranges it specifies multicast.

Incorrect Answers:

B, E. Link level multicast messages, such as those used by routing protocols, use the 224.0.0.0 address range. For example, IGRP uses 224.0.0.10 and OSPF uses 224.0.0.5 and 224.0.0.6.

C, D. This address range should never be seen in the Internet. It is reserved for private use only.

Reference:

Jeff Doyle Volume II chapter on IP Multicast.

QUESTION 300

You wish to implement a multicast video application over your private, internal network.

To do this, you need to use a private multicast range of IP addresses across your network. Which IP range should you use?

- A. 224.0.0.0 - 224.255.255.255
- B. 226.0.0.0 - 226.255.255.255
- C. 241.0.0.0 - 241.255.255.255
- D. 239.0.0.0 - 239.255.255.255
- E. 240.0.0.0 - 254.255.255.255.

Answer: D

Explanation:

The reserved, administratively scoped IPv4 multicast address space is defined to be the range 239.0.0.0 to 239.255.255.255. Administratively scoped multicast addresses are for use only on a private network and are not to be used on the Internet.

Reference:

RFC 2365 - <http://www.faqs.org/rfcs/rfc2365.html>

QUESTION 301

The Certkiller network is implementing IP multicast and they want to ensure that the IP addresses they used are contained within the Certkiller autonomous system. What is the range of limited scope/administrative scope addresses that should be used?

- A. Addresses in the 232.0.0.0/8 range
- B. Addresses in the 239.0.0.0/8 range
- C. Addresses in the 224.0.0.0/8 range
- D. Addresses in the 229.0.0.0/8 range
- E. Addresses in the 234.0.0.0/8 range
- F. None of the above

Answer: B

Explanation

The range of addresses from 239.0.0.0 through 239.255.255.255 contains limited scope addresses or administratively scoped addresses. These are defined by RFC 2365 to be constrained to a local group or organization. Routers are typically configured with filters to prevent multicast traffic in this address range from flowing outside an autonomous system (AS) or any user-defined domain. Within an autonomous system or domain, the limited scope address range can be further subdivided so those local multicast boundaries can be defined. This also allows for address reuse among these smaller domains. These addresses are the IP multicast version of the private, RFC 1918, addresses used for unicast.

Reference: http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/ipmulti.htm

QUESTION 302

In the Certkiller IP multicast network, what range of IP Multicast addresses are used for Scope-Relative Multicast?

- A. The lowest (numerically) 256 multicast addresses of each Administratively scoped address range are automatically reserved for Scope-Relative multicast
- B. Scope-Relative Multicast Addresses must be chosen from the Administratively scoped address range by the network administrator and configured on every router
- C. The highest (numerically) 256 addresses of each administratively scoped address range are automatically reserved for scope relative multicast
- D. The highest (numerically) 32 addresses of each administratively scoped address range are automatically reserved for scope-relative multicast
- E. None of the above

Answer: C

Explanation:

Multicast addresses may be allocated in any of three ways:

Static:

Statically allocated addresses are allocated by IANA for specific protocols that require well-known addresses to work. Examples of static addresses are 224.0.1.1 which is used for the Network Time Protocol and 224.2.127.255 which is used for global scope multicast session announcements.

Scope-relative:

RFC 2365 reserves the highest 256 addresses in every administrative scope range for relative assignments. Relative assignments are made by IANA and consist of an offset which is valid in every scope.


Dynamic:

For most purposes, the correct way to use multicast is to obtain a dynamic multicast address. These addresses are provided on demand and have a specific lifetime.

Reference: <http://www.ietf.org/rfc/rfc2908.txt>

QUESTION 303

The Certkiller network is shown in the following exhibit:



Receivers

Router CK1 is configured as follows:

```
ipmulticast-routing
interface loopback0
ip address 192.168.1.1 255.255.255.0
ip pim send-RP-announce loopback0 scope 16 group-list 1
ip pim send-RP-discovery loopback0 scope 16
access-list 1 permit 239.0.0.0 0.255.255.255
```

Router CK2 is configured as follows:

```
ipmulticast-routing
interface loopback 0
ip address 192.168.11.1 255.255.255.0
ip pim send-RP-announce loopback0 scope 16 group-list 1
ip pim send-RP-discovery loopback0 scope 16
access-list 1 permit 239.0.0.0 0.255.255.255
```

Which of the routers will take on the function of Mapping Agent and source Auto-RP Discovery messages to the 224.0.1.40 group?

- A. Router CK1
- B. Router CK2
- C. Both Router CK1 and Router CK2
- D. Neither, since the access lists configured do not match 224.0.1.40 multicast traffic.

Answer: C

Explanation:

If several RPs announce themselves for a multicast group range, the mapping agent chooses only one, which is the RP with the highest IP address. However, this is for selecting the RP. There is no election process for selecting the mapping agent that will source auto-RP discovery message. Both A and B will source this message.

Incorrect Answers:

A, B. If only one router were elected as a mapping agent, this would adversely affect the other source, since it would not have a mapping agent.

B. This would be the correct choice if the question were related to the RP election, and not the mapping agent election. When multiple routers contend to be the Rendezvous Point, the router with the highest IP address wins the tie-breaker and will be elected as the RP. However, there can be multiple mapping agents in the network, as would be the case in this situation.

D. All PIM-enabled routers automatically join the Cisco RP discovery group (224.0.1.40) that allows them to receive all group-to-RP mapping information. This information is distributed by an entity called RP mapping agent. Therefore, the access list is irrelevant in this case.

QUESTION 304

Which IP address maps to the Ethernet multicast MAC address of 01-00-5e-10-20-02?
(Choose all that apply)

- A. 224.128.10.2
- B. 225.128.10.2
- C. 224.10.20.2
- D. 225.10.20.2
- E. 239.144.32.2
- F. 224.16.32.2
- G. All of the above
- H. None of the above

Answer: E, F

Explanation:

Ethernet interfaces map the lower 23 bits of the IP multicast address to the lower 23 bits of the MAC 0100.5e00.0000. As an example, the IP multicast address 224.0.0.2 is mapped to the MAC layer as 0100.5e00.0002.

1. HEX 01 = 00-5e (all Multicast Addresses);
 2. HEX 10 = 00010000 - could be both 16 and 144 (decimal) due to the fact that we ignore the first bit of the second octet when converting to binary;
 3. HEX 20 = 00100000 = 32;
 4. HEX 02 = 00000010 = 2.
-

QUESTION 305

Certkiller owns a block of Ethernet MAC address that start with 01:00:5E in hexadecimal format. Half of this block is allocated for multicast addresses. The range from 0100.5e00.0000 through 0100.5e7f.ffff is the available range of Ethernet MAC address for IP multicast.

This allocation allow for 23 bits in the Ethernet address to correspond to the IP multicast group address. The mapping places the lower 23 bits of the IP multicast group address into these available 23 bits in the Ethernet address.

Because the upper five bits of the IP multicast address are dropped in this mapping, the resulting address is not unique. In fact, 32 different multicast group IDs map to the same Ethernet address.

225.1.1.1 and 237.1.1.1 have been assigned to map to the same multicast MAC address on a Layer 2 switch. In a multicast network, what will occur from this?

- A. If one user is subscribed to Group A (as designated by 225.1.1.1) and the other user is subscribed to Group B (as designated by 237.1.1.1), they would both receive only the streams meant for them. Group A would go to 225.1.1.1 and group B would go to 237.1.1.1
- B. If one user is subscribed to Group A (as designated by 237.1.1.1) and the other user is subscribed to Group B (as designated by 225.1.1.1), they would both receive only the first stream that reached the network.
- C. If one user is subscribed to Group A (as designated by 225.1.1.1) and the other user is subscribed to Group B (as designated by 237.1.1.1), they would both receive both streams, A and B streams.
- D. If one user is subscribed to Group A (as designated by 225.1.1.1) and the other user is subscribed to Group B (as designated by 237.1.1.1), both of them would not receive A and B streams.
- E. None of the above

Answer: C

Explanation:

Although mathematically there are 32 possibilities for overlap of addresses it is very unlikely to happen in real life. If it does, the impact is that another set of stations receives the multicast traffic. This is still far preferable to ALL stations receiving the traffic. This is always the case where two IP multicast addresses share the same MAC address.

QUESTION 306

Certkiller .com runs a large IP multicast network with thousands of sources and thousands of groups and uses (S, G) entries for forwarding. The applications that are using IP multicast do not require a minimum latency and there is a severe impact on resources on routers and high memory consumption from the size of the multicast routing table.

What would be the right solution in this particular scenario which will decrease the resource issues on the routers, reduce the amount of memory needed by the large multicast routing tables and minimize the amount of state in each router?

- A. Continue using (S, G) entries but add a rendezvous point (RP) in the topology
- B. Use (*,G) entries with source trees and a rendezvous point (RP) in the topology
- C. Use shared trees with a rendezvous point (RP) in the topology
- D. Use combination of source trees and shared trees without rendezvous point (RP) in the topology
- E. Use PIM Sparse mode with (S,G) and (*,G) entries

Answer: C

Explanation:

Shortest path trees have the advantage of creating the optimal path between the source and the receivers. This guarantees the minimum amount of network latency for forwarding multicast traffic. This optimization does come with a price, though: The routers must maintain path information for each source. In a network that has thousands of sources and thousands of groups, this can quickly become a resource issue on the routers. Memory consumption from the size of the multicast routing table is a factor that network designers must take into consideration.

Unlike source trees that have their root at the source, shared trees use a single common root placed at some chosen point in the network. This shared root is called the rendezvous point (RP). Shared trees have the advantage of requiring the minimum amount of state in each router. This lowers the overall memory requirements for a network that allows only shared trees. The disadvantage of shared trees is that, under certain circumstances, the paths between the source and receivers might not be the optimal paths-which might introduce some latency in packet delivery. Network designers must carefully consider the placement of the RP when implementing an environment with only shared trees.

Incorrect Answers:

A, E: Because of the potentially large number of different multicast sources in this particular network, the use of individual (S, G) entries should be avoided.
B, D: The simplest form of a multicast distribution tree is a source tree whose root is the source of the multicast tree and whose branches form a spanning tree through the network to the receivers. Because this tree uses the shortest path through the network, it is also referred to as a shortest path tree (SPT). The shortest-path tree requires more memory than the shared tree, but reduces delay. Because we want to reduce the amount of memory needed, these choices are incorrect.

Reference:

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/ipmulti.htm#xtocid18

QUESTION 307

In the Certkiller IP multicast network where many sources which are also receivers, what protocol is used to allow the use of the same shared tree for traffic from sources towards RP and from RP towards receivers.

- A. PIM Dense Mode (PIM DM)
- B. Multicast Open Shortest Path First (MOSPF)

- C. Distance Vector Multicast Routing Protocol (DVMRP)
- D. Bidirectional PIM
- E. PIM Sparse Mode (PIM SM)
- F. None of the above

Answer: D

Explanation:

Bidirectional-PIM is a variant of the Protocol Independent Multicast (PIM) suite of routing protocols for IP multicast. In PIM, packet traffic for a multicast group is routed according to the rules of the mode configured for that multicast group. The Cisco IOS implementation of PIM supports three modes for a multicast group:

1. Bidirectional mode
2. Dense mode
3. Sparse mode

A router can simultaneously support all three modes or any combination of them for different multicast groups. In bidirectional mode, traffic is routed only along a bidirectional shared tree that is rooted at the rendezvous point (RP) for the group. In bidir-PIM, the IP address of the RP acts as the key to having all routers establish a loop-free spanning tree topology rooted in that IP address. This IP address need not be a router, but can be any unassigned IP address on a network that is reachable throughout the PIM domain.

Figure1: Unidirectional Shared Tree and Source Tree

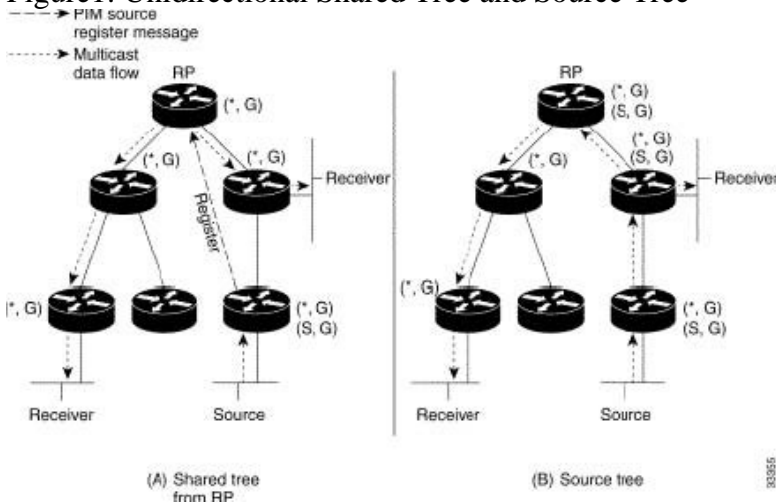



Figure2 Bidirectional Shared Tree



Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1612/products_feature_guide09186a0080080a41.html

QUESTION 308

Certkiller runs their core enterprise network as an ISP network where they have different Autonomous Systems (AS). The BGP core runs OSPF for Intra-connection only. Data center A is in AS 1, data center B is in AS 2, and data center C is in AS 3. The remote locations will be running an IGP and redistribute their routes into BGP core. They would like to enable multicast throughout their network to support multicast applications.

Based upon the scenario, what would be the LEAST EFFECTIVE way to implement IP multicast?

- A. This network runs essentially as an ISP's network with a BGP core and different AS. To implement multicast in this network they can enable MBGP over the BGP backbone.
- B. This is customer's internal network and not a transit provider in the inter-domain SP routing. As long as there is no incongruence (between multicast and unicast topologies), there is no need to run MBGP. They simply run PIM-SM and MSDP for redundancy.

C.

Running MBGP, besides BGP, should present negligible overhead and if done together with the introduction of IP multicast will help to avoid problems later on when the network has grown and some incongruence needs to be supported. At that point, the customer may need to upgrade to MBGP throughout the network to have the transitive nature of incongruence supported correctly, and this may then become an obstacle in deployment. Therefore, MBGP should be implemented.

D. It should be determined what IP multicast applications the customer is intending to run. Source Specific Multicast (SSM) should be recommended to the customer, since it would allow them to overcome MSDP and thus reduce the complexity of IP multicast in their deployment.

E. PIM uses the unicast routing information to perform the multicast forwarding function. They can simply implement Inter AS PIM (IAPIM) to exchange the multicast routing information. This would be the easiest way to implement multicast in the current network where they leverage all the current unicast routing protocol information to populate the multicast routing table, including Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest path First (OSPF), Border Gateway Protocol (BGP), and static routes. This approach would also cause less processing on the routers as PIM does not send and

receive routing updates between routers.

Answer: B

Explanation:

The Multicast Source Discovery Protocol (MSDP) is a mechanism to connect multiple PIM sparse-mode (SM) domains. MSDP allows multicast sources for a group to be known to all rendezvous point(s) in different domains. Each PIM-SM domain uses its own rendezvous points and does not need to depend on them in other domains. A rendezvous point runs MSDP over TCP to discover multicast sources in other domains. MSDP is also used to announce sources sending to a group. These announcements must originate at the domain's Rendezvous Point. MSDP depends heavily on MP-BGP for interdomain operation. Because of this, choice B is the least effective choice since it recommends running MSDP without MGBP.

QUESTION 309

In an IP multicast network, the more sources an application has, the less frequently traffic is sent from each end. Each time a source starts to send packets, protocol operations take place and a forwarding state is established. For applications with a large number of sources, this state can time-out before the source would only create a large number of sources, this state can time-out before sources would not only create a large amount of forwarding state (requiring memory), but they could also require high CPU usage of the routing processor due to the accounting of frequently changing state. In addition, the signaling within the router between the routing processor and forwarding hardware can become another potential bottleneck of continuously large amount of traffic signaling must go to the routing processor and equally large amounts of forwarding state changes must go to the forwarding engine(s).

The Certkiller network is implementing IP multicast, and they wish to avoid the problems described above. Based on this information, what IP multicast technology would you recommend?

Caution: This protocol should avoid maintaining source-specific forwarding state, thereby reducing the amount of memory needed by the number of sources per multicast group, requiring much less traffic signaling in the protocol, preventing the "bursty source" problem, saving on CPU requirements for protocol operations and avoiding potential internal performance limits.

- A. PIM Dense Mode (PIM DM)
- B. PIM Sparse Mode (PIM SM)
- C. Distance Vector Multicast Routing Protocol (DVMRP)
- D. Multicast Open Shortest Path First (MOSPF)
- E. Bi-directional PIM

Answer: E

Explanation:

Bidirectional-PIM is a variant of the Protocol Independent Multicast (PIM) suite of routing protocols for IP multicast. In PIM, packet traffic for a multicast group is routed according to the rules of the mode configured for that multicast group. The Cisco IOS implementation of PIM supports three modes for a multicast group:

1. Bidirectional mode
2. Dense mode
3. Sparse mode

A router can simultaneously support all three modes or any combination of them for different multicast groups.

In bidirectional mode, traffic is routed only along a bidirectional shared tree that is rooted at the rendezvous point (RP) for the group. In bidir-PIM, the IP address of the RP acts as the key to having all routers establish a loop-free spanning tree topology rooted in that IP address. This IP address need not be a router, but can be any unassigned IP address on a network that is reachable throughout the PIM domain. Using this technique is the preferred configuration for establishing a redundant RP configuration for bidir-PIM.

In PIM dense mode (PIM-DM), PIM-SM, and most other multicast routing protocols such as Distance Vector Multicast Routing Protocol (DVMRP) and Multicast Open Shortest Path First (MOSPF), protocol operations and maintenance of packet forwarding state depend on signaling the presence or expiration of traffic (where "signaling" refers to both the packet forwarding engine to routing protocol process within the routers and the packet exchange part of the routing protocol). Triggering PIM assert messages, PIM register messages, and source tree forwarding state are all examples of traffic signaling. There are several advantages to traffic signaling, but they can lead to problems for applications with a large number of sources. For example, the more sources an application has, the less frequently traffic is sent from each sender. Each time a source starts to send packets, protocol operations take place and forwarding state is established. For applications with a large number of sources, this state can time out before the source sends again, resulting in "bursty sources." Therefore, applications with a large number of sources would not only create a large amount of forwarding state (requiring memory), but they also could require high CPU usage on the Route Processor due to the accounting of frequently changing state. In addition, the signaling within the router between the Route Processor and forwarding hardware can become a bottleneck if continuously large amounts of traffic signaling must go to the Route Processor and equally large amounts of forwarding state changes must go to the forwarding engines.

Bidir-PIM solves all these problems. Not only does bidir-PIM avoid maintaining source-specific forwarding state, therefore reducing the amount of memory needed by the number of sources per multicast group, but it also does not require any traffic signaling in the protocol. Thus, bidir-PIM prevents the "bursty source" problem, saving on CPU requirements for protocol operations and avoiding potential router internal performance limits.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1612/products_feature_guide09186a0080080a41.html

QUESTION 310

Source Specific Multicast (SSM) is being used throughout the Certkiller IP Multicast network. Which of the following three statements are correct with regard to SSM?

(Choose three)

- A. There are no RPs to worry about
- B. SSM uses shared Trees only
- C. SSM uses Shortest-Path Trees Only
- D. Is best suited for applications that are of the one-to-many category
- E. The user of SSM is recommended when there are many sources and it is desirable to keep in the amount of mroute state in the routers in the network to a minimum
- F. Is best suited for applications that are of the Many-to-Many category

Answer: A, C, D

Explanation:

A: SSM is easy to install and provision in a network because it does not require the network to maintain information about which active sources are sending to multicast groups. This requirement exists in ISM (with IGMPv1, IGMPv2, or IGMPv3). The current standard solutions for ISM service are PIM-SM and Multicast Source Discovery Protocol (MSDP). Rendezvous point (RP) management in PIM-SM (including the necessity for Auto-RP or bootstrap router [BSR]) and MSDP is required only for the network to learn about active sources. This management is not necessary in SSM. SSM is therefore easier than ISM to install and manage and easier to operationally scale in deployment. Another factor that contributes to the ease of installation of SSM is the fact that it can leverage preexisting PIM-SM networks.

C: SSM is a solution where the knowledge of the source is acquired out of band. SSM uses only a source tree, but there is no flooding of data, because learning the source is out of band. SSM is most useful for applications such as Internet broadcasting or corporate communications

D: SSM is a datagram delivery model that best supports one-to-many applications, also known as broadcast applications. SSM is a core networking technology for the Cisco implementation of IP multicast solutions targeted for audio and video broadcast application environments and is described in RFC 3569.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps5207/products_feature_guide09186a00801a6d6f.html

QUESTION 311

What best describes the Source Specific Multicast (SSM) functionality?

- A. SSM is an extension of the DVMRP protocol that allows for an efficient data delivery mechanism in one-to-many communications.
- B. SSM requires MSDP to discover the active sources in other PIM domains.
- C. In SSM routing of multicast traffic is entirely accomplished with source trees. The RP is used to direct receivers to the appropriate source tree.
- D. Using SSM, the receiver application can signal its intention to join a particular source by using the INCLUDE mode in IGMPv3.
- E. None of the above

Answer: D

Explanation:

The Internet Standard Multicast (ISM) service is described in RFC 1112, Host Extensions for IP Multicasting. This service consists of the delivery of IP datagrams from any source to a group of receivers called the multicast host group. The datagram traffic for the multicast host group consists of datagrams with an arbitrary IP unicast source address S and the multicast group address G as the IP destination address. Systems will receive this traffic by becoming members of the host group. Membership to a host group simply requires signalling the host group through IGMP Version 1, 2, or 3.

In SSM, delivery of datagrams is based on (S, G) channels. Traffic for one (S, G) channel consists of datagrams with an IP unicast source address S and the multicast group address G as the IP destination address. Systems will receive this traffic by becoming members of the (S, G) channel. In both SSM and ISM, no signalling is required to become a source. However, in SSM, receivers must subscribe or unsubscribe to (S, G) channels to receive or not receive traffic from specific sources. In other words, receivers can receive traffic only from (S, G) channels to which they are subscribed, whereas in ISM, receivers need not know the IP addresses of sources from which they receive their traffic. The proposed standard approach for channel subscription signalling utilizes IGMP INCLUDE mode membership reports, which are supported only in IGMPv3.

Incorrect Answers:

A. SSM is associated with PIM in IPv6 multicast networks. It is not associated with DVMP.

B. SSM builds off of PIM-SM, but also requires an update to IGMP. IGMP version 3 includes a larger header, where the source address can be specified, in addition to the group address. This means that a router no longer needs to communicate with an RP in order to locate the source, and also means that MSDP is no longer needed since its only purpose is to pass information among RPs.

C. PIM-SSM is made possible by IGMPv3. Because hosts can now indicate interest in specific sources using IGMPv3, PIM can create state directly along the path to those sources using SSM. SSM does not require a rendezvous point (RP) to operate.

QUESTION 312

The Certkiller network is using IP multicast within to conserve bandwidth during the training video seminars. In this IP multicast network, which of the following correctly describes scoping?

A. Scoping is the restriction of multicast data transport to certain limited regions of the network. There are two types: TTL scoping and administrative scoping.

B. Scoping is used by SSM to locate the sources and receivers in certain limited regions of the network. There are two types: TTL scoping and administrative scoping.

C. Scoping is a process used in MSDP to locate the sources and receivers in different AS.

D. PIM dense mode uses scoping to locate the sources and receivers in order to built shared trees.

E. None of the above

Answer: B

Explanation:

When applications operate in the global Multicast backbone (MBone), it is clear that not all groups should have global scope. Not only is this constraint especially important for performance reasons with flood and prune multicast routing protocols, but it also is true with other routing protocols for application security reasons and because multicast addresses are a scarce resource. Being able to constrain the scope of a session allows the same multicast address to be in use at more than one place as long as the scopes of the sessions do not overlap. This is analogous to the same radio frequency being used by two radio stations operating far apart from one another-each will only be heard locally. Multicast scoping can currently be performed in two ways, known as TTL Scoping and Administrative Scoping . Currently TTL scoping is most widely used, with only a very few sites making use of administrative scoping.

Reference:

http://www.cisco.com/web/about/ac123/ac147/ac174/ac198/about_cisco_ipj_archive_article09186a00800c851e

QUESTION 313

While troubleshooting an IP multicast issue, you issue the "show ip mroute" command:

```
Router#show ip mroute 236.2.3.23
```

IP Multicast Routing table

Flags: D - Dense, S - Sparse, C - Connected, L - Local, P - Pruned

R - RP-bit set, F - Register flag, T - SPT-bit set, J - JOIN SPT

X - Proxy Join Timer Running

Timers: uptime/Expires

Interface state: Interface, next-hop or VCD, State/Mode

(*, 236.2.3.23), 00:09:49/00:04:23 RP 10.1.24.1, flags: SC

Incoming interface: Serial1.708, RPF nbr 10.1.20.2

Outgoing interface list:

Ethernet0, Forward/Sparse, 00:09:50/00:04:12

You are trying to trace this multicast address back to the source of this multicast shared tree. Based on the information above, what is the IP address of the upstream neighbor?

- A. 10.1.24.1
- B. 10.1.24.2
- C. 10.1.20.2
- D. 10.1.20.3
- E. 236.2.3.23

Answer: C

Explanation:

The upstream neighbor is the IP address associated with the Reverse Path Forwarding

Neighbor (RPF nbr), which is 10.1.20.2 in this case.

Incorrect Answers:

A. 10.1.24.1 is the IP address of the Rendezvous Point in this example, not the upstream neighbor.

E. 236.2.3.23 is the IP address of the IP multicast session.

QUESTION 314

The Certkiller network is using multicasting for corporate video training sessions. All routers in the Certkiller network are enabled for IP multicast. How are these video streaming multicast packets forwarded by these routers? (Choose all that apply)

A. When a multicast packet arrives at a router, the router performs a Reverse Path forwarding (RPF) check on the packet. If the RPF check succeeds, the packet is forwarded, otherwise, it is dropped.

B. When traffic is flowing down the source tree the router looks up the source address in the unicast routing table to determine if the packet has arrived on the interface that is on the reverse path back to the source, if the packet has arrived on the interface leading back to source, the RPF check succeeds and the packets is forward. Otherwise, it is dropped.

C. When traffic is flowing down the source tree the router looks up the source address in the multicast routing table to determine if the packet has arrived on the interface that is on the reverse path back to the source. If the packet has arrived on the interface leading back to the source, the RPF check successfully the packets is forwarded. Otherwise, it is dropped.

D. When traffic is flowing down the source tree the router looks up the source address in the multicast routing table to determine if the packet has arrived on the interface that is on the reverse path back to the source and forward path to the receiver. If the reverse path and forward path is found successfully the packet is forwarded. Otherwise, it is dropped.

E. When a multicast packet arrives at a router, the router does not have to perform an RPF check on the packet. The router looks up the source address in the unicast routing table to determine if the destination path is present. If this succeeds the packet is forwarded. Otherwise, it is dropped.

Answer: A, B

Explanation:

Multicast Forwarding

In unicast routing, traffic is routed through the network along a single path from the source to the destination host. A unicast router does not consider the source address; it considers only the destination address and how to forward the traffic toward that destination. The router scans through its routing table for the destination address and then forwards a single copy of the unicast packet out the correct interface in the direction of the destination.

In multicast forwarding, the source is sending traffic to an arbitrary group of hosts that are represented by a multicast group address. The multicast router must determine which direction is the upstream direction (toward the source) and which one is the downstream

direction (or directions). If there are multiple downstream paths, the router replicates the packet and forwards it down the appropriate downstream paths (best unicast route metric)-which is not necessarily all paths. Forwarding multicast traffic away from the source, rather than to the receiver, is called Reverse Path Forwarding (RPF). RPF is described in the following section.

Reverse Path Forwarding (RPF)

PIM uses the unicast routing information to create a distribution tree along the reverse path from the receivers towards the source. The multicast routers then forward packets along the distribution tree from the source to the receivers. RPF is a key concept in multicast forwarding. It enables routers to correctly forward multicast traffic down the distribution tree. RPF makes use of the existing unicast routing table to determine the upstream and downstream neighbors. A router will forward a multicast packet only if it is received on the upstream interface. This RPF check helps to guarantee that the distribution tree will be loop-free.

RPF Check

When a multicast packet arrives at a router, the router performs an RPF check on the packet. If the RPF check succeeds, the packet is forwarded. Otherwise, it is dropped.

For traffic flowing down a source tree, the RPF check procedure works as follows:

1. The router looks up the source address in the unicast routing table to determine if the packet has arrived on the interface that is on the reverse path back to the source.
2. If the packet has arrived on the interface leading back to the source, the RPF check succeeds and the packet is forwarded.
3. If the RPF check in Step 2 fails, the packet is dropped.

Incorrect Answers:

C, D. The RPF lookup is done on the unicast routing table, not the multicast routing table.

E. RPF checks must be done in order to maintain a loop free multicast topology.

QUESTION 315

Router CK1 is the Designated Forwarder (DF) in the Certkiller IP multicast network. Which of the following statements about this Designated Forwarder in a Bidirectional PIM network is true?

- A. It has the best route to the Rendezvous Point and is the only router on the local subnet that may forward multicast traffic down the shared tree
- B. It is responsible for forwarding all multicast traffic onto and off of the local subnet
- C. It is elected based on the highest IP Address of all PIM routers on the local subnet and is the only router on the local subnet that may forward multicast traffic up the shared tree
- D. It has the best route to the Rendezvous Point and is the only router on the local subnet that may forward multicast traffic up the shared Tree
- E. None of the above

Answer: D

Explanation:

To avoid multicast packet looping, bidir-PIM introduces a new mechanism called the

designated forwarder (DF) election, which establishes a loop-free SPT rooted at the RP.

DF Election

On every network segment and point-to-point link, all PIM routers participate in a procedure called DF election. The procedure selects one router as the DF for every RP of bidirectional groups. This router is responsible for forwarding multicast packets received on that network upstream to the RP.

The DF election is based on unicast routing metrics and uses the same tie-break rules employed by PIM assert processes. The router with the most preferred unicast routing metric to the RP becomes the DF. Use of this method ensures that only one copy of every packet will be sent to the RP, even if there are parallel equal cost paths to the RP.

A DF is selected for every RP of bidirectional groups. As a result, multiple routers may be elected as DF on any network segment, one for each RP. In addition, any particular router may be elected as DF on more than one interface.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1834/products_feature_guide09186a00800800d6.html

QUESTION 316

A (*, G) multicast entry is associated with which multicast type of feature within the Certkiller IP multicast network? (Choose three)

- A. Shared tree
- B. PIM bidirectional
- C. Dense Mode
- D. Source tree
- E. Sparse Mode
- F. Dense Tree

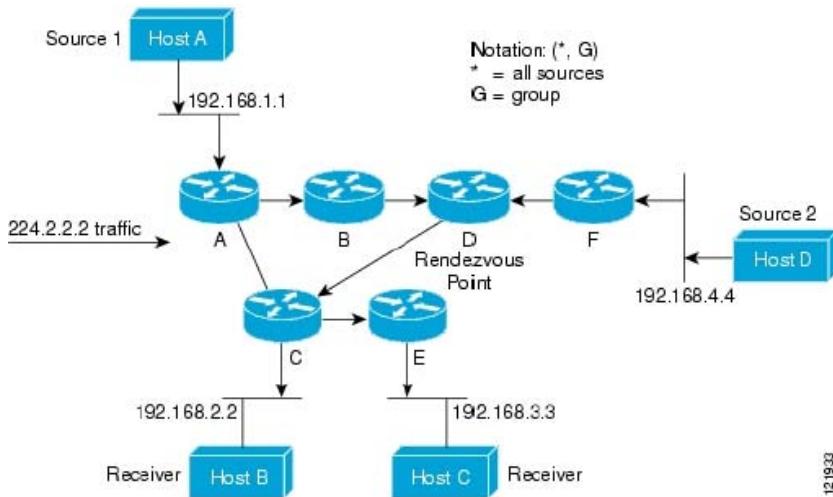
Answer: A, B, E

Explanation:

A: Unlike source trees that have their root at the source, shared trees use a single common root placed at some chosen point in the network. This shared root is called a rendezvous point (RP).

Figure5 shows a shared tree for the group 224.2.2.2 with the root located at Router D. This shared tree is unidirectional. Source traffic is sent towards the RP on a source tree. The traffic is then forwarded down the shared tree from the RP to reach all of the receivers (unless the receiver is located between the source and the RP, in which case it will be serviced directly).

Figure5 Shared Distribution Tree




In this example, multicast traffic from the sources, Hosts A and D, travels to the root (Router D) and then down the shared tree to the two receivers, Hosts B and C. Because all sources in the multicast group use a common shared tree, a wildcard notation written as (*, G), pronounced "star comma G," represents the tree. In this case, * means all sources, and G represents the multicast group. Therefore, the shared tree shown in Figure 5 would be written as (*, 224.2.2.2).

B, E: In bidirectional mode, traffic is routed only along a bidirectional shared tree that is rooted at the RP for the group. In bidir-PIM, the IP address of the RP acts as the key to having all routers establish a loop-free spanning tree topology rooted in that IP address. This IP address need not be a router address, but can be any unassigned IP address on a network that is reachable throughout the PIM domain.

Bidir-PIM is derived from the mechanisms of PIM sparse mode (PIM-SM) and shares many of the shared tree operations. Bidir-PIM also has unconditional forwarding of source traffic toward the RP upstream on the shared tree, but no registering process for sources as in PIM-SM. These modifications are necessary and sufficient to allow forwarding of traffic in all routers solely based on the (*, G) multicast routing entries. This feature eliminates any source-specific state and allows scaling capability to an arbitrary number of sources.

QUESTION 317

Part of the Certkiller IP multicast network is shown below:



Router CK2 sends a (S, G) Prune message to the LAN segment. Will this cause router CK1 to stop the multicast flow to router CK3 ?

- A. No. After seeing the Prune message from router CK2 , Router CK3 will send a Join message to router CK1 to override the Prune.
- B. No. After seeing the Prune message from router CK2 , Router CK3 will send a Join message to router CK2 to override the Prune.
- C. No. After seeing the Prune message from router CK2 , Router CK3 will send a Graft message to router CK2 to override the Prune from router CK2 .
- D. Yes. Router CK3 will need to send a new Join message to re-join the multicast session.
- E. It depends on whether the routers are IGMP version 1 or IGMP version 2.

Answer: A

Explanation:

After a prune, the router waits for joins, if none arrive, then the router drops the Group. In this case, router CK1 will hear the Join message from CK3 to prevent the flow of multicast traffic from being cut off to CK3 .

Incorrect Answers:

- B. Router CK2 will send a Join message to the upstream neighbor, which is CK1 in this case, not CK2 .
- C. No graft messages will be sent in this case.
- E. IGMP versions are irrelevant.

QUESTION 318

What is the primary purpose for the RPF check in IP multicast networks?

- A. To establish reverse flow path of multicast traffic from the receiver to the source.
- B. To prevent multicast traffic looping through the network.
- C. To determine interfaces inclusion in the outgoing interface list.
- D. To prevent the movement of unauthorized multicast traffic.

Answer: B

Explanation:

Reverse Path Forwarding (RPF) provides loop avoidance. It is an algorithm used to forward multicast packets. The RPF rules are: If a router receives a datagram on an interface that it uses to send unicast packets to the source of that packet, then the packet has arrived on the RPF interface. If the packet arrives on the RPF interface, a router forwards the packet out the interfaces that are present in the outgoing interface list of a multicast routing table entry. If the packet does not arrive on the RPF interface, the packet is silently discarded.

QUESTION 319

Which Multicast Protocols use Reverse Path Forwarding (RPF) information when sending multicast traffic streams to the receivers within the Certkiller network?
(Select two)

- A. DVMRP
- B. PIM Sparse Mode
- C. PIM Dense Mode
- D. Multicast OSPF
- E. PIM Sparse-Dense Mode

Answer: A, C

Explanation:

DVMRP uses a technique known as Reverse Path Forwarding. When a router receives a packet, it floods the packet out of all paths except the one that leads back to the packet's source. Doing so allows a data stream to reach all LANs (possibly multiple times). If a router is attached to a set of LANs that do not want to receive a particular multicast group, the router can send a "prune" message back up the distribution tree to stop subsequent packets from traveling where there are no members.

Dense-mode PIM uses Reverse Path Forwarding and looks a lot like DVMRP. The most significant difference between DVMRP and dense-mode PIM is that PIM works with whatever unicast protocol is being used; PIM does not require any particular unicast protocol.

Incorrect Answers:

B. Sparse-mode PIM is optimized for environments where there are many multipoint data streams. Each data stream goes to a relatively small number of the LANs in the internetwork. For these types of groups, Reverse Path Forwarding techniques waste bandwidth. Sparse-mode PIM works by defining a Rendezvous Point. When a sender wants to send data, it first sends to the Rendezvous Point.

D. Multicast OSPF (MOSPF) was defined as an extension to the OSPF unicast routing protocol. OSPF works by having each router in a network understand all of the available links in the network. Each OSPF router calculates routes from itself to all possible destinations.

MOSPF works by including multicast information in OSPF link state advertisements. An MOSPF router learns which multicast groups are active on which LANs.

MOSPF builds a distribution tree for each source/group pair and computes a tree for active sources sending to the group. The tree state is cached, and trees must be recomputed when a link state change occurs or when the cache times out.

QUESTION 320

What single choice listed below best describes PIM functionality?

- A. PIM uses the multicast routing information to perform the multicast forwarding function. PIM is a multicast routing protocol, and uses the multicast routing table to perform the RPF check. Like other routing protocols, PIM sends and receives routing updates between routers.
- B. PIM uses unicast routing protocol information that populates the unicast routing table, including EIGRP, OSPF, BGP, and static routes.
- C. PIM uses the multicast and unicast routing information to perform the multicast forwarding function. PIM uses the multicast routing table to perform the RPF check. Like other routing protocols, PIM does not send and receive routing updates between routers.
- D. PIM uses multicast routing protocols to populate the multicast routing table, including Distance Vector Multicast Routing Protocol (DVMRP); Multicast OSPF (MOSPF), Multicast BGP
- E. PIM uses the unicast routing information to perform the multicast forwarding function. Although PIM is called a multicast routing protocol, it actually uses the unicast routing table to perform the RPF check function instead of building up a completely independent multicast routing table. Unlike other routing protocols, PIM does not send and receive routing updates between routers.

Answer: E

Explanation:

Protocol-independent multicast (PIM) gets its name from the fact that it is IP routing protocol-independent. PIM can leverage whichever unicast routing protocols are used to populate the unicast routing table, including EIGRP, OSPF, BGP, or static routes. PIM uses this unicast routing information to perform the multicast forwarding function, so it is IP protocol-independent. Although PIM is called a multicast routing protocol, it actually uses the unicast routing table to perform the reverse path forwarding (RPF) check function instead of building up a completely independent multicast routing table. PIM does not send and receive multicast routing updates between routers like other routing protocols do.

QUESTION 321

A new Certkiller router is configured to participate in the classic PIM-SM IP multicast network. When running classic PIM-SM (aka PIM-ASM or PIM any-source Multicast) on this router, which one of the following statements about PIM Joins is correct?

- A. PIM Joins are sent every 60 seconds to refresh the upstream router's mroute state for

the multicast tree

- B. Routers send a PIM Join-Ack in response to each PIM Join received from a downstream router
- C. PIM Joins are only sent when the multicast distribution tree is first being established
- D. PIM Joins are sent every 3 minutes to refresh the upstream router's mroute state for the multicast tree
- E. None of the above

Answer: A


Explanation:

In a typical PIM-SM implementation, each neighboring router that supports a downstream multicast group member sends a PIM Join/Prune Message every 60 seconds, so it can take up to 60 second for a rebooted router to learn the forwarding states for any downstream multicast group members after learning the identity of the multicast group RP router.

Reference: <http://www.patentstorm.us/patents/6631420-description.html>

QUESTION 322

The Certkiller network is shown in the following exhibit:



While troubleshooting a problem with the IP multicast network, you see the following on router CK1 :

(* , 224.1.1.1), 00:00:03/00:00:00, RP 171.68.28.140, flags: SP

Incoming interface: Serial0, RPF nbr 171.68.28.191,

Outgoing interfaces list: Null

(171.68.37.121/32, 224.1.1.1), 00:00:03/00:02:56, flags FPT

Incoming interface: Ethernet0, RPF nbr 0.0.0.0, Registering

Outgoing interface list: Null

Which of the following could be the cause of the "Registering" condition on CK1 ?

(Choose all that apply)

- A. Router CK1 has incorrectly calculated the RPF interface for the source (171.68.37.121) as Serial1.
- B. Router CK3 (RP) failed to send a "Register-Stop" message to Router CK2 .
- C. Router CK2 is IGMP version 1 while Router CK1 is an IGMP version 2 speaker.
- D. PIM is not enabled on Router CK2 .
- E. Registering is the normal operational status of an operational multicast session.

Answer: B, D

Explanation:

The Rendezvous Point will need to send a "Register Stop" in order to clear the registration process, and all routers in between the RP and the multicast source must be multicast enabled.

Incorrect Answers:

A. The output shows that this is not the problem, as router CK1 is correctly calculating the incoming interface as Ethernet 0.

C. IGMP is used by hosts, and IGMP version 2 is backwards compatible with version 1.

Reference:

Developing IP Multicast Networks, (from page 259, PIM register process).

QUESTION 323

You are a technician at Certkiller . Your newly appointed Certkiller trainee wants to know which IP protocol is used to send PIMv2 control messages.

What would your reply be?

- A. UDP
- B. TCP
- C. BGP
- D. Protocol number 107
- E. Protocol number 103

Answer: E

Explanation:

All PIM control messages have protocol number 103.

Reference:

<http://www.ietf.org/proceedings/99mar/I-D/draft-ietf-pim-v2-dm-01.txt>

QUESTION 324

Which of the following PIMv2 Sparse mode control messages are also used in PIM Dense mode? (Choose all that apply.)

- A. Graft
- B. Join
- C. Prune
- D. Register
- E. Assert
- F. Hello
- G. Register

Answer: B, C, E, and F

Explanation:

PIM-DM uses the following PIMV2 messages.

- Hello

- Join/Prune
- Graft
- Graft-Ack
- Assert

PIM-SM uses the following PIMV2 messages

- Hello
- Bootstrap
- Candidate-RP-Advertisement
- Join/Prune
- Assert
- Register
- Register-Stop

Reference:

'CCIE Professional Development Routing TCP/IP Volume 2' in the section
'Understanding IP Multicast Routing' pages 475 and 488.

QUESTION 325

In the Certkiller IP multicast network, what best describes the MDT role in MVPN operations?

- A. PE routers that have CE routers who are intended recipients of the data only join data MDT. PE routers signal use of data-MDT via a UDP packet on port 3232, which is sent via the default MDT. This packet contains an all-PIM routers message, indicating the group is joined if required.
- B. CE routers do not have a PIM adjacency across the provider network with remote CE routers, but rather have an adjacency with their local routes on the PE router. When the PE router receives an MDT packet. It performs an RPF check. During the transmission of the packet through the Provider network, the normal RPF rules apply. However, at the remote's PE, the router needs to ensure that the originating PE router was the correct one for that CE. It does this by checking the BGP next hop address of the customer's packet's source address. This next hop address should be the source address of the MDT packet. The PE also checks that there is a PIM neighbor relationship with the remote PE.
- C. A unique Group address is required to be used as MDT for each particular customer. A unique source address for the Multicast packet in the provider network is also required. This source address is recommended to be the address of the loopback interface, which is used as the source for the IBGP, as this address is used for the RPF check at remote PE.
- D. PE routers are the only routers that need to be MVPN aware and able to signal to remote PE's information regarding the MVPN. It is therefore fundamental that all PE routers have a BGP relationship with each other. Either directly or via a Route Reflector. The source address of the Default-MDT will be the same address used to source the IBGP sessions with the remote PE routers that belong to the same VPN and MVRF.
- E. All of the above.

Answer: E

Explanation:

Cisco MVPN Details:

While there are significant deployment obstacles to each of the preceding MVPN solutions, Multicast Domains is the most attractive alternative because:

1. The provider must configure a native IP multicast network within their core network; this includes both the P and PE routers.
2. IP Multicast is a mature technology that has been deployed since Cisco IOS Software 10.0. This minimizes risk for the provider network, because a new feature will not have to be introduced into its core to support MVPNs.


Multicast Domain Solution

This method originally had less than optimal performance, because it requires that all PE routers connected to a customer receive all of that customer's Multicast data regardless of the presence of an interested receiver in that location. When enhancements resolved this characteristic with a new methodology, it became a truly attractive solution.

Default MDT Concept

These trees distribute multicast group data that exceeds a certain configured threshold of Bandwidth (BW) to only those PE who have joined this new tree. These are trees are called MDT-data trees. The word data is appended as these groups are designed to be used for groups that will require a higher amount of bandwidth to deliver their data.

Data MDT Concept



This diagram indicates that the Data MDT is only joined by those PE routers that have CE routers who are intended recipients of the data.


PE routers signal use of Data-MDT via a UDP packet on port number 3232, which is sent via the default MDT. This packet contains an all-PIM routers message, indicating the group to be joined if required.

Interaction of Customer and Providers Multicast Network

It is important to remember that the customer's IP Multicast network has no relationship to the provider's multicast network. From the perspective of the provider, the customer's IP Multicast packets are merely data to the provider's distinctive IP Multicast network.

It is important to understand that PIM, and in particular PIM-SM, are the only supported multicast protocols for MVPN. Bi-Dir PIM may be supported in the future, when it is deemed stable enough for the core of a provider network.

Customer PIM Adjacencies



CE routers do not have a PIM adjacency across the provider network with remote CE routers, but rather have an adjacency with their local routers and the PE router. When the PE router receives an MDT packet, it performs an RPF check. During the transmission of the packet through the Provider network, the normal RPF rules apply. However, at the remote's PE, the router needs to ensure that the originating PE router was the correct one for that CE. It does this by checking the BGP next hop address of the customer's packet's source address. This next hop address should be the source address of the MDT packet. The PE also checks that there is a PIM neighbourhood with the remote PE.

Currently, only a single MVRF is supported per customer. This limitation precludes the customer also receiving Internet or any other outside domain's Multicast traffic. A unique Group address is required to be used as MDT for each particular customer. A unique source address for the Multicast packet in the provider network is also required. This source address is recommended to be the address of the loopback interface, which is used as the source for the iBGP, as this address is used for the RPF check at remote PE. If the provider uses MDT-data groups, then these will also need to be configured. These MDT-data groups must be unique for each customer.

The PE routers must have a PIM adjacency to each other. No other routing protocols may use these MTIs.

Provider's PIM Adjacencies



BGP Requirements

PE routers are the only routers that need to be MVPN aware and able to signal to remote PE's information regarding the MVPN. It is therefore fundamental that all PE routers have a BGP relationship with each other. Either directly or via a Route Reflector.

The source address of the Default-MDT will be the same address used to source the iBGP sessions with the remote PE routers that belong to the same VPN and MVRF.

When PIM-SSM is used for transport inside the provider core, it is via this BGP

relationship that the PEs indicate that they are MVPN capable and provide for source discovery. This capability is indicated via the updated BGP message.

Reference:

http://www.cisco.com/en/US/tech/CK8_28/technologies_white_paper09186a00800a3db6.shtml

QUESTION 326

The Certkiller IP multicast network utilizes classic PIM-SM. Which two statements are correct with regard to this classic PIM-SM? (Choose two)

- A. The Cisco IOS default is for last-hop routers to trigger a switch to the shortest-path tree as soon as new source is detected on the shared tree
- B. The default behavior of switching to the Shortest-path tree as soon as a new source is detected on the shared Tree can be disabled by setting the "ip pim spt-threshold" to infinity
- C. The IOS default is for ANY routers on the shared Tree to trigger a switch to the Shortest-Path Tree as soon as a new source is detected on the shared Tree
- D. The default behavior of switching to the shortest-path Tree as soon as a new source is detected on the shared Tree can be disabled by setting the 'ip pim spt-threshold' to "zero"

Answer: A, B

Explanation:

PIM-SM Shortest-Path Tree Switchover:

- PIM-SM has the capability for last-hop routers (i.e. routers with directly connected members) to switch to the Shortest-Path Tree and bypass the RP if the traffic rate is above a set threshold called the "SPT-Threshold".

The default value of the SPT-Threshold" in Cisco routers is zero. This means that the default behaviour for PIM-SM leaf routers attached to active receivers is to immediately join the SPT to the source as soon as the first packet arrives via the (*,G) shared tree.

With the "ip pim spt-threshold" command, specifying the infinity keyword causes the router never to move to the shortest-path tree; it remains on the shared tree. This keyword applies to a multicast environment of "many-to-many" communication.

Reference: <ftp://ftp-eng.cisco.com/ipmulticast/training/Module1.pdf> page 82

QUESTION 327

Certkiller is suing Sparse Mode PIM within their IP multicast network. In PIM-SM, what control plane signaling must a multicast source perform before it begins to send multicast traffic to a group?

- A. The source must perform a "Request to send " (RTS) and "Clear to send " (CTS) handshake with the PIM Designated Router (DR) on the local subnet
- B. The source sends a PIM register message to the Rendezvous Point (RP)
- C. The source must first join the multicast group using IGMP before sending
- D. None, the source just begins to send on the local subnet
- E. None of the above

Answer: D

Explanation:

PIM-SM was designed to support the following goals:

Maintain the traditional IP multicast service model of receiver-initiated multicast group membership. In this model, sources simply put packets on the first-hop Ethernet, without any signaling. Receivers signal to routers in order to join the multicast group that will receive the data.

Leave the host model unchanged. PIM-SM is a router-to-router protocol, which means that the hosts don't have to be upgraded, but that PIM-SM-enabled routers must be deployed in the network.

Reference:

<http://www.microsoft.com/technet/prodtechnol/windows2000serv/maintain/featusability/pimsm2.msp#EADAC>

QUESTION 328

Which of the following is used to calculate the upstream neighbor interface for a multicast route entry in a PIMv2 Sparse Mode network?

- A. The address of the Mapping Agent.
- B. The address of a directly connected member of the multicast group.
- C. The address of the currently active Rendezvous Point for the multicast group.
- D. The address of the PIM neighbor that sent the PIM Join message.
- E. The address of the PIM neighbor that sent the PIM Hello message.

Answer: C

Explanation:

The address of the upstream neighbor in any PIMv2 Sparse Mode network is always calculated via the neighbor closest to the Rendezvous Point (RP).

Incorrect Answers:

- A. The upstream neighbor for a multicast group is calculated from the RP, not the mapping agent.
- B. The directly connected multicast neighbor would only be used if it were the nearest upstream neighbor toward the RP, which will not always be the case.
- D, E. The neighbor that sends the PIM messages is not necessarily going to be the same neighbor that is upstream toward the RP, so these choices are also incorrect.

Reference:

CCIE Professional Development Routing TCP/IP Volume II by Jeff Doyle and Jennifer De Haven Carroll, Page 492.

QUESTION 329

The Certkiller network administrator has just issued the "ip pim autorp listener" command on router CK1 . This command is used to:

- A. Allow AutoRP Packets in the groups 224.0.1.39 and 224.0.1.40 to Dense Mode flooded out interfaces configured as 'ip pim sparse-mode'
- B. Configures the router as an AutoRP Mapping Agent
- C. Enable a Cisco Router to "Passively" listen to AutoRP packets without the router actively sending/forwarding any AutoRP packets
- D. Enables the use of AutoRP on the router
- E. None of the above.

Answer: A

Explanation:

To cause IP multicast traffic for the two Auto-RP groups 224.0.1.39 and 224.0.1.40 to be Protocol Independent Multicast (PIM) dense mode flooded across interfaces operating in PIM sparse mode, use the `ip pim autorp listener` command in global configuration mode. Use the `ip pim autorp listener` command with interfaces configured for PIM sparse mode operation in order to establish a network configuration where Auto-RP operates in PIM dense mode and multicast traffic can operate in sparse mode, bidirectional mode, or source specific multicast (SSM) mode.

Example:

The following example enables IP multicast routing and the Auto-RP listener feature on a router. It also configures the router as a Candidate RP for the multicast groups 239.254.2.0 through 239.254.2.255.

```
ip multicast-routing
ip pim autorp listener
ip pim send-rp-announce Loopback0 scope 16 group-list 1
ip pim send-rp-discovery Loopback1 scope 16
access-list 1 permit 239.254.2.0 0.0.0.255
```

Reference:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fiprmc_r/mult/1rfmult2.htm#wp1090

QUESTION 330

Routers on the Certkiller IP multicast network are configured to forward Auto RP messages. These messages are forwarded via:

- A. IP Multicast Using IANA registered groups 224.0.1.39 and 224.0.1.40
- B. Hop-by-Hop flooding of Auto-RP control packets via the 224.0.0.2 [ALL-Routers] link local multicast group
- C. Unicast Messages between candidate RPs and the Mapping Agent
- D. Hop-by-Hop flooding of Auto-RP control packets via the 224.0.0.13 [PIM-ROUTERS] local link multicast group
- E. None of the above

Answer: A

Explanation:

Auto-RP automates the distribution of group-to-RP mappings in a network supporting sparse mode PIM. It supports the use of multiple RPs within a network to serve different group ranges, and allows configurations of redundant RPs for reliability purposes. In comparison, manual configuration of RP information is prone to inconsistency which can cause connectivity problems. In summary, the automatic distribution of Group-to-RP mappings simplifies configuration task, and guarantees consistency.

In a nutshell, the Auto-RP mechanism operates on two basic components, the candidate RPs and the RP mapping agents:

1. All candidate RPs advertise their willingness to be an RP via "RP-announcement" messages.

These messages are periodically sent to a reserved well-known group 224.0.1.39 (CISCO-RP-ANNOUNCE). The default interval is 60 seconds (tunable via CLI).

1. The RP mapping agents join group 224.0.1.39 and select consistently an RP for each group address range. The pair (group range --> RP) is called an RP-mapping.

The RP mapping agents advertise the authoritative RP-mappings to another well-known group address 224.0.1.40 (CISCO-RP-DISCOVERY). All PIM routers join 224.0.1.40 and store the RP-mappings in their private cache.

With Auto-RP, Multiple RPs can be used to serve different group ranges, or as hot-backups of each other.

Reference: <ftp://ftpeng.cisco.com/ipmulticast/autorp.html>

QUESTION 331

What interface command must be configured for auto-rp to function properly in the Certkiller IP multicast network?

- A. ip pim dense-mode
- B. ip pim sparse-dense-mode
- C. ip pim sparse-mode
- D. ip multicast helper
- E. None of the above

Answer: B

Explanation:

RPs are used by senders to a multicast group to announce their existence and by receivers of multicast packets to learn about new senders.

Auto-RP is a feature that automates the distribution of group-to-RP mappings in a PIM network. Auto-RP has the following benefits:

Configuring the use of multiple RPs within a network to serve different group ranges is easy.

Auto-RP allows load splitting among different RPs and arrangement of RPs according to the location of group participants.

Auto-RP avoids inconsistent, manual RP configurations that can cause connectivity problems.

Example configuration using Auto-RP:

ip multicast-routing

```
interface ethernet 0/0
ip pim sparse-dense-mode
ip pim send-rp-announce ethernet 0 scope 16 group-list 1
ip pim rp-address 10.8.0.20 1
```

Incorrect Answers:

- A. Rendezvous points are used in sparse mode multicasts, not dense mode.
- C. If router interfaces are configured in sparse mode only a static RP address must also be configured.
- D. This is an invalid command.

QUESTION 332

In order to configure anycast RP between two different Certkiller routers, which of the following minimum requirements must be satisfied?

- A. Multicast Source Discovery Protocol (MSDP), mesh-groups must be configured between in two Anycast RPs
- B. The RPs must be within the same IGP domain
- C. Multicast Source Discovery Protocol (MSDP) must be configured between the two Anycast RPs
- D. The two Anycast RPs must be iBGP peers
- E. None of the above

Answer: C

Explanation:

Anycast-RP is an extension of the Static RP technique that also allows multiple Rendezvous Points for a group range to be deployed. This allows the network to continue to operate if a Rendezvous Point fails. The idea is to configure two or more routers in the network to be the Rendezvous Point. Each of these Anycast-RP routers will be configured with the same Rendezvous Point address (in this case 10.1.1.1) on one of their Loopback interfaces. Each router also will advertise this address (the Rendezvous Point address) as a /32 host route. This will result in the other routers in the network using the closest Anycast-RP as their Rendezvous Point based on the unicast routing metrics. Normally, this would split the network into multiple PIM-SM domains that would not talk to each other. However, the Multicast Source Discovery Protocol (MSDP) is used to communicate active source information from one Anycast-RP to the other in Source Active (SA) messages. This allows active sources in one half of the network to be learned and joined by the Rendezvous Point in the other half of the network.

Reference:

http://www.cisco.com/en/US/tech/CK8_28/technologies_white_paper09186a00802d4643.shtml

QUESTION 333


Anycast Rendezvous Points (RPs) have been implemented in the Certkiller IP Multicast network. Which statement below is correct with regard to these Anycast RPs?

- A. After a failure of one of the Anycast RPs, the PIM network will reconverge on the remaining Anycast RP(s) in approximately the same time as it takes the unicast routing table to reconverge
- B. The Anycast RPs must reside within the same IGP domain
- C. Anycast RPs can't be used in conjunction with AutoRP
- D. After a failure of one of Anycast RPs, the PIM network will reconverge on the remaining Anycast RP(s) in less than 1 second
- E. None of the above

Answer: A

Explanation:

Anycast-RP:



If an Anycast-RP were to fail, its host route would cease being advertised to the network and the unicast routing will reconverge on the remaining Anycast-RP. This will cause the routers in the network to rejoin and reregister receivers and sources to the remaining Anycast-RP to maintain multicast traffic flow. This process occurs in approximately the time that it takes unicast routing to converge which means that Anycast-RP has one of the fastest Rendezvous Point failover times of all of the Rendezvous Point configuration methods.

Reference:

http://www.cisco.com/en/US/tech/CK828/technologies_white_paper09186a00802d4643.shtml

QUESTION 334

The Certkiller network is utilizing IP multicast technology. Along with this, router CK1 is configured as an anycast Rendezvous Point (RP). What best describes the functionality of Anycast RP?

- A. Anycast RP is a useful application of MSDP, MBGP and SSM that configures a multicast sparse mode network to provide for fault tolerance and load sharing within a single multicast domain.
- B. Only a maximum of two RPs are configured with the same IP address (for example, 10.0.0.10) on loopback interfaces. The loopback address should not be configured as a host address (with a 32-bit mask). All the downstream routers are configured so that they know that 10.0.0.10 is the IP address of their local RP.

C. IP routing automatically selects the topologically closest RP for each source and receiver. Because some sources use only one RP and some receivers a different RP, MBGP enables RPs to exchange information about active sources. All the RPs are configured to be MSDP peers of each other.

D. Each RP will know about the active sources in its own area. If RP fails, IP routing converges and backup RP would become the active RP of this area using HSRP.

E. Anycast RP is an implementation strategy that allows load sharing and redundancy in PIM sparse mode (PIM-SM) networks by configuring two or more RPs that have the same IP address and use Multicast Source Discovery Protocol to share active source information.

Answer: E

Explanation:

IP multicast is deployed as an integral component in mission-critical networked applications throughout the world. These applications must be robust, hardened, and scalable to deliver the reliability that users demand.

Using Anycast RP is an implementation strategy that provides load sharing and redundancy in Protocol Independent Multicast sparse mode (PIM-SM) networks.

Anycast RP allows two or more rendezvous points (RPs) to share the load for source registration and the ability to act as hot backup routers for each other. Multicast Source Discovery Protocol (MSDP) is the key protocol that makes Anycast RP possible.


The main purpose of an Anycast RP implementation is that the downstream multicast routers will "see" just one address for an RP.

Reference:

http://www.cisco.com/en/US/tech/CK828/technologies_white_paper09186a00800d6b60.shtml#57583

QUESTION 335

The Certkiller network consists of network X and Y that are connected via Router CK1 and Router CK2. The Certkiller network is shown in the following exhibit:



You wish to set up an IPSec VPN between routers CK1 and CK2. Now, which of the following crypto access-lists must be configured on Router CK1 in order to send LAN to LAN traffic across the encrypted VPN tunnel?

A. access-list 101 permit ip host 192.168.1.1 host 192.168.1.2

- B. access-list 101 permit ip 10.1.1.0.0.0.0.255 host 192.168.1.2
- C. access-list 101 permit ip 10.1.1.0.0.0.0.255 10.1.2.0.0.0.0.255
- D. access-list 101 permit ip 10.1.1.0.0.0.0.255 10.1.2.0.0.0.0.255
access-list 101 permit ip 10.1.2.0.0.0.0.255 10.1.1.0.0.0.0.255
- E. access-list 10 permit ip 10.1.1.0.0.0.0.255 10.1.2.0.0.0.0.255

Answer: C

Explanation:

The format of the command for configuring IPSec is shown below:

access-list 101 permit "Source Network Addresses on X" "Destination Network Subnets on Y"

Incorrect Answers:

- A. You define the traffic that is to be sent over the encrypted tunnel, which is all traffic from subnet X to subnet Y, not the serial interfaces.
- B. This would only be useful for traffic going from subnet X to the serial interface of CK2 , not for LAN to LAN traffic.
- D. You only need to specify the traffic from X to Y on router CK1 , as this is the traffic that will be encrypted. The second line of this access list would need to be applied to router CK2 only.
- E. Access list 100 or higher must be used, as this is an extended access list.

QUESTION 336

You try to perform a traceroute to an Internet destination from your PC, but the Traceroute hangs when it reaches the router. Currently, there is an inbound access-list applied to the serial interface on the Internet router with a single line: "access-list 101 permit tcp any any".

What access-list entry may you need to be added to the access-list in order to get traceroute to work?

- A. access-list 101 permit tcp any any
- B. access-list 101 permit icmp any any time-exceeded
access-list 101 permit icmp any any port-unreachable
- C. access-list 101 permit icmp any any time-exceeded
access-list 101 permit icmp any any echo-reply
- D. access-list 101 permit icmp any any echo
access-list 101 permit icmp any any net-unreachable
- E. access-list 101 permit udp any any
access-list 101 permit icmp any any protocol-unreachable

Answer: B

Explanation:

Port-unreachable and time-exceeded are the ICMP messages that Cisco traceroute uses, so these ports must be permitted to allow the traceroute to go through.

Incorrect Answers:

A, C, D, E. None of these options give us both the time-exceeded and port-unreachable ICMP ports that need to be opened in the access list to allow traceroute through.

QUESTION 337

You are writing an access list on a router to prevent users on the Ethernet LAN connected to Ethernet interface 0 from accessing a TFTP server (10.1.1.5) located on the LAN connected to Ethernet interface 1. Which of the following would be the correct configuration change if applying the ACL inbound on the Ethernet 0 interface?

- A. access-list 1 deny tcp 0.0.0.0 255.255.255.255 10.1.1.5 0.0.0.0 eq 69
- B. access-list 100 deny tcp 0.0.0.0 255.255.255.255 10.1.1.5 0.0.0.0 eq 69
- C. access-list 100 deny tcp 0.0.0.0 255.255.255.255 10.1.1.5 0.0.0.0 eq 68
- D. access-list 100 deny tcp 10.1.1.5 0.0.0.0 0.0.0.0 255.255.255.255 eq 69
- E. access-list 100 deny tcp 0.0.0.0 255.255.255.255 10.1.1.5 0.0.0.0 eq port 68
- F. None of the above

Answer: F

Explanation:

TFTP uses UDP port 69, so choice F would be the correct access list entry. An extended access list is needed when filtering based on source and destination address, as well as layer 4 port information. However, all of the choices listed are filtering based on TCP ports, and since TFTP uses UDP none are correct.

Incorrect Answers:

- A. This is an invalid command, since using source and destination information along with port numbers requires an extended access list.
 - B. This would be the correct choice if UDP was specified as the transport layer protocol instead of TCP.
 - C, E. In addition to incorrectly specifying TCP instead of UDP, the port number of 68 is also incorrect.
 - D. The order of the IP address arrangement is incorrect. This access list will block all TCP port 69 traffic sourced from the TFTP server, not destined to it. This choice is also incorrectly using TCP instead of UDP.
-

QUESTION 338

You wish to allow only telnet traffic to a server with an IP address 10.1.1.100. You add the following access list on the router:

```
access-list 101 permit tcp any host 10.1.1.100 eq telnet
```

```
access-list 101 deny ip any any
```

You then apply this access list to the inbound direction of the serial interface. Which types of packets will be permitted through the router after this change? (Choose all that apply)

- A. A non-fragment packet en route to the server on port 21.

- B. A non-initial fragment packet en route to the server on port 23.
- C. A non-initial fragment packet passing through to another host that's not 10.1.1.100.
- D. A non-initial fragment packet going to the server on port 21.
- E. An initial-fragment or non-fragment packet en route to the server on port 23.

Answer: B, D, E

Explanation:

B, E: Telnet (port 23) is permitted by ACL.

D: A non initial fragment destined to the server will indeed be permitted. The reason for this is that the first line of ACL has some L3 and some L4 information which needs to be matched for a packet to be permitted.

Since a non initial frame matches the L3 information it will pass the layer 3 check.

Moreover, since it is a non initial frame it will contain no L4 information in it. Hence the packet will be permitted.

Incorrect Answers:

A, C. For non-initial fragments, only telnet packets going to the 10.1.1.100 address will be allowed.

QUESTION 339

The following access list is configured on router CK1 :

```
access-list 100 deny udp 0.0.0.0 255.255.255.255 10.1.1.5 0.0.0.0 eq 69
```

What does the access-list accomplish?

Note: Assume that all other traffic is permitted with a permit all statement at the end of the access list.

- A. It blocks all incoming traffic arriving on E0 from accessing any FTP server.
- B. It blocks all incoming traffic, except traffic addresses to 10.1.1.5, from accessing any FTP servers.
- C. It blocks all incoming traffic arriving on E0 from accessing the FTP server with an address of 10.1.1.5.
- D. It blocks all incoming UDP traffic.
- E. This access list is trying to block traffic from accessing a TFTP server. However, this is only half of what is needed to accomplish that. You would also need the following:

```
access-list deny tcp 0.0.0.0 255.255.255.255 10.1.1.5 0.0.0.0 eq 69
```

Answer: E

Explanation:

The access list shown above is designed to block UDP port 69 traffic from all sources to the destination device with the IP address of 10.1.1.5. Port 69 is used for TFTP. Both TCP and UDP ports are used with the TFTP application, so in order to block all TFTP traffic another access list block TCP port 69 should also be applied.

Incorrect Answers:

A, B: TFTP traffic is being blocked, not FTP. In addition, this traffic is being blocked only for traffic destined to a single server, not all traffic.


C. TFTP uses port 69, not FTP. FTP uses ports 20 and 21. Since TFTP uses both TCP and UDP, both ports will need to be filtered.

D. Only UDP port 69 traffic destined to a single server is being filtered, not all UDP traffic.

Reference: <http://www.ibiblio.org/security/articles/ports.html>

QUESTION 340

Network Topology Exhibit:



In this Certkiller network segment, you want to block all Smurf attacks that originate on the 192.168.2.0 network from being sent into the 192.168.1.0 network. However, all other traffic must be permitted. No access lists currently exist on the router. Which of the following configuration excerpt would accomplish this task when applied to E0 on router Certkiller 0 as an input filter?

- A. access-list 1 permit 192.168.2.0 0.0.0.255
access-list 1 deny any
- B. access-list 1 deny 192.168.1.0 0.0.0.255
access-list 1 permit any
- C. access-list 100 permit ip any 192.168.1.0 0.0.0.255
access-list 100 deny ip any any
- D. access-list 100 deny icmp any 192.168.1.255 0.0.0.0 echo
access-list 100 permit icmp any 192.168.1.0 0.0.0.255 echo
access-list 100 permit ip any any
- E. access-list 100 deny icmp any 192.168.1.255 0.0.0.0 echo-reply
access-list 100 permit icmp any any echo-reply
access-list 100 permit ip any any
- F. None of the above

Answer: D

Explanation:

A SMURF attack (named after the program used to perform the attack) is a method by which an attacker can send a moderate amount of traffic and cause a virtual explosion of

traffic at the intended target. The method used is as follows:

1. The attacker sends ICMP Echo Request packets where the source IP address has been forged to be that of the target of the attack.
2. The attacker sends these ICMP datagrams to addresses of remote LANs broadcast addresses, using so-called directed broadcast addresses. These datagrams are thus broadcast out on the LANs by the connected router.
3. All the hosts which are "alive" on the LAN each pick up a copy of the ICMP Echo Request datagram (as they should), and sends an ICMP Echo Reply datagram back to what they think is the source. If many hosts are "alive" on the LAN, the amplification factor can be considerably (100+ is not uncommon).
4. The attacker can use largish packets (typically up to ethernet maximum) to increase the "effectiveness" of the attack, and the faster network connection the attacker has, the more damage he can inflict on the target and the target's network.

Not only can the attacker cause problems for the target host, the influx of traffic can in fact be so great as to have a seriously negative effect on the upstream network(s) from the target. In fact, those institutions being abused as amplifier networks can also be similarly affected, in that their network connection can be swamped by the Echo Reply packets destined for the target.

In this example, answer choice D is correct as it prevents all ICMP messages destined to the broadcast IP address.

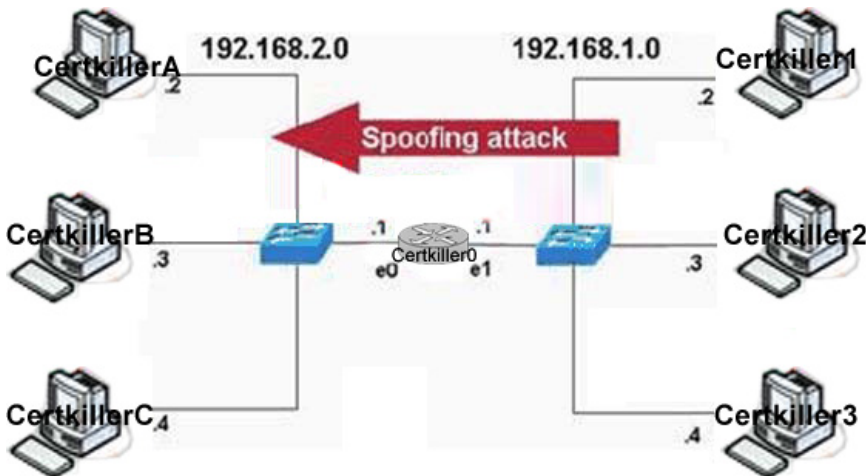
Note: The Cisco IOS command "no ip directed-broadcasts" is also an effective way to prevent smurf and fraggle attacks on the network.

Incorrect Answers:

- A. This will permit all traffic sourced from the 192.168.2.0/24 network, including the smurf attack packets.
- B. This choice will deny all traffic sourced from the 192.168.1.0 incoming on the e0 interface. Although this is probably a good choice, as it will effectively prevent all spoofed IP traffic (as the 192.168.1.0/24 network should never be a source IP address in the incoming direction of this interface) we wish to only prevent the smurfed traffic, so E is a better choice.
- C. This choice will only permit traffic that is destined to the 192.168.1.0 network. If additional networks exist behind the 192.168.1.0 network, such as traffic to the Internet, it will not be allowed through the CK1 router.
- E. It would be preferable to stop the attack before the replies are sent, rather than simply filtering the replies.

QUESTION 341

Part of a Certkiller network is shown below:



In the Certkiller network shown above, you want to block all IP spoofing attacks that originate on the 192.168.1.0/24 subnet using a spoofed address outside the 192.168.1.0/24 range from being sent into the 192.168.2.0/24 network. However, normal traffic must be permitted. No access lists currently exist on the router. Which of the following configuration excerpts would accomplish this task when applied to E1 on Certkiller 0 as an incoming filter?

- A. access-list 1 permit 192.168.1.0 0.0.0.255
- B. access-list 100 permit ip any 192.168.2.0 0.0.0.255
- C. access-list 1 deny 192.168.2.0 0.0.0.255
access-list 1 permit any
- D. access-list 1 deny 192.168.2.0 0.0.0.255
access-list 1 deny 192.168.1.0 0.0.0.255
access-list 1 permit any
- E. access-list 100 deny ip 192.168.2.0 0.0.0.255 any
access-list 100 permit ip any any

Answer: A

Explanation:

The access list in choice A will prevent all incoming traffic sourced from the 192.168.2.0/24 network from interface Ethernet 1 of router CK1 due to the implicit deny all. In the diagram above, hosts on the 192.168.2.0 network should only be used as a destination for traffic coming from this interface. Only traffic sourced from 192.168.1.0/24 should be seen in the input direction of this interface on CK1. If any traffic does not match the access list on choice A it could only be the result of a spoofed IP address and should be dropped.

Incorrect Answers:

- B. This will allow all traffic (from any source) to reach the 192.168.2.0 network. This will not prevent spoofed IP addresses from the network on E1 to go through.
- C. This would prevent spoofed packets that were spoofed only from the 192.168.2.0/24 network. This will not prevent all spoofed addresses outside of the 192.168.1.0 network, as required.
- D. This will prevent the two networks from communicating at all.

E. This choice could also be used to prevent the spoofed traffic as required, but it will only prevent spoofed traffic that is IP based. Therefore, the access list in choice C is a better fit for this situation.

QUESTION 342

The Certkiller network administrator is concerned with spoofing based attacks. Which IOS feature can be used to defend against these spoofing attacks?

- A. TCP Intercept
- B. IP Source Guard and/or Unicast RPF
- C. Auth-proxy
- D. Lock and Key ACL and/or Reflexive ACL
- E. IOS IPS
- F. IOS Firewall (CBAC)
- G. None of the above

Answer: B

Explanation:

The proper deployment and configuration of Unicast RPF provides the most effective means of anti-spoofing protection against attacks with spoofed source IP addresses. IP source guard provides the most effective means of anti-spoofing protection against attacks with spoofed source MAC addresses. Deployment as close to all traffic sources as possible provides maximum effectiveness.

Reference: <http://tools.cisco.com/security/center/getDocument.x?id=442>

QUESTION 343

The Certkiller security administrator is determining security policies for deploying Unicast Reverse Path Forwarding (uRPF) on the network. Which of the following is NOT an action that uRPF can perform in this network?

- A. You can have an ACL configured and combined with your uRPF configuration
- B. uRPF checks to see if any packet received at a router interface arrives on the best return path
- C. uRPF inspects IP Packets encapsulated in tunnels, such as GRE
- D. uRPF examines all packets received to make sure that the source address and source interface appear in the routing table and match the interface where the packets was received
- E. uRPF events can be logged by specifying the logging option for the ACL entries used by the unicast RPF command
- F. None of the above

Answer: C

Explanation:

Consider the following points in determining your policy for deploying Unicast RPF:

Unicast RPF must be applied at the interface downstream from the larger portion of the network, preferably at the edges of your network.

The further downstream you apply Unicast RPF, the finer the granularity you have in mitigating address spoofing and in identifying the sources of spoofed addresses. For example, applying Unicast RPF on an aggregation router helps mitigate attacks from many downstream networks or clients and is simple to administer, but it does not help identify the source of the attack. Applying Unicast RPF at the network access server helps limit the scope of the attack and trace the source of the attack; however, deploying Unicast RPF across many sites does add to the administration cost of operating the network.

The more entities that deploy Unicast RPF across Internet, intranet, and extranet resources, the better the chances of mitigating large-scale network disruptions throughout the Internet community, and the better the chances of tracing the source of an attack.


Unicast RPF will not inspect IP packets encapsulated in tunnels, such as GRE, LT2P, or PPTP. Unicast RPF must be configured at a home gateway so that Unicast RPF processes network traffic only after the tunneling and encryption layers have been stripped off the packets.

Reference:

http://www.cisco.com/en/US/products/ps6350/products_configuration_guide_chapter09186a00804fdef9.html

QUESTION 344

The Certkiller network is connected to the Internet as shown in the diagram below:



Note: Private addressing is only used for reference

Certkiller 1 is currently configured and passing traffic. You want to block all IP spoofing attacks that originate in the Internet from being sent into the 192.168.1.0/24 network. However, normal traffic must be permitted. No access lists currently exist on the router. What configuration excerpt would accomplish this task when applied to Certkiller 1?

- A.

```
access-list 100 permit ip any 192.168.1.0 0.0.0.255
access-list 100 deny any any
interface Ethernet 0
access-group 100 in
```
- B.

```
ip cef
```

```
interface Ethernet 0
ip verify unicast reverse-path
C. ip cef
interface Ethernet 1
ip verify unicast reverse-path
D. access-list 100 permit icmp 192.168.1.0 0.0.0.255 any echo
access list 100 deny ip any any
interface Ethernet 1
access-group 100 out
E. access-list 100 permit icmp 192.168.1.0 0.0.0.255 any echo
access list 100 deny ip any any
interface Ethernet 0
access-group 100 in
```

Answer: B

Explanation:

Use the ip verify unicast reverse-path interface command to mitigate problems caused by malformed or forged (spoofed) IP source addresses that pass through a router. Malformed or forged source addresses can indicate denial-of-service (DoS) attacks based on source IP address spoofing.

When Unicast RPF is enabled on an interface, the router examines all packets received on that interface. The router checks to make sure that the source address appears in the routing table and matches the interface on which the packet was received. This "look backwards" ability is available only when Cisco Express Forwarding (CEF) is enabled on the router because the lookup relies on the presence of the Forwarding Information Base (FIB). CEF generates the FIB as part of its operation.

Unicast RPF is an input function and is applied only on the input interface of a router at the upstream end of a connection.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_command_reference_chapter09186a00800c

QUESTION 345

In order to enhance security on the Certkiller network, unicast Reverse Path Forwarding (uRPF) has been implemented. Regarding how uRPF works to help in preventing malformed or forged IP source addresses, which of the following options is true?

- A. It is applied only on the input interface of a router
- B. It is applied only on the output interface of a router
- C. It can be configured either on the input or output interface of a router
- D. It can't be configured on a router interface
- E. It is configured under any routing protocol process
- F. None of the above

Answer: A

Explanation:

When Unicast RPF is enabled on an interface, the router examines all packets received on that interface. The router checks to make sure that the source address appears in the routing table and matches the interface on which the packet was received. This "look backwards" ability is available only when Cisco Express Forwarding (CEF) is enabled on the router because the lookup relies on the presence of the Forwarding Information Base (FIB). CEF generates the FIB as part of its operation.

Unicast RPF is an input function and is applied only on the input interface of a router at the upstream end of a connection.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_command_reference_chapter09186a00800c

QUESTION 346

A new Certkiller user connects her PC to the LAN where IP Source Guard has been configured. What will be the results of enabling IP Source guard on an untrusted switch port that does not have DHCP snooping enabled?

- A. The switch will run out of the ACL hardware resources
- B. DHCP requests will be switched in the software which may result in lengthy response times
- C. The DHCP server reply will be dropped and the client will not be able to obtain an IP Address
- D. All DHCP requests will pass through the switch untested
- E. None of the above

Answer: C

Explanation:

IP source guard prevents IP spoofing by allowing only the IP addresses that are obtained through DHCP snooping on a particular port. Initially, all IP traffic on the port is blocked except for the DHCP packets that are captured by DHCP snooping. When a client receives a valid IP address from the DHCP server, a port access control list (PACL) is installed on the port that permits the traffic from the IP address. This process restricts the client IP traffic to those source IP addresses that are obtained from the DHCP server; any IP traffic with a source IP address other than that in the PACLs permit list is filtered out. This filtering limits the ability of a host to attack the network by claiming a neighbor host's IP address.

IP source guard uses source IP address filtering, which filters the IP traffic that is based on its source IP address. Only the IP traffic with a source IP address that matches the IP source binding entry is permitted.

A port's IP source address filter is changed when a new DHCP-snooping binding entry for a port is created or deleted. The port PACL is modified and reapplied in the hardware to reflect the IP source binding change. By default, if you enable IP source guard without

any DHCP-snooping bindings on the port, a default PACL that denies all IP traffic is installed on the port. When you disable IP source guard, any IP source filter PACL is removed from the port.

Reference:

http://www.cisco.com/en/US/products/hw/switches/ps708/products_configuration_guide_chapter09186a008022f

QUESTION 347

Certkiller is using the firewall feature in Cisco IOS routers at the remote office locations. Which three statements are true regarding IOS firewall Configuration? (Choose three)

- A. The ip inspection rule can be applied in the inbound direction on the secured interface
- B. The ip inspection rule can be applied in the outbound direction on the unsecured interface
- C. The ACL applied in the outbound direction on the unsecured interface must be an extended ACL
- D. The ACL applied in the inbound direction on the unsecured interface must be an extended ACL
- E. For temporary openings to be created dynamically by Cisco IOS Firewall, the access-list for the returning traffic must be a standard ACL
- F. For temporary openings to be created dynamically by Cisco IOS firewall, the ip inspection rule must be applied to the secured interface

Answer: A, B, D

Explanation:

The below excerpt is from the Cisco Press book "The Cisco IOS Firewall Feature Set" By Anupam Tewari. Note the places in italics where the correct answers can be found: For CBAC (IOS Firewall) to function properly, it is essential that the access list be configured appropriately on the interfaces. An extended access list must be used for the creation of temporary openings.

The five steps involved in configuring CBAC are as follows:

1. Choose an interface. CBAC can identify any interface as an internal or external interface. Unlike Firewall, CBAC has no concept of inside or outside; instead, it is concerned with the direction of the first packet initiating the conversation. Sessions originating from the external side are not permitted. For example, when user X in ABC Company tries to connect to the Internet, the direction of the conversation is from the ABC Company to the Internet. The router interface that connects to user X is considered internal, and the interface connected to the Internet is considered external.
2. Configure IP access list at the interface. CBAC permits less traffic than necessary to get similar functionality with static access lists. When configuring an access list on the internal interface, the inbound access list (at the internal interface) or outbound (at the external interface) can be standard or extended. These access lists permit the CBAC to inspect the traffic. The outbound access list (on internal interface) and the inbound access list (at the external interface), on the other hand, should always be extended.

On the external interface, the outbound access list can be standard or extended, but the inbound access list must be an extended list. The inbound access list will deny the traffic to be inspected by CBAC. This denied traffic would be permitted in by the temporary openings created by the CBAC.

3. Configure global timeouts and thresholds.

Global timeouts are used to configure the duration for which a hole in the firewall is maintained to allow in the return traffic. Thresholds are configured to shield the network from denial-of-service (DoS) attacks. The sessions that are not established under the configured parameters are dropped.

For example, the `ip inspect tcp synwait-time 30` command says to drop all the TCP sessions that are not established in 30 seconds. Similar timeouts can be set up for FIN-exchange, TCP or UDP idle timeouts, and DNS timeouts.

4. Define an inspection rule. An inspection rule defines which application-layer protocol will be inspected by the CBAC. After configuring the inspection rule for an application-layer protocol, all the packets for that protocol are permitted out and are allowed back in. Each protocol packet is inspected to maintain the session information; the same session information is used to determine whether the packet is the part of valid session.

`ip inspect name inspection-name protocol [timeout seconds]` is a global command used to configure the inspection rule. Protocol keywords can be `tcp`, `udp`, `ftp-cmd`, or `http`. `timeout` refers to the period of protocol inactivity before dropping the connection.

5. Apply the inspection rule to the interface. The `ip inspect inspection-name {in | out}` command is used to apply the inspection rule to an interface. The keyword `in` is used for inbound traffic when the CBAC is applied on the internal (trusted, or secure) interface. The keyword `out` is used for outbound traffic when the CBAC is applied on the external, unsecured interface.

Reference: <http://www.ciscopress.com/articles/article.asp?p=26533&seqNum=5&rl=1>

QUESTION 348

Router CK1 is an MPLS LSR router in the Certkiller service provider network. A labeled packet is received by this router for which there is not a label entry present in the LFIB. The router performs which of the following actions?

- A. It uses a default label for forwarding
- B. It strips the label and does a lookup in the FIB using the IP destination address
- C. It drops the packet
- D. It uses LDP to create an LSP and a new entry in the LFIB for that label
- E. None of the above

Answer: C

Explanation:


In normal operation, an LSR should receive only a labeled packet with a label at the top of the stack that is known to the LSR, because the LSR should have previously advertised that label. However, it is possible for something to go wrong in the MPLS network and the LSR to start receiving labeled packets with a top label that the LSR does not find in

its LFIB. The LSR can theoretically try two things: strip off the labels and try to forward the packet, or drop the packet. The Cisco LSR drops the packet. This is the right thing to do, because this LSR did not assign the top label, and it does not know what kind of packet is behind the label stack

Reference: <http://www.ciscopress.com/articles/article.asp?p=680824>

QUESTION 349

Router CK1 , CK2 , and CK3 are MPLS provider edge routers as shown below:



How many VRF tables will be on router CK2 ?

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

Answer: B

Explanation:


Router CK2 must keep track of two separate VPN customers, VPN C and VPN D. Each VPN will have its own VRF table.

Incorrect Answers:

A. This would be the correct choice if the question were asking how many routing tables existed on CK2 , as there would only be one routing table.

QUESTION 350

Certkiller is an ISP providing MPLS VPNs to their customers. Part of the Certkiller network is shown in the following exhibit:



You are trying to verify that the correct paths are configured within the CK routers. Which networks should show up in the same Forwarding Equivalence Class (FEC) on Router CK1 ?

- A. 172.16.3.0/24 and 192.1.1.0/28
- B. 172.16.1.0/24, 172.16.2.0/24, and 172.16.3.0/24
- C. 172.16.1.0/24 and 193.1.1.0/28
- D. 172.16.1.0/24 and 172.16.2.0/24
- E. 172.16.0.0/16 and 193.0.0.0/8

Answer: C

Explanation:

Both networks in answer C are along the same path, so they would both be in the same FEC.

Reference:

For a better understanding of MPLS TE and FEC see the white paper link below:

http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/mwglp_wp.htm

QUESTION 351

In the Certkiller MPLS network, which of the following parameters are used to determine a Forwarding Equivalent Class, or FEC? (Choose two.)

- A. IP Prefix
- B. Layer 2 Circuit
- C. BGP MED Value
- D. RSVP Request from CE for bandwidth reservation

Answer: A, B

Explanation:

Forwarding Equivalency Class (FEC) is a set of packets which will be forwarded in the same manner (e.g., over the same path with the same forwarding treatment). Typically packets belonging to the same FEC will follow the same path in the MPLS domain. While assigning a packet to an FEC the ingress LSR may look at the IP header and also some other information such as the interface on which this packet arrived. The FEC to which a packet is assigned is identified by a label.

One example of an FEC is a set of unicast packets whose network layer destination address matches a particular IP address prefix. A set of multicast packets with the same source and destination network layer addresses is another example of an FEC. Yet

another example is a set of unicast packets whose destination addresses match a particular IP address prefix and whose Type of Service bits are the same. Layer 2 circuits can also be used, as in layer 2 MPLS VPN's such as those defined by the Martini draft or through the notion of "pseudo-wire" networks.

Reference: <http://www.mplsrc.com/faq2.shtml>

QUESTION 352

What is the signaling protocol used for the MPLS fast reroute (FRR) feature in the Certkiller Service Provider network?

- A. B-ISUP
- B. LDP
- C. RSVP
- D. SS7
- E. TDP
- F. None of the above

Answer: C

Explanation:

Cisco FRR utilizes MPLS label stacking with RSVP signaling to create a backup tunnel around the link or node that needs to be protected. On detection of loss of signal from the link, the MPLS FRR application in Cisco IOS Software starts forwarding the traffic onto the backup tunnel, transparent to end users or applications such as VoIP or video, in 50 ms or less (actual failover time may be greater or less than 50ms, depending on the hardware platform, the number of TE Tunnels and/or Network prefixes).

Incorrect Answers:

A. The ITU-T's broadband ISDN user part (B-ISUP) is based on signaling system no. 7 (SS7) and is used for signaling between the nodes of a public ATM network, that is, across an NNI.

B, E. LDP is the Label Distribution Protocol used to distribute label information across the MPLS network. TDP is the Tag Distribution Protocol, which is the Cisco proprietary method of distributing tags across the network in tag switching. MPLS is founded on Cisco's tag switching.

D. SS7 is a signaling protocol normally found in Voice circuits. It is not related to MPLS fast reroute.

Reference:

White Paper, Deploying Guaranteed-Bandwidth Services with MPLS

http://www.cisco.com/en/US/tech/CK436/CK428/technologies_white_paper09186a00800a3e69.shtml

QUESTION 353

In the Certkiller MPLS network, you want the ability to send some traffic around less congested links. To do this, you want to bypass the normal routed hop-by-hop paths.

What technology should you implement?

What should you use?

- A. Traffic engineering
- B. Traffic tunneling
- C. Traffic policing
- D. Traffic shaping
- E. Traffic routing

Answer: A

Explanation:

Traffic engineering allows you to bypass the routing protocol information to send traffic over alternative paths.

Incorrect Answers:

- B. Using tunnels will not force the traffic over the tunnels to bypass the normal hop by hop routed topology.
- C, D. Traffic policing and traffic shaping are methods of QoS.
- E. Traffic routing is not a well defined Cisco term.

QUESTION 354

Certkiller is and ISP providing MPLS standards according to RFC 4364 standards. By using a unique route distinguisher per VRF, this Certkiller RFC 4364 L3 VPN provides the ability for which of the following?

- A. Unique IGP per VRF
- B. Multi-homed access
- C. Traffic Engineering
- D. Overlapping IP Address space
- E. None of the above

Answer: D

Explanation:

RFC 4364 describes a method by which a Service Provider may use an IP backbone to provide IP Virtual Private Networks (VPNs) for its customers. This method uses a "peer model", in which the customers' edge routers (CE routers) send their routes to the Service Provider's edge routers (PE routers). Border Gateway Protocol (BGP) [BGP, BGP-MP] is then used by the Service Provider to exchange the routes of a particular VPN among the PE routers that are attached to that VPN. This is done in a way that ensures that routes from different VPNs remain distinct and separate, even if two VPNs have an overlapping address space.


The BGP Multiprotocol Extensions [BGP-MP] allow BGP to carry routes from multiple "address families". We introduce the notion of the "VPN-IPv4 address family". A VPN-IPv4 address is a 12-byte quantity, beginning with an 8-byte Route Distinguisher (RD) and ending with a 4-byte IPv4 address. If several VPNs use the same IPv4 address prefix, the PEs translate these into unique VPN-IPv4 address prefixes (route distinguisher). This ensures that if the same address is used in several different VPNs, it

is possible for BGP to carry several completely different routes to that address, one for each VPN.

Reference: <http://www.networksorcery.com/enp/rfc/rfc4364.txt>

QUESTION 355

Router CK1 , CK2 , and CK3 are MPLS provider edge routers as shown below:



How many routing tables should there be on Router CK2 ?

- A. 1
- B. 2
- C. 3
- D. 4
- E. 0

Answer: A

Explanation:

There should always be only one routing table within an MPLS network.

Incorrect Answers:

B. In this case there would be 2 VRF (Virtual Routing and Forwarding) tables, but still only 1 routing table.

QUESTION 356

The Certkiller network is being used to provide MPLS VPN services to its customers. In an MPLS label within this network, if the stack bit is set to 1, which of the following is correct?

- A. The stack bit is a Cisco-only implementation and is only used when TDP is the label distribution protocol
- B. The stack bit is reserved for future use
- C. The stack bit is only used when LDP is the label distribution protocol
- D. It signifies that this is the top entry in the label stack and remains set to 1 until only the last entry, the bottom label is reached
- E. It signifies the last entry in the label stack

packet is removed by a Label Switched Router (LSR) before the packet is passed to an adjacent Label Edge Router (LER).

The process is important in a Layer 3 MPLS VPN (RFC2547) environment as it reduces the load on the LER. If this process didn't happen, the LER would have to perform at least 2 label lookups:

The outer label, identifying that the packet was destined to have its label stripped on this router.

The inner label, to identify which Virtual Routing/Forwarding (VRF) instance to use for the subsequent IP routing lookup.

In a large network this can result in the CPU load on the LER reaching unacceptable levels. By having PHP for an LER done on the LSRs connected to it, the load is effectively distributed among its neighbor routers.

PHP functionality is achieved by the LER advertising a label with a value of 3 to its neighbors. This label is defined as implicit-null and informs the neighboring LSR(s) to perform PHP.

QUESTION 358

In the Certkiller service provider network, what can be said of TTL Propagation when IP propagation has been disabled (MPLS)?

- A. The TTL field of the MPLS label header is set to 255
- B. TTL field true the IP packet is copied into the TTL field of the MPLS control header at the ingress edge LSR
- C. The TTL field of the packet is set to 0
- D. TTL propagation only disabled in the ingress edge LSR
- E. TTL propagation cannot be disabled in MPLS domain
- F. None of the above

Answer: A

Explanation:

Tunnel Ingress Processing (Push):

For each pushed Uniform model label, the TTL is copied from the label/IP-packet immediately underneath it.

For each pushed Pipe model label, the TTL field is set to a value configured by the network operator. In most implementations, this value is set to 255 by default.

Reference: "TTL Processing in MPLS Networks" February 2001

<http://quimby.gnus.org/internet-drafts/draft-agarwal-mpls-ttl-00.txt>

QUESTION 359

Certkiller, Inc is an ISP providing MPLS services to its customers. Which fundamental modifications, related to traffic forwarding for these customers, does MPLS introduce? (Choose two)

- A. For multicast routing, labels are assigned to IPMC groups

- B. IP destination routing is reduced to label lookup with the MPLS network
- C. For unicast routing, labels are assigned to FECs (i.e IP prefixes)
- D. IP Lookup is done on every hop with the MPLS core

Answer: B, C


Explanation:

MPLS works by tagging packets with an identifier (a label) to distinguish the LSPs.

When a packet is received, the router uses this label (and sometimes also the link over which it was received) to identify the LSP. It then looks up the LSP in its own forwarding table to determine the best link over which to forward the packet, and the label to use on this next hop.

A different label is used for each hop, and it is chosen by the router or switch performing the forwarding operation. This allows the use of very fast and simple forwarding engines, as the router can select the label to minimize processing.

Ingress routers at the edge of the MPLS network use the packet's destination address to determine which LSP to use. Inside the network, the MPLS routers use only the LSP labels to forward the packet to the egress router.



In the diagram above, LSR (Label Switched Router) A uses the destination IP address on each packet to select the LSP, which determines the next hop and initial label for each packet (21 and 17). When LSR B receives the packets, it uses these labels to identify the LSPs, from which it determines the next hops (LSRs D and C) and labels (47 and 11). The egress routers (LSRs D and C) strip off the final label and route the packet out of the network.

As MPLS uses only the label to forward packets, it is protocol-independent, hence the term "Multi-Protocol" in MPLS. Packet forwarding has been defined for all types of layer-2 link technologies, with a different label encoding used in each case.

QUESTION 360

In the Certkiller MPLS network, which protocol is used to distribute traffic engineering information? (Choose all that apply)

- A. EIGRP or IGRP
- B. BGP MEDs

- C. OSPF Opaque LSAs or IS-IS TLVs
- D. RTP or RTCP packets
- E. None of the above

Answer: C

Explanation:

OSPF Opaque LSA provides a generalized mechanism for OSPF to carry additional information. The new information can be used directly by OSPF or indirectly by other applications, which use OSPF to distribute information. This document defines how to use opaque LSA to carry additional information for traffic engineering.

Reference:

<http://www.watersprings.org/links/mlr/id/draft-yeung-ospf-traffic-00.txt>

Incorrect Answers:

- A. Interior routing protocols are used to distribute the interior network routes, but are not used to carry TE information.
- B. BGP MEDs are used to influence the incoming traffic from a BGP neighbor.
- D. RTP is the Real Time Protocol, used for setting up voice and video traffic.

QUESTION 361

Certkiller is using IP version 6 addressing within their network. You have been tasked with summarizing the addressing so that they can be expressed in the shortest form possible. The IPv6 address of 2013:0000:130F:0000:0000:09c0:876A:130B can be expressed most efficiently as which of the following:

- A. 2013:0:130F:0:0:09C0:876A:130B
- B. 2013::130F::9C0:876A:130B
- C. 2013:0:130F:0:0:9c0:876A:130B
- D. 213::13F::9C:876A:13B
- E. 2013:0:130F::9C0:876A:130B
- F. None of the above

Answer: E

Explanation:

IPv6 addresses are written as eight sets of four hexadecimal digits:

FEDC:BA98:0000:0000:0000:0000.7654.3210

To make writing the addresses easier, groups of zeros that appear in the address may be replaced with double colons:

FEDC:BA98::7654:3210

Only one group of contiguous zeros may be condensed within an address.

Reference: <http://www.netcordia.com/tnm/tnm34/ipv6.html>

QUESTION 362

Which of the following are legal representations of the IPv6 prefix 12AB00000000CD3? (Choose Two)

- A. 12AB:0000:0000:CD30:0000:0000:0000:0000/60
- B. 12AB:0:0:CD3/60
- C. 12AB:: CD3/60
- D. 12AB:0:0:CD30::/60
- E. 12AB::CD3::/60

Answer: A, D

Explanation:

IPv6 Address Compaction:

In IPv6, the leading zeroes in a 16-bit segment can be compacted.

Example:

fe80:0210:1100:0006:0030:a4ff:000c:0097 becomes fe80:210:1100:6:30:a4ff:c:97

All zeroes in one or more contiguous 16-bit segments can also be represented with a double colon(::).

Example:

ff02:0000:0000:0000:0000:0000:0000:0001 Becomes ff02::1.

However, double colons can only be used once in any one address.

Example:

2001:0000:0000:0013:0000:0000:0b0c:3701

Can be:

2001::13:0:0:b0c:3701 or 2001:0:0:13::b0c:3701, but not 2001::13::b0c:3701

For IPv6 Prefix Representation, CIDR-like notations are used to specify the address prefix length. Examples of this include:

3ffe:0:0:2300:ce21:233:fea0:bc94/60 and 201:468:1102:1::1/64

Finally, an IPv6 Prefix Compaction example is:

2002:0000:0000:18d0:0000:0000:0000:0000/60

Can be represented as:

2002::18d0:0:0:0/60

2002:0:0:18d0::/60

In the example shown in this question, 12AB00000000CD3 can be represented either as:

12AB:0000:0000:CD30:0000:0000:0000:0000/60 or 12AB:0:0:CD30::/60

QUESTION 363

IPv6 is being implemented in the Certkiller network. Which of the following is a valid IPv6 Address Type? (Choose Three)

- A. Broadcast
- B. Multicast
- C. Anycast
- D. Unicast

Answer: B, C, D

Explanation:

With IPv6 the IETF sought to carve the new address space into functional categories, each of which would enable more-efficient routing through a more-sophisticated hierarchy. These functional categories are known as anycast, unicast, and multicast. Noticeably absent was the broadcast address type. IPv6 doesn't use broadcast addresses, but it satisfies that function through the multicast address type.

IPv6 defines three address types:

Unicast:

Identifies an interface of an individual node.

Multicast:

Identifies a group of interfaces, usually on different nodes. Packets that are sent to the multicast address go to all members of the multicast group.

Anycast:

Identifies a group of interfaces, usually on different nodes. Packets that are sent to the anycast address go to the anycast group member node that is physically closest to the sender.

QUESTION 364

A new Certkiller LAN segment is using IPv6 and a host is sending solicitation messages. While doing the IPv6 address resolution, a node does NOT send a Neighbor Solicitation (NS) message in order to: (Choose three)

- A. Discover the layer 2 multicast address of the destination node
- B. Discover the layer 2 address of the destination node based on the destination IPv6 address
- C. Discover the IPv6 address of the destination node based on the destination layer 2 address
- D. Discover the solicited node multicast address of the destination node

Answer: A, C, D

Explanation:

Whenever a system needs to find out the link address for another system residing on the same link, it sends a neighbor solicitation to the solicited node address to which the IPv6 address of the remote system maps. The source host includes its own MAC address in the neighbor solicitation, so the neighbor knows where to send the reply. This message is sent based on the link address, not the IPv6 address. A good summary article called "IPv6 Internals" discusses this and can be found at the link below.

Reference:

http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_9-3/ipv6_internals.html

QUESTION 365

The Certkiller network is migrating from IPv4 to IPv6. What IPv6 header field has a similar function as the "Type of Service" field in an IPv4 header?

- A. Flow Label
- B. Version
- C. Next Header
- D. Traffic Class
- E. None of above

Answer: D

Explanation:

The IPV6 header is shown below:

Packet Fields

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31																						
Version				Traffic Class								Flow Label																																									
Payload Length																Next Header								Hop Limit																													
Source Address																																																					
Destination Address																																																					

Version

The version field exists for the same purpose as in IPv4, namely to differentiate between different versions of the protocol. Obviously it carries the value of "6" for all IPv6 packets.

Traffic Class

This is similar to the Type of Service (ToS) bits in IPv4, except that it tries to describe the type of traffic rather than the type of service.

Flow Label

This is used to label a "flow" of packets. One problem with a packet switching network is that there's no good way to tell if a bunch of separate packets are related to each other.

Payload Length

This is the length of the packet not including the IPv6 header. It replaces the Total Length field in IPv4. Unlike IPv4, IPv6 headers are of a fixed length, and there's no point in adding a constant to the length field.

Next Header

This is like the Protocol field from IPv4 - it identifies the higher level protocol. However, unlike IPv4, this is not always a transport layer protocol like TCP or UDP. It takes on special values which are used to implement the IPv6 extension and option mechanisms.

Hop Limit

This is like the TTL (Time To Live) field from IPv4 - it's used to minimize the resources that a packet consumes in the event of a routing loop. Its name has been changed to reflect the actual usage of the field. Namely, it is decremented every hop, rather than every second as was originally intended.

Source/Destination Address

These have the same purpose in IPv4, but have been expanded to 128 bits.

Reference: <http://www.mit.edu/~elliot/internet/1998/ipv6-notes.html>

QUESTION 366

Router CK1 needs to be configured with RIP to support IPv6. Which of the following are the minimum required tasks to configure IPv6 RIP on a Cisco router? (Choose Two)

- A. Customizing IPv6 RIP
- B. Configuring Tags for RIP routes
- C. Enable IPv6 RIP on the interface
- D. Configuring IPv6 Multicast routing
- E. Enable IPv6 on the router

Answer: C, E

Explanation:

Before configuring the router to run IPv6 RIP, two things must be done. First, globally enable IPv6 using the "ipv6 unicast-routing" global configuration command. Secondly, enable IPv6 on any interfaces on which IPv6 RIP is to be enabled. These are the required minimum configuration prerequisites.

Incorrect Answers:

A, B, D: These choices describe optional IPv6 RIP configurations. A list of optional components for configuring RIP for IPv6 is shown below:

Customizing IPv6 RIP (optional)

Redistributing Routes into an IPv6 Routing Process (optional)

Configuring Tags for RIP Routes (optional)

Filtering IPv6 RIP Configuration Updates (optional)

Verifying IPv6 RIP Configuration and Operation (optional)

Reference:

http://www.cisco.com/en/US/products/ps6350/products_configuration_guide_chapter09186a00801d6601.html

QUESTION 367

The Certkiller network is implementing IP multicast with their IPv6 network. Which statement below is true in reference to IPv6 multicast?

- A. MSDP is required for IPV6 multicast
- B. PIM dense mode is not part on IPV6 multicast
- C. IPv6 multicast uses Multicast Listener Discovery (MLD)
- D. The first 8 bits of an IPv6 multicast address are always FF (1111 1111)
- E. None of the above

Answer: C

Explanation:

The Multicast Listener Discovery Protocol (MLD) is used by IPv6 routers to discover the

presence of multicast listeners (i.e., nodes that wish to receive multicast packets) on their directly attached links, and to discover specifically which multicast addresses are of interest to those neighboring nodes. Note that a multicast router may itself be a listener of one or more multicast addresses; in this case it performs both the "multicast router part" and the "multicast address listener part" of the protocol, to collect the multicast listener information needed by its multicast routing protocol on the one hand, and to inform itself and other neighboring multicast routers of its listening state on the other hand.

Reference: <http://www.ietf.org/rfc/rfc3810.txt>

QUESTION 368

What is a main difference between the IPv6 and IPv4 multicast?


- A. IPv6 has significantly more address space (128 bits), so overlapping addresses are less likely.
- B. Multicast Listener Discovery (MLD) replaces IGMP in IPv6 multicasts.
- C. MSDP and dense mode multicast is not part of IPv6 multicast.
- D. The first 8 bits of IPv6 Multicast address are always FF (1111 1111).
- E. All of the above

Answer: E

Explanation:

A multicast address identifies multiple interfaces, and is used for one-to-many communication. With the appropriate multicast routing topology, packets addressed to a multicast address are delivered to all interfaces that are identified by the address. IPv6 multicast addresses have the Format Prefix of 1111 1111. An IPv6 address is simple to classify as multicast because it always begins with FF. Multicast addresses cannot be used as source addresses.

Multicast addresses include additional structure to identify their flags, scope, and multicast group, as shown in the following illustration:



MLD is used to exchange membership status information between IPv6 routers that support multicasting and members of multicast groups on a network segment. Membership in a multicast group is reported by individual member hosts, and membership status is periodically polled by multicast routers. MLD is defined in RFC 2710, "Multicast Listener Discovery (MLD) for IPv6." IGMP is not used in IPv6.

QUESTION 369

Within the Certkiller IP multicast network, what best describes the functionality of the Multicast Listener Discovery (MLD)?

- A. IPv6 routers use MLD to discover multicast listeners on directly attached links.
- B. For each Unicast and Anycast addresses configured on an interface of the node or a

router, a corresponding entry is automatically enabled.

C. The MLD addresses is scoped to the local link.

D. Since the ARP is not used in the IPv6, the MLD is used by nodes and routers to learn the link layer address of the neighbor nodes and routers on the same local link.

E. MLD is used to verify if the IPv6 address is already in use on it's local link, before it configure it's own IPv6 address with stateless auto-configuration.

F. None of the above

Answer: A

Explanation:

The purpose of Multicast Listener Discovery (MLD) is to enable each IPv6 router to discover the presence of multicast listeners (that is, nodes wishing to receive multicast packets) on its directly attached links, and to discover specifically which multicast addresses are of interest to those neighboring nodes. This information is then provided to whichever multicast routing protocol is being used by the router, in order to ensure that multicast packets are delivered to all links where there are interested receivers.

MLD is an asymmetric protocol, specifying different behaviors for multicast listeners and for routers. For those multicast addresses to which a router itself is listening, the router performs both parts of the protocol, including responding to its own messages.

If a router has more than one interface to the same link, it need perform the router part of MLD over only one of those interfaces. Listeners, on the other hand, must perform the listener part of MLD on all interfaces from which an application or upper-layer protocol has requested reception of multicast packets.

QUESTION 370

What best describes the IPv6 Solicited-node Multicast address?

A. For each unicast and anycast addresses configured on an interface of the node or a router, a corresponding solicited-node multicast addresses is automatically enabled.

B. The solicited-node multicast address is scoped at the local link.

C. Since ARP is not used the in IPv6, the solicited-node multicast addresses is used by nodes and router to learn the link layer address of the neighbor nodes and routers on the same local link.

D. Duplicate Address Detection (DAD) is used to verify if the IPv6 address is already in used on it's local link, before it configure it's own IPv6 address with stateless auto-configuration, Solicited-node multicast addresses probe the local link to make sure.

E. All of the above

F. None of the above

Answer: E

Explanation:

In IP version 6, the solicited-node multicast address facilitates efficient querying of network nodes during address resolution. IPv6 uses the Neighbor Solicitation message to perform address resolution. In IPv4, the ARP Request frame is sent to the MAC-level

broadcast, disturbing all nodes on the network segment regardless of whether a node is running IPv4. For IPv6, instead of using ARP requests and disturbing all IPv6 nodes on the local link by using the local-link scope all-nodes address, the solicited-node multicast address is used as the Neighbor Solicitation message destination.

The solicited-node multicast address consists of the prefix FF02::1:FF00:0/104 and the last 24-bits of the IPv6 address that is being resolved.

The following steps show an example of how the solicited-node address is handled for the node with the link-local IPv6 address of FE80::2AA:FF:FE28:9C5A, and the corresponding solicited-node address is FF02::1:FF28:9C5A:

1. To resolve the FE80::2AA:FF:FE28:9C5A address to its link layer address, a node sends a Neighbor Solicitation message to the solicited-node address of

FF02::1:FF28:9C5A.

2. The node using the address of FE80::2AA:FF:FE28:9C5A is listening for multicast traffic at the solicited-node address FF02::1:FF28:9C5

A. For interfaces that correspond

to a physical network adapter, it has registered the corresponding multicast address with the network adapter.

As shown in this example, by using the solicited-node multicast address, address resolution that commonly occurs on a link can occur without disturbing all network nodes. In fact, very few nodes are disturbed during address resolution. Because of the relationship between the network interface MAC address, the IPv6 interface ID, and the solicited-node address, in practice, the solicited-node address acts as a pseudo-unicast address for efficient address resolution.

Reference:

<http://msdn.microsoft.com/library/default.asp?url=/library/en-us/wcetcpip/html/cmconmulticastipv6addresses.asp>

QUESTION 371

Two different islands of the Certkiller IPv6 network, which are both running ISIS (IPv6 IGP), need to connect via a tunnel over an IPv4 network. Which of the following tunneling methods could be used to achieve this goal?

- A. 6to4 Tunnels
- B. Manual Tunnels (RFC 2893)
- C. ISATAP Tunnels
- D. GRE Tunnels
- E. None of the above

Answer: D

Explanation:

IPv6 traffic can be carried over IPv4 GRE tunnels using the standard GRE tunneling technique that is designed to provide the services necessary to implement any standard point-to-point encapsulation scheme. As in IPv6 manually configured tunnels, GRE tunnels are links between two points, with a separate tunnel for each link. The tunnels are not tied to a specific passenger or transport protocol, but in this case carry IPv6 as the

passenger protocol with the GRE as the carrier protocol and IPv4 or IPv6 as the transport protocol.

The primary use of GRE tunnels is for stable connections that require regular secure communication between two edge routers or between an edge router and an end system.

The edge routers and the end systems must be dual-stack implementations.

GRE has a protocol field that identifies the passenger protocol. GRE tunnels allow Intermediate System-to-Intermediate System (IS-IS) or IPv6 to be specified as a passenger protocol, which allows both IS-IS and IPv6 traffic to run over the same tunnel. If GRE did not have a protocol field, it would be impossible to distinguish whether the tunnel was carrying IS-IS or IPv6 packets. The GRE protocol field is why it is desirable that you tunnel IS-IS and IPv6 inside GRE.

Explanation:

http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/products_configuration_guide_chapter09186a00

QUESTION 372

You need to troubleshoot an OSPF issue on the Certkiller IPv6 network. What IPv6 address would you ping to determine if OSPFv3 is able to send and receive unicast packets across a link?

- A. Link Local Address
- B. Site local Multicast
- C. Unique Local Address
- D. Anycast Address
- E. Global Address of the link
- F. None of the above

Answer: A

Explanation:

Troubleshooting OSPFv3 for IPv6 should be handled in the same way as OSPFv2 for IPv4. The major difference will be the addressing as OSPFv3 uses the link-local addresses as the source and, when sending directly to a neighbor, destination of packets.

Reference:

<http://www.networkworld.com/subnets/cisco/050107-ch9-ospfv3.html?page=9>

QUESTION 373

Router CK1 is running OSPF in an IPv6 network. What kind of information is carried within the OSPFv3 Intra-Area-Prefix LSA on this router?

- A. IPv6 prefixes and topology information
- B. Solicited mode multicast address
- C. Link-local address
- D. IPv6 Prefixes
- E. All of the above

Answer: D

Explanation:

The OSPFv3 LSA types and their OSPFv2 counterparts:

OSPFv3 LSAs		OSPFv2 LSAs	
LS Type	Name	Type	Name
0x2001	Router LSA	1	Router LSA
0x2002	Network LSA	2	Network LSA
0x2003	Inter-Area Prefix LSA	3	Network Summary LSA
0x2004	Inter-Area Router LSA	4	ASBR Summary LSA
0x4005	AS-External LSA	5	AS-External LSA
0x2006	Group Membership LSA	6	Group Membership LSA
0x2007	Type-7 LSA	7	NSSA External LSA
0x0008	Link LSA		No Corresponding LSA
0x2009	Intra-Area Prefix LSA		No Corresponding LSA

OSPFv3 removes the prefix advertisement function from Router and Network LSAs, and puts it in the new Intra-Area Prefix LS

A. Now Router and Network LSAs only represent the router's node information for SPF and are only flooded if information pertinent to the SPF algorithm changes.

Reference:

<http://www.networkworld.com/subnets/cisco/050107-ch9-ospfv3.html?page=3>

QUESTION 374

The Certkiller EIGRP network is migrating to IPv6. When using IPv6, which of the following describes the correct configuration tasks for IPv6 EIGRP? (Choose 3)

- A. EIGRP for IPv6 is directly configured on the interfaces over which it runs
- B. EIGRP for IPv6 is not configured on the interfaces over which it runs but if a user uses passive-interface configuration, EIGRP for IPv6 does need to be configured on the interface that is made passive
- C. When a user uses a passive-interface configuration, EIGRP for IPv6 does not need to be configured on the interface that is made passive

- D. There is no network statement configuration in EIGRP for IPv6
- E. There is a network statement configuration in EIGRP for IPv6 same as IPv4
- F. None of the above.

Answer: A, C, D

Explanation:

Restrictions for Implementing EIGRP for IPv6:

This section lists ways in which EIGRP for IPv6 differs from EIGRP IPv4 as well as EIGRP for IPv6 restrictions.

EIGRP for IPv6 is directly configured on the interfaces over which it runs. This feature allows EIGRP for IPv6 to be configured without the use of a global IPv6 address. There is no network statement in EIGRP for IPv6.

In per-interface configuration at system startup, if EIGRP has been configured on an interface, then the EIGRP protocol may start running before any EIGRP router mode commands have been executed.

An EIGRP for IPv6 protocol instance requires a router ID before it can start running.

EIGRP for IPv6 has a shutdown feature. The routing process should be in "no shutdown" mode in order to start running.

When a user uses passive-interface configuration, EIGRP for IPv6 does not need to be configured on the interface that is made passive.

EIGRP for IPv6 provides route filtering using the distribute-list prefix-list command. Use of the route-map command is not supported for route filtering with a distribute list.

Reference:

http://www.cisco.com/en/US/products/sw/iosswrel/ps5187/products_configuration_guide_chapter09186a00805f

QUESTION 375


What keywords do you need to the access-list to provide to the logging message like source address and source mac address?

- A. logging
- B. log
- C. log-input
- D. log-output

Answer: C

QUESTION 376

Network topology exhibit:



Certkiller 4 exhibit:
router eigrp 100
network 0.0.0.0 0.0.0.0
distribute-list 10 out

...

access-list 10 permit 192.168.1.0 0.0.0.255

You work as a network administrator at Certkiller .com. Study the exhibits carefully. For this network, assume all routers have been configured to run EIGRP in AS 100, and have also been configured to run EIGRP on all connected links. If the link between Certkiller 3 and Certkiller 4 fails, how many queries will Certkiller 5 and Certkiller 6 receive?

- A. Neither Certkiller 5 and Certkiller 6 will receive any queries for either 192.168.1.0/24 or 192.168.2.0/24.
- B. Certkiller 6 will receive two queries, one for 192.168.1.0/24 and one for 192.168.2.0/24. Certkiller 5 will receive one query, for 192.168.1.0/24.
- C. Certkiller 5 will receive one query, for 192.168.1.0/24, and Certkiller 6 will receive no queries.
- D. Both Certkiller 5 and Certkiller 6 will receive any queries for either 192.168.1.0/24 or 192.168.2.0/24.

Answer: B

QUESTION 377

An OSPF adjacency will not form correctly across a point-to-point link in the same area. What is the most likely reason for this problem?

- A. Each interface has different MTU size.
- B. Each interface is configured with the ip unnumbered loopback 0 command.
- C. Each interface is configured with secondary addresses as well as primary addresses.
- D. Each interface has a different OSPF cost.


Answer: A

Explanation:

The question states "OSPF adjacency" ... an incorrect MTU size will cause this problem ... this is the "best" answer. Secondary addresses is more a concern for EIGRP not OSPF.

QUESTION 378

Network topology exhibit:



You work as a network administrator at Certkiller .com. Study the exhibit carefully. Both multicast sources are sending to the same multicast address. How does a user specify which multicast stream they would like to receive?

- A. The multicast streams must be separated from each other by specifying a scope for each. This means that each user (Certkiller A and Certkiller B) can only get multicast traffic from one of the sources.
- B. Dense mode is used to forward the multicast streams to the end users, Certkiller A and Certkiller B, allowing them to pick up the desired stream when it arrives.
- C. Routers Certkiller 1 and Certkiller 2 are set up as rendezvous points. The user joins a multicast group by sending an IGMP request to their local router. The local router then connects to the correct RP and receives the desired multicast stream.
- D. The user must know the source address and group address of the desired multicast stream and explicitly join that stream.

Answer: D

Explanation:

The user must know the source address and group address of the desired stream and explicitly join it.

This is a feature of igmpv3. the group add is 232.1.1.1 and this is in the ssm range specified for igmpv3. this is a security behaviour incorporated to mitigate dos that can be caused whn rogue multicast group add diverts stream from clients. this is the case with version 1 and 2.

QUESTION 379

Which of these statements about OSPF external LSAs (type 5) is correct?

- A. External Network LSAs (type 5) redistributed from other routing protocols into OSPF

are not permitted to flood into a stub area.

B. External LSAs (type 5) are automatically changed to type 1 LSAs at ASBRs.

C. OSPF external LSAs can be flooded into an NSSA area if redistributed from other routing protocols into OSPF and if the subnets parameter is used with the redistributed command.

D. OSPF external LSAs are automatically flooded into all OSPF areas, unlike type 7 LSAs, which require that redistribution be configured.

Answer: A

Explanation: type 5 not flooded into stub area

QUESTION 380

If you have overlapping IP address between two different networks or routing domains, what two commands do you need to globally configure NAT to get this to work?

A.

ip nat outside source static tcp x.x.x.x.y.y.y and ip nat outside source tcp x.x.x.x.y.y.y.y.

B. ip nat outside source list 1 interface x and ip nat inside source list 1 interface x

C. ip nat outside source static udp x.x.x.x.y.y.y and ip nat outside source udp x.x.x.x.y.y.y.y.

D. ip nat outside source static x.x.x.x.y.y.y and ip nat outside source x.x.x.x.y.y.y.y.

Answer: B

Explanation: Both source and destination address has to be translated for this solution.

QUESTION 381

If you have multiple DHCP pools configured on the same router, what IOS command has to be entered in the DHCP configuration to be processed be using the other DHCP pool configuration?

A. default-.gateway

B. ip helper

C. host

D. network

Answer: C

Explanation: If you want to use the other DHCP pools you should do it statically.

Not D: The question is: << If you have multiple DHCP pools configured in the same Router>> That's mean you don't need the network to configure pools, because they are already configured

QUESTION 382

Network topology exhibit:

```
interface Serial0
 ip address 172.16.47.161 255.255.255.240
 ip nat inside
!
interface Serial1
 ip address 172.16.47.146 255.255.255.240
 ip nat inside
!
no ip classless
ip route 0.0.0.0 0.0.0.0 172.16.47.145
ip route 171.68.200.0 255.255.255.0 172.16.47.162

TestKing4 #show ip nat translation
Pro Inside global      Inside local      Outside local      Outside global
--- 172.16.47.150      171.68.200.48      ---                ---

TestKing4 #debug ip packet detail
TestKing4 #debug ip nat
TestKing4 #IP: s=171.68.200.48 (Serial0), d=172.16.47.142, len 100, unroutable
ICMP type=8, code=0
TestKing4 #IP: s=172.16.47.161 (local), d=171.68.200.48 (Serial0), len 56, sending
ICMP type=3, code=1
TestKing4 #IP: s=171.68.200.48 (Serial0), d=172.16.47.142, len 100, unroutable
ICMP type=8, code=0
TestKing4 #IP: s=171.68.200.48 (Serial0), d=172.16.47.142, len 100, unroutable
ICMP type=8, code=0
TestKing4 #IP: s=172.16.47.161 (local), d=171.68.200.48 (Serial0), len 56, sending
ICMP type=3, code=1
TestKing4 #IP: s=171.68.200.48 (Serial0), d=172.16.47.142, len 100, unroutable
ICMP type=8, code=0
TestKing4 #IP: s=171.68.200.48 (Serial0), d=172.16.47.142, len 100, unroutable
ICMP type=8, code=0
TestKing4 #IP: s=172.16.47.161 (local), d=171.68.200.48 (Serial0), len 56, sending
ICMP type=3, code=1
```

You work as a network administrator at Certkiller .com. Study the exhibit carefully. The exhibit shows the NAT configuration on Certkiller 4 and the output for a ping issued from device 171.68.200.48 and destined to 172.16.47.142. Based on this information, what change must be made on Certkiller 4 in order for the ping to work?

- A. load a newer IOS image
- B. reload the router
- C. add a static route
- D. configure IP as classless
- E. clear the route cache

Answer: D

QUESTION 383

Which three LMI types can be configured for use with Frame Relay on a Cisco router? Select three.

- A. Q.931 - Annex B
- B. Cisco
- C. ANSI - Annex D
- D. Q.933 - Annex A

Answer: B, C, D

Explanation:

The three LMI types for use with Frame Relay on a CISCO Router are:

CISCO

ANSI

Q.933

QUESTION 384

Frame Relay traffic shaping is enabled on a WAN interface with the following settings:

CIR=768 kb/s

Bc=2000

Be=7680

What is the time interval Tc?

A. 12.6 ms

B. 2.6 ms

C. 7.4 ms

D. 10 ms