

**QUESTION 201**

The Certkiller network is using EIGRP as the routing protocol, and the EIGRP topology information for router CK4 is displayed below:

CK4 #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 28009856

via 116.16.34.3 (28009856/2297856), Serial1/0

P 190.190.0.0/16, 1 successors, FD is 2297856

via 116.16.34.2 (2297856/128256), Serial1/0

P 130.130.0.0/16, 1 successors, FD is 2297856

via 116.16.34.3 (2297856/129256), 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 116.16.34.0/28, 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 116.16.34.2 are EIGRP neighbors to CK4 .
- B. The 116.16.37.0 network is reachable via the 116.16.34.2 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 202**

The CertK king 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 203**

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 209.224.111.4 255.255.255.252 secondary
!
router eigrp 200
 network 192.168.1.0
 network 209.224.111.0

Hostname Certkiller2
!
interface Ethernet0/0
 bandwidth 1000
 ip address 209.224.111.62 255.255.255.252
!
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 204

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

## 350-001

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 205**

Which statement is true regarding IS-IS?

- A. Within one area of ISIS, all links must have ip router isis configured. At least one link in that area must have this command, if the IP routing is to work correctly.
- B. In ISIS the pseudonode is the router that is responsible for originating the LSP for the LAN.
- C. If a new router boots on the LAN with a higher interface priority, the new router waits for the active DIS to go down before it becomes the DIS.
- D. You do not need to configure a NET on the router to run ISIS for IP.

Answer: A

Explanation:

All interfaces need to have the "ip router isis" command configured under each interface

that is to participate in ISIS routing.

Incorrect Answers:

B. The Designated Intermediate System is responsible for the LAN LSP, not the pseudonode

C. The DIS election is preemptive (unlike OSPF). If a new router boots on the LAN with a higher interface priority, the new router becomes the DIS.

D. A NET is required for all IS-IS configurations.

Reference:

<http://www.cisco.com/warp/public/97/DIS-LSP-1.html>

---

**QUESTION 206**

What are the states that ISIS neighbors go through during the process of building adjacencies?

A. Init, Up

B. Init, 2-way, Full

C. Loading, Start, Synchronizing, Up

D. Init, 2-way, Exstart, Exchange, Up

E. Adjacency formation will be automatic

Answer: A

Explanation:

Init and UP are the only two ISIS states.

---

**QUESTION 207**

The Certkiller network is running IS-IS as the routing protocol. Router CK1 is a non-pseudonode router. How many LSPs will originate from router CK1 ?

A. 1 per link and 1 per external route.

B. 1 for L1 routers and 2 for L2 routers

C. Always 1.

D. Sometimes 1.

E. Always none.

Answer: C

Explanation:

Non-pseudonode routers always produce only one LSP.

---

**QUESTION 208**

2 ISIS routers are each configured with the default priority. Which router will be elected as the Designated Intermediate System (DIS)?

A. The router with the highest router-ID.

B. The router with the lowest system-ID.

- C. The router with the lowest router-ID.
- D. The router with the highest system-ID.
- E. The router with the highest SNPA.

Answer: E

Explanation:

The default L1 and L2 priority is 64 in Cisco routers, with a value range of 0-127. A value of 0 means the router is ineligible to become the DIS. The router with the highest priority is chosen as the DIS, and in the case of a tie, the router with the numerically highest SNPA becomes the DIS.

DIS is the highest Subnetwork Points of Attachment (SNPA) address on the LAN segment. The SNPA address refers to the data link address, and in this case is the MAC address.

On a LAN, one of the routers elects itself the DIS, based on interface priority (the default is 64). If all interface priorities are the same, the router with the highest subnetwork point of attachment (SNPA) is selected. The SNPA is the MAC address on a LAN, and the local data link connection identifier (DLCI) on a Frame Relay network. If the SNPA is a DLCI and is the same at both sides of a link, the router with the higher system ID becomes the DIS. Every IS-IS router interface is assigned both a L1 priority and a L2 priority in the range of 0 to 127.

The DIS election is preemptive (unlike OSPF). If a new router boots on the LAN with a higher interface priority, the new router becomes the DIS. It purges the old pseudonode LSP and floods a new set of LSPs.

Incorrect Answers:

A, B, C, D. ISIS routers do not use router ID's or system ID information.

Reference:[http://www.cisco.com/en/US/tech/CK365/technologies\\_tech\\_note09186a0080094b42.shtml](http://www.cisco.com/en/US/tech/CK365/technologies_tech_note09186a0080094b42.shtml)

---

### **QUESTION** 209

The Certkiller network uses ISIS as its routing protocol. You notice periodic CSNP and PSNP packets going across the network. What are the PSNP and CSNP packets used for?

- A. PSNP are used to acknowledge the receipt or to request the retransmission of the latest version of an LSP while the CSNP are used for synchronizing the LS Database on adjacent neighbors.
- B. CSNP are used to acknowledge the receipt or to request the retransmission of the latest version of an LSP while the PSNP are used for synchronizing the LS Database on adjacent neighbors.
- C. PSNP are used to acknowledge the receipt of the latest version of an LSP while the CSNP are used to synchronize the LS Database of adjacent neighbors or to request the retransmission of an LSP.
- D. CSNP are used to acknowledge the receipt of the latest version of an LSP while the PSNP are used to synchronize the LS Database of adjacent neighbors or to request the retransmission of an LSP.

Answer: A

Explanation:

CSNP (Complete Sequence Number PDU) is sent by the DR to maintain DB synchronization. PSNP (Partial Sequence Number PDU) are used to acknowledge or request one or more LSPs.

---

**QUESTION 210**

You have configured all routers in your Certkiller network for IS-IS. You want to verify that the correct paths to all your Intermediate systems are being used. Which command should you use?

- A. show isis topology
- B. show ip route isis topology
- C. show ip route isis
- D. show clns neighbor
- E. show ip route

Answer: A

Explanation:

Show ISIS topology is the command used to view the ISIS paths to Intermediate Systems as shown below:

```
Certkiller #show isis ?
database IS-IS link state database
hostname IS-IS Dynamic hostname mapping
mesh-groups IS-IS mesh groups
mpls IS-IS MPLS
nsf Non-stop forwarding state
route IS-IS level-1 routing table
spf-log IS-IS SPF log
topology IS-IS paths to Intermediate Systems
```

```
Certkiller #show isis topology
```

Cisco definitions

```
show isis topology
```

To display a list of all connected routers in all areas, use the show isis topology command in user EXEC or privileged EXEC mode.

Incorrect Answers:

D. show clns neighbors is used to display end system (ES), intermediate system (IS), and multitopology Integrated Intermediate System-to-Intermediate System (M-ISIS) neighbors.

---

**QUESTION 211**

What answer best describes the concept of the backbone in integrated IS-IS?

- A. Integrated IS-IS uses L3 routers for backbone inter-domain routing.
- B. The backbone is a contiguous chain of L2 or L1/L2 IS-IS routers.
- C. Since integrated IS-IS routers belong to at most one area, no centralized backbone area is required.
- D. The IS-IS backbone is a centralized area to which all other IS-IS areas must be connected.
- E. Integrated IS-IS assumes a single domain and thus no backbone is required.

Answer: B

Explanation:

L2 routers are analogous to OSPF backbone routers, and the set of L2 routers (including L1/L2 routers) and their interconnecting links make up the IS-IS backbone, similar to area 0 in an OSPF network.

Reference: Jeff Doyle, "Routing TCP/IP, Volume 1" page 598-599.

Incorrect Answers:

- A. IS-IS does not use level 3 routers.
- C. Some routers, called L1/L2 routers, belong to both area types.
- D. Unlike OSPF, IS-IS routers are not required to be connected to a contiguous backbone area. In fact, the backbone area can also be segmented in IS-IS.
- E. IS-IS uses the concepts of router levels, which is similar to OSPF areas. L2 routers are analogous to OSPF backbone routers, and L1/L2 routers are analogous to OSPF ABRs.

---

**QUESTION 212**

Given the following integrated IS-IS NSAP address:

47.1052.cb01.000c.3ab6.029c.00

To what area will a router configured with this NSAP address belong?

- A. 47
- B. 47.1052
- C. 47.1052.cb01
- D. 1052.cb01
- E. 000c.3ab6.029c

Answer: D

Explanation:

Any NSAP address contains the following items:

- \* AFI - Authority Format Identifier - This is a 1 byte (two hex character) field to denote the 'type' of NSAP. The AFI information is used to determine how the Area ID is to be interpreted and its formatting.
- \* Area ID - This is the part that generates much of the confusion since the value is listed as 2-12 bytes in length. So the area ID can be anywhere from 4 to 24 hex characters.
- \* System ID or ESI-Address is the next field. This must be a 6-byte identifier that is used



to uniquely denote a system within an area. Many times we see MAC addresses used (since they should be unique), other times we see IP addresses.

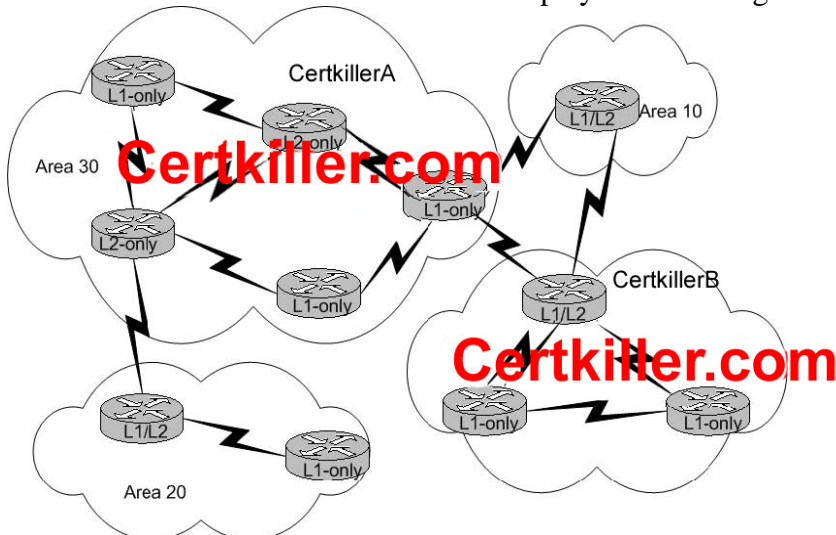
\* N-Selector - For lack of a better analogy, this is like a port number in TCP and UDP. If you have multiple NSAPs on a device pointing to different applications/protocols/insertion points, then the N-Selector is used to vary between them and save your CPU cycles a bit. Typically this will be .00, at least when we only have one NSAP.

Since the Area ID is a variable length value, the problem becomes determining where the area ID ends and the system ID part of the NSAP address begins. When presented with an NSAP address, the AFI may tell you things to begin with. But beyond that, look at the known quantities: The AFI is always one byte, the system ID is always six bytes and the N-Selector is always one byte. Therefore, counting off seven bytes from the right-hand side and not counting the initial byte, anything leftover is the area ID. And this value should be between two and twelve bytes in length. Therefore, in this case the area ID is 1052.cb01

---

**QUESTION 213**

The Certkiller IS-IS routed network is displayed in the diagram below:



Which type of IS-IS adjacency is formed between Certkiller A and Certkiller B?

- A. L1 adjacency only
- B. L2 adjacency only
- C. Both L1 and L2 adjacency
- D. No adjacencies will be formed

Answer: D

Explanation:

Adjacency Building In IS-IS

Neighbors on point-to-point networks always become adjacent unless they do not see themselves in their neighbors' hello PDU and match on certain parameters. On broadcast networks and nonbroadcast multiaccess (NBMA) networks, the DIS (Designated

Intermediate System) will become adjacent with its neighbors.

Two routers will become neighbors if the following parameters are agreed upon:

- Level 1-The two routers sharing a common network segment must have their interfaces configured to be in the same area if they are to have a Level 1 adjacency.
- Level 2-The two routers sharing a common network segment must be configured as Level 2 if they are in different areas and want to become neighbors.
- Authentication-IS-IS allows for configuration of a password for a specified link, for an area, or for an entire domain.

Incorrect Answers:

A, C: L1 neighbor adjacencies can only be formed by routers within the same area.

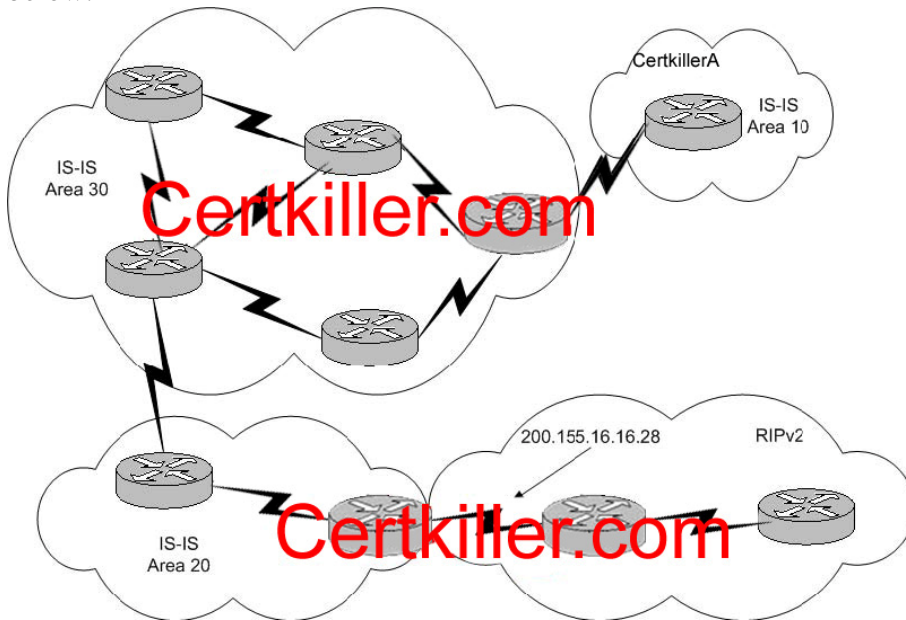
Even though both routers are configured for L1 routing, they are in different areas and so router Certkiller A will need to go through a L2 router to reach the area in Certkiller B.

B, C: L2 neighbor adjacencies can only be formed if both routers are configured for L2 routing. In this case Certkiller A is only configured for L1 routing.

---

**QUESTION 214**

The Certkiller network is using IS-IS and RIP version 2 as displayed in the diagram below:



Network 200.155.16.16/28 is being redistributed from RIPv2 into IS-IS with a metric of 50 and a metric-type Internal. Given no modifications have been made to IS-IS metrics in the network, what is the total cost of the path from Certkiller A to network 200.155.16.16?

- A. 50
- B. 55
- C. 100
- D. 370

Answer: C

Explanation:

The default IS-IS metric type is 10 for all links, regardless of the interface type. Since there are 5 links between the 200.155.16.16/28 network and router Certkiller A, the metric is 50. When we add the redistributed route, which came from RIP version 2 with a metric of 50, the total metric is  $50+50 = 100$ .

---

**QUESTION 215**

You are tasked with comparing the functionality of IS-IS and OSPF to determine the routing protocol to use with the Certkiller network. Which of the following is true regarding the IS-IS and OSPF routing protocols? (Select two)

- A. Integrated IS-IS uses 3 levels of hierarchy where OSPF only supports 2.
- B. IS-IS allows for easier backbone extension as the backbone does not have to be a centralized area.
- C. OSPF area addresses are fixed at 4 bytes where IS-IS uses an area address fixed at 2 bytes encoded in the NET.
- D. IS-IS is slightly more efficient as IP-only changes trigger partial route calculations instead of a full SPF recalculation.
- E. IS-IS and OSPF are both link state protocols, but IS-IS uses a modified version of the Dijkstra algorithm to improve its efficiency.

Answer: B, D

Explanation:

With IS-IS, there is no restriction that all backbone routers (level 2 routers) be contiguous such as the backbone area of OSPF. In OSPF all areas must be directly linked to area 0, and the backbone area must also not be segmented. With IS-IS, the backbone area can be more easily extended since all L2 routers need not be linked directly together.

With regard to CPU use and the processing of routing updates, IS-IS is more efficient than OSPF. In IS-IS, one LSP is sent per IS-IS router in each area (including redistributed prefixes [routes]), compared to the many OSPF LSAs that would be sent. Not only are there fewer LSPs to process, but the mechanism by which IS-IS installs and withdraws prefixes is less processor intensive. In IS-IS, the entire SPF table is not refreshed periodically like OSPF, which does so every 30 minutes by default.

Incorrect Answers:

- A. IS-IS routers only use two levels of hierarchy; level 1 or level two. Note that a router can act as both (L1/L2) but there still exists only 2 different types.
  - C: The IS-IS area field is a variably lengthed value. In addition, OSPF areas are not fixed at 4 bytes, since they can also have single digit values for the area (such as area 0 or area 1, etc).
  - E: Although IS-IS is slightly more efficient than OSPF, the Dijkstra algorithm remains the same for both protocols.
-

**QUESTION 216**

The Certkiller network is using IS-IS for the interior routing protocol. Which of the following correctly describe the characteristics of level-1 routing in this IS-IS network? (Select two)

- A. Level-1 routers maintain a single copy of the link state database.
- B. Level-1 intermediate systems are responsible for intra-area routing.
- C. Level-1 routing in integrated IS-IS is based on the area address of contained in the NET.
- D. An ES may be adjacent to a level-1 router only if they both share a common Area Address.
- E. Level-1 routers uses Intermediate System Hello (ISH) packets to form adjacencies with other IS-IS level-1 routers in the same area.

Answer: B, E.

Explanation:

Router performs only Level 1 (intra-area) routing. This router only learns about destinations inside its area. Level 2 (interarea) routing is performed by the closest Level 1-2 router.

IS routers use hello PDUs to establish and maintain adjacencies. IS-IS forms separate adjacencies for L1 and L2 neighbors, so a L1 router will never become adjacent to a L2 router. By default, these hello packets are exchanged every 10 seconds.

Incorrect Answers:

- A. All IS-IS routers maintain both a CLNS and an IP routing link state database.
- C. Level 1 routers use the information contained in the system ID of the NET, while level 2 routers use the area ID information.
- D. An ES will become adjacent with a L1 router if their respective capabilities and interface parameters match. Information regarding the System ID is used in L1 routers, not the area address ID.

---

**QUESTION 217**

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:

- Classify traffic based on extended access list criteria. Access lists, then, establish the match criteria.
- Set IP Precedence bits, giving the network the ability to enable differentiated classes of service.
- 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 218**

The router CK1 is being configured to filter BGP routes. In a BGP peering relationship with a customer where routing information is exchanged, which prefix list filter(s) will ensure that only class-B address space networks are accepted by the 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 ge 16  
ip prefix-list list-D permit 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 be match. Therefore, only addresses that fall in the class B range will pass through the filter.

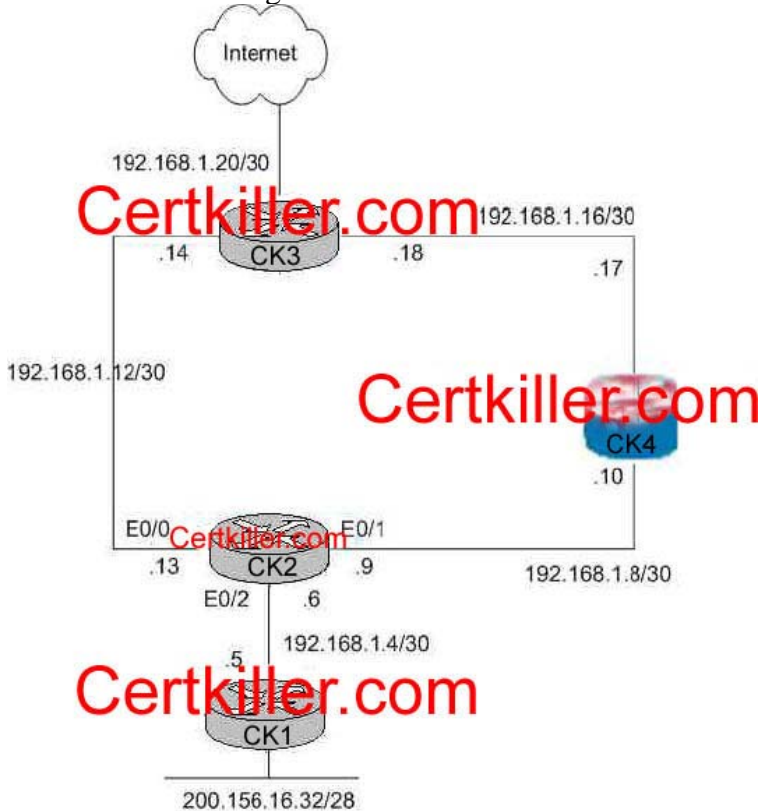
Incorrect Answers:

- 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 219**

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 200.155.16.32 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/CK3 65/technologies\\_tech\\_note09186a008009481d.shtml#configforfire](http://www.cisco.com/en/US/partner/tech/CK3_65/technologies_tech_note09186a008009481d.shtml#configforfire)

---

**QUESTION 220**

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 policy route acl mark-em-up
|
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 172.168.1.55, 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 221

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 "passiveinterface 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 222**

You have an ISDN backup link installed to be used only as a fail over should your primary link go down. You want to use a static route to accomplish this. The rest of you network uses a dynamic routing protocol. How should you configure the static route?

- A. Configure the static route with a lower administrative distance than the dynamic protocol.
- B. Configure the static route with a higher administrative distance than the dynamic protocol.
- C. Configure the static route with a lower metric than the dynamic protocol.
- D. Configure the static route with the floating-static keyword.
- E. Configure the static route with the default administrative distance.

Answer: B

Explanation:

With a higher administrative distance, the dynamic routing protocol will always be the preferred route. This is normally known as a floating static route.

Incorrect Answers:

- A. The route with the lowest administrative distance will be the preferred path.
- C. Metrics are only compared for tie-breaking routes within the same protocol, not from different protocols.
- D. There is no such command. Although we are indeed trying to configure a floating static route, this is not the command used.
- E. By default, the AD of a static route is 1, which will be lower than the dynamic routing protocol's AD.

Reference:

<http://www.cisco.com/univercd/cc/td/doc/cisintwk/ics/cs002.htm#xtocid1902134>

---

### **QUESTION 223**

Two routers are connected via a Frame Relay link, with ISDN used as a backup. OSPF is configured on both routers, with OSPF demand circuit is configured on the ISDN interfaces. Which of the following is true?

- A. The calling router's ISDN interface must be network type non-broadcast.
- B. OSPF demand circuit will not trigger the link if an OSPF interface goes down.
- C. The calling router's ISDN interface must be network type point-to-point.
- D. OSPF demand circuit will bring up the link if the topology of the network changes.

Answer: D

Explanation:

An ISDN circuit interface can be either point-to-point or point-to-multipoint. Using the OSPF demand circuit options, which suppresses Hello and LSA refresh functions, OSPF can establish a demand link to form an adjacency and perform initial database synchronization. The adjacency remains active even after Layer 2 of the demand circuit goes down.

Unlike the OSPF demand circuit feature, flooding reduction is usually configured on leased lines. Flooding reduction uses same technique as demand circuits to suppress the periodic LSA refresh.

When an OSPF demand circuit is configured on a link, the periodic OSPF Hellos are suppressed. Periodic Hellos are suppressed only on a point-to-point and point-to-multipoint network type. On any other network type, OSPF Hellos are still sent over the interface.

There are only two scenarios where the periodic LSA refresh occurs when using the OSPF demand circuit feature:

- If there is a change in network topology
- If there is a router in the OSPF domain that can not understand demand circuits

In the first case, not much can be done to stop the LSA refresh because the router has to send the new LSA information to update the neighbor about the topology change.

For additional information on OSPF demand circuit see:

[http://www.cisco.com/en/US/tech/CK365/CK480/technologies\\_tech\\_note09186a0080094a8f.shtml](http://www.cisco.com/en/US/tech/CK365/CK480/technologies_tech_note09186a0080094a8f.shtml)

Incorrect Answers:

- A. The OSPF demand circuit interface must be either point to point or point to multipoint. When an OSPF demand circuit is configured on a link, the periodic OSPF hellos are suppressed. Periodic hellos are suppressed only on a point-to-point and point-to-multipoint network type. On any other network type, OSPF hellos are still sent over the interface.
- B. When an OSPF interface goes down, the topology of the OSPF network changes, which will indeed trigger the ISDN call.
- C. Although the interface can indeed be point to point, it could also be point to multipoint. Therefore, the interface does not necessarily need to be configured as point to point.

---

**QUESTION 224**

You are the network administrator at Certkiller . On the Certkiller network there is a point-to-point ISDN link routers CK1 and CK2 . Because toll charges are cheaper where CK1 resides you want CK1 to be able to dial CK2 , but not vice versa. What can you use to prevent CK2 from ever dialing to CK1 ?

- A. Use an IP access-list with the access-group command on the interface.
- B. Remove the dial string from CK2 .
- C. Restrict the amount of interesting traffic on CK2 .
- D. Use the no isdn-dial keyword on the ISDN interface of CK2 .
- E. No need to take action.

Answer: B

Explanation:

Without a dial-string, the router will never be able to initiate a call, since the router would not be aware of the number to call even if there is interesting traffic.

Incorrect Answers:

- A. Using an access list would restrict the traffic going from CK2 to CK1 , but an ISDN call could still be initiated.
- C. Restrict the amount of interesting traffic would not be enough. We would have to define all traffic as un-interesting traffic as an alternative way to stop CK2 from dialing CK1 .
- D. This is an invalid command.

---

**QUESTION 225**

In a DDR implementation using an ISDN BRI circuit, you wish to utilize both B channels simultaneously. When an ISDN call is triggered, a connection can be made but only one of the B channels appears to be utilized. Which are likely causes of this behavior? (Choose two)

- A. PPP authentication using CHAP has been improperly configured.
- B. The dialer-load threshold command is missing or improperly configured.

## 350-001

- C. Interesting traffic has not been properly identified using the dialer-list command.
- D. The load-interval has not been properly configured on the dialing BRI interface.
- E. The BRI interfaces have not been configured for multilink PPP.

Answer: B, E

Explanation:

In order for both ISDN B channels to be utilized at the same time, the encapsulation type must be set to PPP so that multilink PPP can be used. Multilink PPP provides a means to bond multiple physical channels into one larger virtual channel.

The dialer load-threshold command sets a load that will trigger a concurrent call to be placed on the second B channel. Multilink Point-to-Point Protocol (PPP) has been configured (ppp multilink) so that both ISDN B channels can be bundled together as one Virtual-Access interface for aggregate bandwidth.

Reference: [www.cisco.com/warp/public/125/fr\\_isdn\\_backup.html](http://www.cisco.com/warp/public/125/fr_isdn_backup.html)

Incorrect Answers:

- A. If the CHAP authentication was improperly configured, the ISDN connection would fail completely. In this case, the initial ISDN connection succeeds using only one of the B channels.
- C. Again, if this were the case, then the ISDN circuit would have difficulty connecting at all. We need to get the second B channel to bond to the already working B channel in this specific case.
- D. The dialer load-interval command is used to adjust the frequency that the load of the interface is calculated. By default, the interface load is calculated as an exponential average over the last 5 minutes. Adjusting this timer could indeed cause the second B channel to be brought up in rare circumstances, but it would not be a likely cause for the second B channel not being utilized.

---

### **QUESTION** 226

Routers CK1 and CK2 are in the same LAN, and both are running RIP version 2.

During a troubleshooting session you place a sniffer on the LAN network. Using the sniffer you see routers CK1 and CK2 sending routing updates to each other every 30 seconds. Which IP address should you expect to see these updates destined to?

(Choose all that apply)

- A. 224.0.0.10
- B. 255.255.255.255
- C. 224.0.0.13
- D. 224.0.0.5
- E. 224.0.0.9
- F. 224.0.0.6

Answer: E

Explanation:

RIPV2 sends periodic route updates sent every 30 seconds to multicast address 224.0.0.9.

Incorrect Answers:

- A. 224.0.0.10 is used by IGRP
- D, F. These are the multicast addresses used by OSPF

Reference:

CCIE Routing and Switching Exam Certification Guide Page 338

---

**QUESTION 227**

What is the destination IP address of routing update packets used by RIPv2?

What would your reply be?

- A. 224.0.0.1
- B. 224.0.0.10
- C. 224.0.0.5
- D. 224.0.0.9
- E. 255.255.255.255

Answer: D

Explanation:

224.0.0.9 is the RIP v2s multicast address.

Incorrect Answers:

- A. This is the multicast address destined for all hosts on the subnet.
  - B. This is the multicast address used by IGRP.
  - C. This address is used by OSPF.
  - E. This is the all hosts broadcast address.
- 

**QUESTION 228**

The router CK1 is using RIPv2 as the routing protocol, and the partial configuration file is displayed below:

```
interface Ethernet 1
ip address 10.1.1.1 255.255.255.0
ip summary-address rip 10.2.0.0 255.255.0.0
ip split-horizon
!
router rip
network 10.0.0.0
```

What is a result of the configuration shown for router CK1 ?

- A. The 10.2.0.0 network overrides the auto summary address of 10.1.1.1.
- B. The 10.2.0.0 network is advertised out interface E1, and the auto summary address is not advertised.
- C. The auto summary address of 10.1.1.1 will be advertised out interface E1 and the interface summary-address is not advertised.

- D. Neither the auto summary address nor the interface summary-address is advertised because split horizon is enabled.
- E. Both the auto summary address and the interface summary-address are advertised out of interface E1.

Answer: D

Explanation:

If split horizon is enabled, neither auto-summary nor interface summary addresses (those configured with the ip summary-address rip command) are advertised. The split horizon mechanism blocks information about routes from being advertised by a router out of any interface from which that information originated. It is enabled on all interfaces by default.

Reference:[http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps1830/products\\_feature\\_guide09186a0080087ad1.html#xtocid2483014](http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps1830/products_feature_guide09186a0080087ad1.html#xtocid2483014)

---

**QUESTION 229**

A customer has a frame-relay network with 2 sites - a headquarters site and a remote site - with a single PVC connecting the 2 sites. The network is running RIP version II. The company 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. The customer 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.

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 230**

A RIP Version 2 router is sending RIP updates to its neighbor that include several contiguous IP subnet routes in the 10.1.1.0/24 space. What command should be configured to aggregate the routes into a single route in the update to the RIP neighbor?

- A. summary-address rip 10.1.1.0 255.255.255.0, configured under the RIP process or the interface
- B. summary-address 10.1.1.0 255.255.255.0, configured under the RIP process
- C. ip summary-address rip 10.1.1.0 255.255.255.0, configured under the interface
- D. rip summary-address 10.1.1.0 255.255.255.0, configured under the interface
- E. ip rip summary-address 10.1.1.0 255.255.255.0, configured under the interface
- F. None of the above

Answer: C

Explanation:

The "ip summary-address rip" command causes the router to summarize a given set of routes learned via RIP version 2 or redistributed into RIP version 2. Host routes are especially applicable for summarization. To configure IP summary addressing, use the following commands in global configuration mode:

	Command	Purpose
<b>Step 1</b>	Router(config)# interface ethernet1	Enters interface configuration mode.
<b>Step 2</b>	Router(config-if)# ip summary-address rip <i>ip_address ip_network_mask</i>	Specifies the IP address and network mask that identify the routes to be summarized.

Example:

```
router rip
network 10.0.0.0
exit
interface ethernet1
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
exit
```

Reference:[http://www.cisco.com/en/US/products/sw/iosswrel/ps1831/products\\_configuration\\_guide\\_chapter09186a00800d97f7.html#52880](http://www.cisco.com/en/US/products/sw/iosswrel/ps1831/products_configuration_guide_chapter09186a00800d97f7.html#52880)

---

**QUESTION 231**

What is the default seed metric for routes redistributed into RIPv2?



- A. 1
- B. 15
- C. 16
- D. 120
- E. Infinity

Answer: A

Explanation:

When routes are redistributed into RIP, the default metric applied to the route is 1. Because RIP (both version 1 and version 2) uses hop count as the metric, the routes will be viewed as being 1 hop away.

Note: When connected routes are redistributed into RIP, the default seed metric is 0.

Incorrect Answers:

B, C. Since RIP considers a route with a hop count of 16 as unreachable (infinity) using these values as the default metric will make all routes unreachable when advertised to RIP peer routers.

D. This is the default AD of RIP.

E. This would make all routes unreachable. Note that choices C and E are effectively the same answer.

---

**QUESTION 232**

Router CK1 is running RIP Version II and has 2 interfaces. CK1 has received RIP routing updates from its neighbors on both interfaces. The first interface receives a routing update for network 10.1.1.0/24 with a metric of 3 while the second interface also receives a routing update for network 10.1.1.0/24 with a metric of 5. Which interface(s) will router CK1 select to forward packets to network 10.1.1.0/24?

- A. The router will choose the first interface because it has the lowest metric.
- B. The router will load share across both interfaces in a weighted fashion, sending the first 3 packets out of the first interface, and the next 5 packets out of the second interface.
- C. The router will choose the second interface because it has the highest metric.
- D. The router will equally load share packets across both interfaces in a round robin fashion, because both are valid RIP Version II routes.
- E. The router will ignore the RIP metrics and compare the administrative distance of each route, and choose the interface with the lowest administrative distance.

Answer: A

Explanation:

Although load sharing occurs when there are equal cost paths to a destination, it does not occur when the paths are not equal. Sometimes the router must select a route from among many learned via the same routing process with the same administrative distance. In this case, the router chooses the path with the lowest cost (or metric) to the destination.

RIP version 2 uses hop count as the metric, like version 1. In this case, only the path with a metric of 3 will be chosen over the path with a metric of 5 hops.

**QUESTION 233**

The relevant configuration and ip route information on router CE11A is displayed below:

```
! CE11A 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

CE11A#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 configuration and the show ip route connected output shown in the exhibit, which RIPv2 updates will be sent out of the Serial2/0.101 sub interface from router CE11A? (Select all that apply)

- A. 192.168.1.0/24
- B. 192.168.1.80/30
- C. 192.168.1.84/30
- D. 192.168.1.80/32
- E. 192.168.1.81/32
- F. 192.168.1.82/32
- G. 192.168.1.83/32
- H. 192.168.1.84/32

Answer: B, H

**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 234**

Part of the configuration file for router Certkiller 1 is displayed 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
```

**Certkiller.com**

You have configured RIPv2 summarization on Certkiller 1 interface Ethernet 0/0 but the routes are still not being summarized. Looking at the Certkiller 1 partial configuration above, what could be causing the problem?

- A. You need also to enable the auto summarization under the RIP process.
- B. You need also to disable the auto summarization under the RIP process.
- C. RIP does not support summarization on interface basis.
- D. Split horizon is enabled on an interface basis.
- E. The mask configured on the "ip summary-address" command must be /24 bits.

Answer: D

Explanation:

Cisco routers can summarize routes in two ways:

- Automatically, by summarizing subprefixes to the classful network boundary when crossing classful network boundaries (autosummary)
- 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:

- 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).
- 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 Certkiller 1, 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. Supernet 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\\_chapter09186a00800d97f7.html#1001151](http://www.cisco.com/en/US/products/sw/iosswrel/ps1831/products_configuration_guide_chapter09186a00800d97f7.html#1001151)

---

### QUESTION 235

To display the routing table of router CK1 , the "show ip route" command was issued. Router CK1 is running OSPF. Which of the following statements are 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.

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 236

While troubleshooting an issue with a serial interface on the Certkiller 4 router, you issue the "show interface" command as shown below:

```
Certkiller4#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 (10 sec)
LMI enq sent 154467, LMI stat rcvd 154468, LMI upd rcvd 0, DTE LMI up
LMI enq rcvd 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: weighted fair
Output queue: 0/1000/64/0 (size/max total/threshold/drops)
  Conversations 0/1/256 (active/max active/max data)
  Reserved Conversations 0/0 (all thresholds 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 runts, 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 collisions, 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
```

Certkiller.com

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](http://www.cisco.com/univercd/cc/td/doc/cisintwk/itg_v1/tr1915.htm#xtocid7)

---

**QUESTION 237**

The show ip bgp regexp [regexp] command is most useful when performing which type of BGP troubleshooting?

- A. To verify and troubleshoot BGP Prefix-list filtering configurations.
- B. To verify and troubleshoot BGP AS-Path filtering configurations.
- C. To verify and troubleshoot BGP route-maps configurations.
- D. To verify and troubleshoot BGP synchronization problems.
- E. To verify and troubleshoot BGP AS-path prepending configurations.

Answer: B

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 238**

You want traffic on your frame relay link to conform to specific policies. Because of this, you configure traffic shaping 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
ip address 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 239

What are the primary reasons to implement traffic shaping on a 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.

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](http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4224/sw_config/traffic.htm)

---

**QUESTION 240**

You have set up priority queuing on the serial interface of your router as follows:

```
priority-list 1 protocol ip high list 101
priority-list 1 protocol ip medium list 102
priority-list 1 protocol ip normal list 103
priority-list 1 protocol ip low list 104
priority-list 1 default low
access-list 101 permit ip any any precedence critical
access-list 102 permit ip any any precedence flash
access-list 103 permit ip any any precedence priority
access-list 104 permit ip any any precedence network
```

A packet reaches the router with an IP Precedence value of 4. What priority will this packet be assigned by the router?

- A. Low
- B. Normal
- C. Medium
- D. High
- E. Critical
- F. Flash

Answer: A

Explanation:

The IP precedence values are shown below:

```
0 : routine
1 : priority
2 : immediate
3 : flash
4 : flash Override
5 : critical
6 : internet
7 : network
```

In this example, IP precedence 4 (flash override) was not explicitly defined in the priority

list, so it will be handled by the default queue. In this case, the default queue is given a priority of low.

---

**QUESTION 241**

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.

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/framerelay\\_ts\\_cmd.html](http://www.cisco.com/warp/public/125/framerelay_ts_cmd.html)

---

**QUESTION 242**

Consider the following scenario: An interface has been configured for custom queuing across a DS3 interface. Bandwidth has been allocated for three application flows: A, B and C. The average packet for each application is as follows:

Application A= 2000

Application B= 1000

Application C= 500

You wish to configure the router to allow for 20% of the bandwidth to be allocated to flow A, 50% for flow B, and the remaining 30% for flow C. If only one packet is serviced for flow A per pass, how many packets need to be allowed on flow C to maintain the 20:50:30 ratio?

- A. 3
- B. 4
- C. 5
- D. 6
- E. 500
- F. More information needed

Answer: D

Explanation:

If flow A uses 20% of the bandwidth and flow C uses 30%, then C uses 1.5 times the bandwidth as

A. The average byte size of A is 2000 bytes. 2000 times 1.5 is 3500 bytes.

That would give the bandwidth of 3000 bytes for flow C. Since the average packet size of C is 500 bytes,  $3000/500 = 6$ .

---

**QUESTION 243**

Your VOIP network needs to give priority to the VOIP traffic across the serial interface of a router. You wish to support this by implementing a solution that enables the router to service the Voice traffic in a strict priority queue. All other non-voice traffic should be serviced using the weighted fair queuing mechanism. Which command should you enable on this serial interface?

- A. fair-queue
- B. ip cef
- C. priority-group
- D. ip rtp priority
- E. priority-queuing

Answer: D

Explanation:

The "ip rtp priority" command creates a strict priority queue for voice packets while providing WFQ for non-voice traffic. To reserve a strict priority queue for a set of Real-Time Transport Protocol (RTP) packet flows belonging to a range of User Datagram Protocol (UDP) destination ports, use the ip rtp priority command in interface configuration mode. This command is most useful for voice applications, or other applications that are delay-sensitive.

This command extends and improves on the functionality offered by the ip rtp reserve command by allowing you to specify a range of UDP/RTP ports whose voice traffic is

guaranteed strict priority service over any other queues or classes using the same output interface. Strict priority means that if packets exist in the priority queue, they are dequeued and sent first—that is, before packets in other queues are dequeued.

Incorrect Answers:

A. This command can be used for the default class (commonly known as the classdefault class) only. You can use it in conjunction with either the queue-limit command or the random-detect command. The class-default class is the default class to which traffic is directed if that traffic does not satisfy the match criteria of other classes whose policy is defined in the policy map.

B. This will enable Cisco Express Forwarding, which will not fulfill the requirements of this question.

C, C, E. This will enable priority queuing, which could indeed be used to give RTP packets priority over other protocols, but used alone will not provide the mechanism for having the other traffic types serviced in a WFQ manner as described in this question.

Reference:

[http://www.cisco.com/en/US/products/sw/iosswrel/ps5207/products\\_command\\_reference\\_chapter09186a00801a7edf.html#wp1128299](http://www.cisco.com/en/US/products/sw/iosswrel/ps5207/products_command_reference_chapter09186a00801a7edf.html#wp1128299)

---

**QUESTION 244**

A serial interface with flow-based WFQ is carrying 25 flows in the following fashion:

- Twelve flows are marked as IP Precedence 0.
- Ten flows are marked as IP Precedence 1.
- Three flows are marked as IP Precedence 5.

Based on the above information, how much interface bandwidth is allocated to one of flows that are marked as IP Precedence 5?

- A. 1%
- B. 4%
- C. 12%
- D. 15%
- E. 33%
- F. Cannot tell from the information given

Answer: C

Explanation:

The total parts are found by adding one to each Precedence value, multiplying by the number of flows in that Precedence, and then totaling the parts (weights). Interface bandwidth is allocated to one of flows that are marked as IP Precedence is calculated as shown below.

$$12\% \frac{(5+1)}{[12*(0+1)]+ [10*(1+1)]+ [3*(5+1)]} = \frac{6}{12+20+18} = \frac{6}{50} =$$

Incorrect Answers:

B. If we took one flow and divided it by the number of total flows, the answer would be 4% (1/25). However, the correct answer is found using the formula above.

---

**QUESTION 245**

Using the 3 layer hierarchical approach to a network, What QoS functions are performed at the access layer? (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 246**

You need to give your new VOIP traffic priority over other traffic types in your network. To do this, you plan to implement custom queuing. What statement is FALSE about custom queuing?

- A. Custom queuing defines up to 16 queues.
- B. Custom queuing has one preemptive priority queue. This can be extended to multiple priority queues by configuring the 'lowest-custom' queue in the 'queuelist'.
- C. In custom queuing there is a weight assigned to each queue which specifies how each queue is treated.
- D. With custom queuing you cannot specify a minimum bandwidth guarantee per queue.
- E. In custom queuing you can classify based on the incoming interface.

Answer: B

Explanation:

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.

Custom Queuing (CQ)

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

Why Use CQ?

You can use the Cisco IOS QoS CQ feature to provide specific traffic guaranteed bandwidth at a potential congestion point, assuring the traffic a fixed portion of available bandwidth and leaving the remaining bandwidth to other traffic. For example, you could reserve half of the bandwidth for SNA data, allowing the remaining half to be used by other protocols.

If a particular type of traffic is not using the bandwidth reserved for it, then unused bandwidth can be dynamically allocated to other traffic types.

Restrictions

CQ is statically configured and does not adapt to changing network conditions. With CQ enabled, the system takes longer to switch packets than FIFO because the packets are classified by the processor card.

---

**QUESTION 247**

Which of the following is a required configuration parameter for setting up NBAR?

- A. match protocol IP
- B. match nbar type 1
- C. match ftp session passive
- D. match protocol http
- E. match url www.cisco.com

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 248**

The Certkiller network is using Class of Service to prioritize the traffic throughout the network. Setting the CoS IP Precedence bits can be done in what situation?

- A. For ATM CLP traffic only
- B. To set the frame-relay DE bit
- C. When we receive HDLC frame with DEADBEEF pattern
- D. On a router on ISL or DOT1Q trunks in the output direction only
- E. None of the above
- F. All of the above

Answer: E

#### Explanation:

CoS refers to the ability of a network to provide differentiated service to selected network traffic over packet networks and cell networks. By default, the Cisco IOS software leaves the IP Precedence value untouched, preserving the precedence value set in the header, allowing all internal network devices to provide service based on the IP Precedence setting. This policy follows the standard approach stipulating that network traffic should be sorted into various types of service at the basic perimeter of the network and that those types of service should be implemented in the core of the network. Routers in the core of the network can then use the precedence bits, for example, to determine the order of transmission, the likelihood of packet drop, and so on.

You can use any of the following features to set the IP precedence in packets:

- Policy-Based Routing
- QoS Policy Propagation via Border Gateway Protocol
- Committed Access Rate

#### Incorrect Answers:

- A, B. The COS bit can be set for any FR/ATM traffic, using policy based routing. Frame relay and ATM networks can be configured to adjust traffic settings based on the Cell Loss Priority (CLP) and frame DE values, but the CoS bit is not limited to this type of traffic only.
- C. The CoS bits can not be used based on layer 2 information in PPP and HDLC links, since specific identifiers are needed.
- D. The CoS IP Precedence bits can be set based on both incoming and outgoing traffic.

---

#### **QUESTION 249**

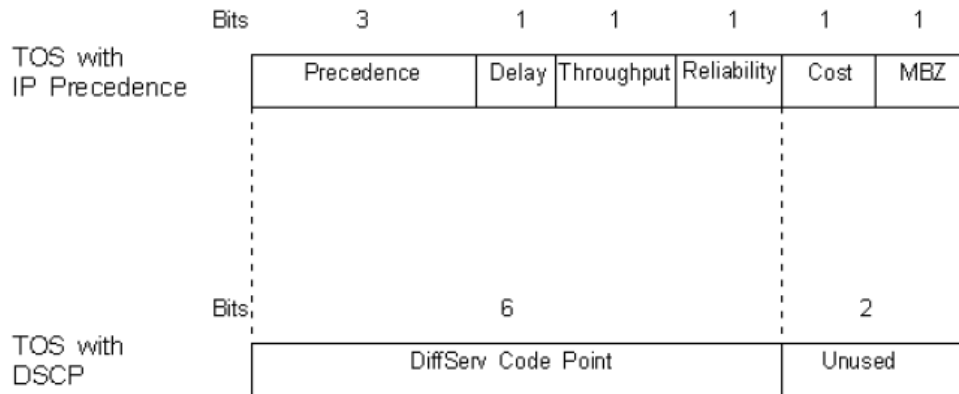
The Certkiller network plans to implement some method of quality of service using DSCP information. In comparing the different options which of the following statements is TRUE?

- A. The IP precedence and DSCP have no overlapping fields.
- B. The DSCP contains class selectors for backward compatibility with the IP precedence.
- C. The DSCP is exactly the same as IP precedence; the name change is merely as marketing naming convention.
- D. The last 2 bit of the DSCP overlaps with the IP precedence.
- E. DSCP is only for TCP; IP precedence is for UDP
- F. None of the above.

Answer: A

Explanation:

DiffServ introduces the concept of the DiffServ Code Point (DSCP) that uses the first 6 bits of the TOS field thereby giving  $2^6 = 64$  different values. RFC 2474 describes the Differentiated Services (DS) field and the DiffServ Code Point (DSCP). A comparison of these two is displayed below:



As you can see from the comparison of the two packet formats, there are no overlapping fields.

---

**QUESTION 250**

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 confirm-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:

- Define a traffic class with the class-map command.
- Create a service policy by associating the traffic class with one or more QoS features (using the policy-map command).
- 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 { <i>bandwidth-</i> <i>kbps</i> percent <i>percent</i> }	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.html](http://www.cisco.com/en/US/products/sw/iosswrel/ps5014/products_feature_guide_chapter09186a008008813a.html)

---

**QUESTION 251**

Priority queuing is being configured on router CK1 to give mission critical traffic priority over the WAN link. What statement is true with regard to priority queuing?

- A. There are 4 priority queues: high, medium, normal, low.
- B. The high and medium queues have precedence over the default queue.
- C. The classification is configurable via the command 'priority-list'
- D. The default queue is the normal queue, by default.
- 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 252**

The IP precedence of a packet can be determine from:

- A. All 8 bits of the ToS byte
- B. Bits 3, 4 and 6 of the ToS byte.
- C. The three most significant bits of the ToS byte.
- D. The three least significant bits of the ToS byte.

Answer: C

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 253**

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
  match-all
int fastethernet0/0
 ip address 10.1.1.1 255.255.255.0
 service-policy input mark
```

Based upon the MQC configuration shown in the exhibit, 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 254

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 extend VoIP-RTP
  permit udp any any range 16384 32767
```

**Certkiller.com**

Classed-Based marking has been configured 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 endusers as CS3 and EF are not remarked to DSCP 00000. Which of the following issues could be the cause of the problem? (Select two).

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.

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 255**

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 0
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

<i>name</i>	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 256**

The following are 3 separate queues for a router configured to prioritize traffic:

350-001

Queue 2

500	1500	500	500	1500
-----	------	-----	-----	------

Queue 1

500	1500	500	1500
-----	------	-----	------

Queue 0

1500	500	1500	500
------	-----	------	-----

Queue 2 is a low-latency queue running in alternate-priority mode. The interface MTU is 1500. The queue weights are as follows:

Weight of 1 for Queue 2

Weight of 2 for Queue 1

Weight of 1 for Queue 0

Assume that all the default counters are currently zero (0) and Queue 2 will be serviced first, how many packets will be left in Queue 2 after both of the other queues have been serviced once?

- A. 0
- B. 1
- C. 3
- D. 500
- E. 1500

Answer: B

Explanation:

Since the example states that Q2 is serviced alternately with Q1 and Q0, the order goes Q2, Q1, Q2, Q0, Q2, etc. The quantum values calculated as  $MTU + (weight-1)*512$  per queue are: Q2 = 1500, Q1 = 2012 ( $MTU + (weight-1)*512$ ) & Q0 = 1500. Therefore immediately after Q0 is serviced there is still a single packet in Q2.

Further clarification (step by step):

MTU is 1500.

Weight of 1 for Queue 0 => 1500 bytes will be de-queued in first round

Weight of 2 for Queue 1 =>  $1500 + (2-1)*512 = 2012$  bytes will be de-queued per round

Weight of 1 for Queue 2 => 1500 bytes will be de-queued in first round

Sequence of de queuing operation...

Q 2 first (500 byte) and second (1500 byte) packet offloaded. Deficit

$1500 - 2000 = -500$

Q 1 first (500 byte) second (1500 byte) third (500 byte) packet

offloaded Deficit  $2012 - 2500 = -488$

Q 2 1500 - 500 = 1000 bytes can be off loaded third (500 byte) and fourth (500 byte)

packet are removed. last packet 1500 byte still remains.

Q 0 First (1500 byte) packet is off loaded.

Q 2 Is being served after both the other queues have had their initial pass.

How many packets remain in Q 2 after first run is what the question is asking.

The answer is 1.

---

**QUESTION 257**

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 `framerelay fair-queue` command), a strict priority queue with WFQ (configured by the `framerelay 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:

- 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).
- 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.
- 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\\_chapter09186a00800bd8ed.html#1001175](http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800bd8ed.html#1001175)

---

**QUESTION 258**

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 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

Answer: A

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-shape adaptive [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 259**

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 or mark incoming traffic.
- C. Class Based Shaping uses class-map statements. A set of hierarchical policy maps are not required for configuring CBS.

---

**QUESTION 260**

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 token 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 are true.

Answer: A

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 7000 bits, but it can indeed be set to 0, which means that no traffic will be able to burst above the CIR.

Incorrect Answers:

- 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 Bc value for priority queuing on frame relay links is 56000, but it is 7000 for regular FRTS.
- E. Bc is the committed burst rate, not the total burst.

---

**QUESTION 261**

In weighted fair queuing (WFQ), one 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.
- D. The CDT defines a value from when IOS starts to account all messages in the WFQ system in conjunction with Netflow.
- E. CDT means 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 262**

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 servicepolicy has to be applied on output.
- D. CBWFQ can only be configured in a hierarchical policy-map. The parent policymap 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 263**

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 264**

You wish to enable the Resource Reservation protocol on one of the interfaces of a router. Which of the following commands will accomplish this?

- A. ip rsvp sender
- B. ip rsvp enable
- C. ip rsvp bandwidth
- D. rsvp enable
- E. ip rsvp reservation
- F. RSVP is enabled in global configuration mode, not in interface configuration mode.

Answer: C

Explanation:

ip rsvp bandwidth is the command that enables RSVP

Incorrect Answers:

- F. RSVP is configured on a per interface basis, not in global configuration mode.

---

**QUESTION 265**

Which of the following statements is valid regarding Custom Queuing?

- A. Custom queuing always services the highest priority traffic first before servicing

the lower priority traffic.

- B. Custom queuing looks at groups of packets from the similar source-destination pairs.
- C. Custom queuing processes the queue based on the number of packets sent.
- D. Custom queuing will not proceed to a next queue unless the current queue is empty.
- E. Custom queuing can prevent one type of traffic from saturating the entire link.

Answer: E

Explanation:

CQ allows fairness not provided with priority queuing (PQ). With CQ, you can control the available bandwidth on an interface when it is unable to accommodate the aggregate traffic enqueued. Associated with each output queue is a configurable byte count, which specifies how many bytes of data should be delivered from the current queue by the system before the system moves on to the next queue. When a particular queue is being processed, packets are sent until the number of bytes sent exceeds the queue byte count defined by the queue-list queue byte-count command, or until the queue is empty. With custom queuing, all queues will be serviced. With priority queuing, a bandwidth hog can dominate the link.

Incorrect Answers:

- A, D Custom queue uses a round robin mechanism, ensuring that one type of traffic (even ones with the highest priority) does not completely starve out the lower priority queues. Once the byte count for that queue is fulfilled, the next queue is serviced.
- B, C. Custom queues are serviced based on the number of bytes sent for each queue, not on the number of packets sent. This prevents traffic with bigger packets (such as FTP) from dominating a link with smaller packets (such as a RTP session).

---

**QUESTION 266**

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:

- Extended burst parameter value
- Compounded debt. Compounded debt is computed as the sum over all  $a_i$ .
  - o  $i$  indicates the  $i$ th packet that attempts to borrow tokens since the last time a packet was dropped.
  - o  $a_i$  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](http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/qos_c/qcpart4/qcpolts.htm)

---

### QUESTION 267

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](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](http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_tech_note09186a00800fb50a.shtml)

---

**QUESTION 268**

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 269**

You issue the following configuration change on router CK1 :

```
ip rsvp sender 225.1.1.1 192.1.2.1 UDP 3030 192.1.2.1 serial0  
20 1
```

What is the effect of this change?

A. The router will simulate receiving RSVP PATH messages destined to multicast address 225.1.1.1 from source 192.1.2.1.

The previous hop of the PATH message is 192.1.2.1, and the message was received on interface serial 0.

B. The router will simulate generating RSVP RESV messages destined to multicast address 225.1.1.1 from source 192.1.2.1.

The next hop of the PATH message is 192.1.2.1, and the message was received on interface serial 0.

C. The router will act as if it was sending RSVP PATH messages destined to multicast address 225.1.1.1 from source 192.1.2.1.

The next hop of the PATH message is 192.1.2.1, and the message was received on interface serial 0.

D. The router will act as if it was receiving RSVP RESV messages destined to multicast address 225.1.1.1 from source 192.1.2.1.

The previous hop of the PATH message is 192.1.2.1, and the message was received on interface serial 0.

Answer: A

Explanation:

This command causes the router to act as if it were receiving PATH messages destined to multicast address 225.1.1.1 from a source 12.1.2.1. The previous hop of the PATH message is 12.1.2.1, and the message was received on interface Serial 0.

To enable a router to simulate receiving and forwarding Resource Reservation Protocol (RSVP) PATH messages, use the ip rsvp sender global configuration command. To disable this feature, use the no form of this command.

```
ip rsvp sender session-ip-address sender-ip-address {tcp | udp | ip-protocol} sessionport  
sender-sport previous-hop-ip-address previous-hop-interface bandwidth burst-size
```

Incorrect Answers:



- B. This answer describes the "ip rsvp reservation-host" command.
  - C. This answer describes the "ip rsvp sender-host" command
  - D. The "ip rsvp sender" command simulates a host that is receiving PATH messages, not RESV messages.
- 

**QUESTION 270**

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 271**

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#47832](http://www.cisco.com/en/US/products/sw/iosswrel/ps1834/products_feature_guide09186a0080080232.html#47832)

---

**QUESTION 272**

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](http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a00801b2406.html)

---

**QUESTION 273**

Which of the following terminal types of ISDN is used by an AT&T 5ESS NI1 switch?

- A. Terminal type A.
- B. Terminal type B.
- C. Terminal type S.
- D. Terminal Type T.
- E. All of the above.

Answer: A

Explanation:

Switch Hardware: 5ESS; Software Variant: National (All NIs)

Cisco IOS configuration command: `isdn switchtype basic-ni`

Terminal Type = A

Two B-channels for voice and data

Two Directory numbers, assigned by service provider

Two SPIDs are required, assigned by service provider; format will vary

Can have directory number 1 hunt to directory number 2, it does cost a little extra money

Incorrect Answers:

C, D. S and T ISDN reference types, not terminal types. Reference point S is the term for the connection between the CPE and the NT2, and reference point T is for the connection between the NT1 and the ISDN network.

Reference:

<http://www.cisco.com/warp/public/129/10.html>

---

**QUESTION 274**

Your ISDN equipment is a TE2 device. What type of device is this?

- A. A device that manages the switching functions.
- B. A device that use the standard native ISDN interface.
- C. A device that marks the boundary between the end user-CPE and the NT2.
- D. A device that does not use the standard ISDN interface, and require a terminal adapter.
- E. A device that converts the BRI signals into a form used by the ISDN digital line.

Answer: D

Explanation:

A TE2 (Terminal Equipment Type 2 designates a device that does not include a native

ISDN interface and requires a TA for its ISDN signals.

Incorrect Answers:

- A. ISDN switches manage the switching functions.
- B. TE1 designates a device with a native ISDN interface.
- C. The S reference point is used to designate this boundary.
- E. This describes what a NT1 device does.

---

**QUESTION 275**

Which ISDN reference point is only applicable in North America?

- A. R
- B. U
- C. T
- D. A
- E. S
- F. None of the above.

Answer: B

Explanation:

The U reference point is used in North America only. U is not specified in any ITU-T standard.

Incorrect Answers:


A, C, E, R, S, and T are all reference points are specified in I.411 and I.412 and are used in North America as well as internationally.

D. A is not a reference point, it is an ISDN switch type used by AT&T.

---

**QUESTION 276**

Six separate provisioning options for an ISDN circuit are displayed below:



The Certkiller network needs new ISDN lines to be installed to be used as backup links. Which of the choices above depict valid ISDN interface provisioning options for a single local loop?

- A. A

- B. B
- C. C
- D. D
- E. E
- F. F

Answer: F

Explanation:

The only valid ISDN channel configurations are BRI and PRI. BRI comes with 2 64K bearer channels and 1 16K data channel. PRI comes with 23B channels and 1 64K D channel for the T1 version and 30B plus 1 64K D channel for the E1 version. Only choice F lists one of these choices. More information can be found below:

**ISDN Channels**

Two types of ISDN interfaces are focused on in IOS documentation: Basic Rate Interface (BRI) and Primary Rate Interface (PRI). Both BRI and PRI provide multiple digital bearer channels over which temporary connections can be made and data can be sent. The result is digital dial access to multiple sites concurrently. Table 8-24 summarizes the features of BRI and PRI.

Table 8-24 *BRI and PRI Features*

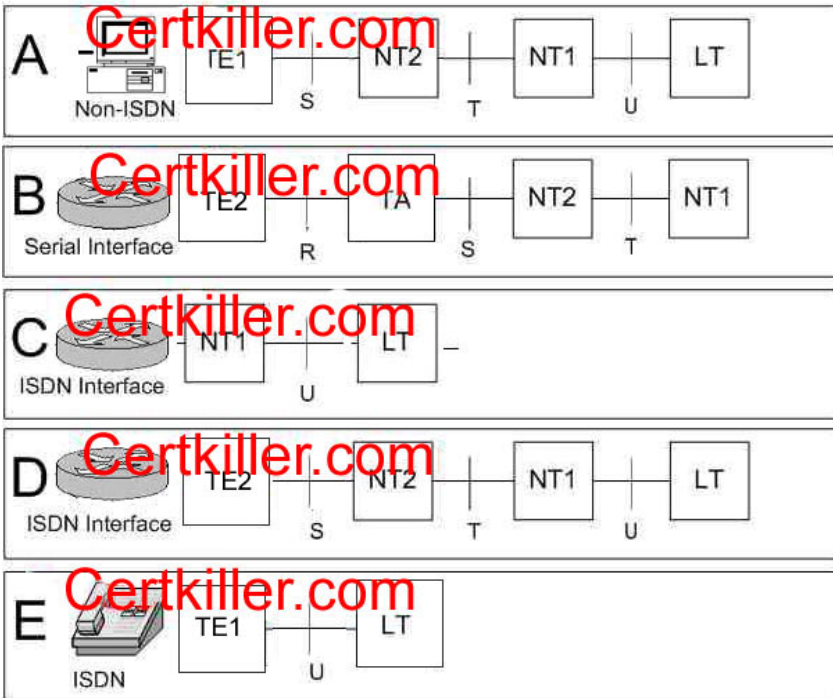
Type of Interface	Number of Bearer Channels (B Channels)	Number of Signaling Channels (D Channels)
BRI	2	1 (16kbps)
PRI (T1)	23	1 (64kbps)
PRI (E1)	30	1 (64kbps)

Bearer channels (B channels) are used to transport data. B channels are called bearer channels because they bear the burden of transporting the data. B channels operate at up to 64kbps, although the speed might be lower depending on the service provider. The section "ISDN Configuration," later in the chapter, discusses how to configure the correct speed for the bearer channels. D channels are used for signaling.

---

**QUESTION 277**

The exhibit below depicts five different ISDN equipment configurations



Which two ISDN Local Loop diagrams correctly depict valid ISDN connectivity between devices and ISDN Interfaces?

- A. Diagram A
- B. Diagram B
- C. Diagram C
- D. Diagram D
- E. Diagram E

Answer: B, C

Explanation:

The following describe the ISDN components and their reference points:

TE1 (Terminal Equipment Type 1): Designates a device with a native ISDN interface.

TE2: Designates a device that does not include a native ISDN interface and requires a TA (Terminal Adapter) for its ISDN signals.

NT1 (Network Termination 1): Designates a device that converts the BRI signals into a form used by the ISDN digital line, and is the boundary between the carrier ISDN network and the CPE.

NT2: Designates a device that aggregates and switches all ISDN lines at a customer location.

TA: Used by a TE2 device to convert signals into BRI signals.

Reference Points:

R: The connection between a non-ISDN compatible device (TE2) and a TA.

S: The connection between the CPE and the NT2

T: The connection between the NT2 and the NT1

U: The connection between the NT1 and the carrier ISDN network.

Based on these definitions, choices B and C are correct.

Incorrect Answers:

A. A TE1 device is designated for a native ISDN interface. The equipment shown in choice A is not a native ISDN type.


D. A TE2 device is used for a non native ISDN interface. The router in this choice shows a native ISDN interface.

E. Reference point U is used to describe the connection between the ISDN cloud and an NT1, not the ISDN cloud and a TE1.

---

**QUESTION 278**

Six separate provisioning options for an ISDN circuit are displayed below:



New ISDN lines need to be provisioned to back up the Certkiller data network.

Which of the choices above depict valid ISDN configuration options? (Select three.)

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

Answer: A, D, E

The only valid ISDN channel configurations are BRI and PRI. BRI comes with 2 64K bearer channels and 1 16K data channel. PRI comes with 23B channels and 1 64K D channel for the T1 version and 30B plus 1 64K D channel for the E1 version. More

information can be found below:

### ISDN Channels

Two types of ISDN interfaces are focused on in IOS documentation: Basic Rate Interface (BRI) and Primary Rate Interface (PRI). Both BRI and PRI provide multiple digital bearer channels over which temporary connections can be made and data can be sent. The result is digital dial access to multiple sites concurrently. Table 8-24 summarizes the features of BRI and PRI.

Table 8-24 *BRI and PRI Features*

Type of Interface	Number of Bearer Channels (B Channels)	Number of Signaling Channels (D Channels)
BRI	2	1 (16kbps)
PRI (T/1)	23	1 (64kbps)
PRI (E/1)	30	1 (64kbps)

Bearer channels (B channels) are used to transport data. B channels are called bearer channels because they bear the burden of transporting the data. B channels operate at up to 64kbps, although the speed might be lower depending on the service provider. The section "ISDN Configuration," later in the chapter, discusses how to configure the correct speed for the bearer channels. D channels are used for signaling.

---

### QUESTION 279

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 280

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](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 281**

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  $(1500-13)/5 = \text{max DLCI's} = 297.4$ . 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:  $(\text{MTU} - 13) / 5 = \text{max DLCI's}$ .

For Cisco, the formula is  $(\text{MTU} - 13) / 8 = \text{max DLCI's}$ .

---

**QUESTION 282**

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 283**

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

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 trafficshaping 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](http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide09186a0080087b91.html)

---

**QUESTION 284**

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 three)

- A. A
- B. B
- C. C

- D. D
- E. E

Answer: 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.

You might expect your equipment to notify end stations to stop sending data to keep additional frames from being discarded or hitting a congested network. In practice, however, this doesn't happen: Most routers, bridges and frame relay access devices (FRADs) do nothing when these bits get set. Instead, they expect higher layer protocols, such as TCP/IP, to know how to react implicitly to the packet loss.

Incorrect Answers:

A, 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 285**

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.

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 286**

The Certkiller frame relay network is displayed in the following diagram:



```
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 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-config
interface Ethernet0
 ip address 123.123.123.1 255.255.255.0
!
interface Serial
 ip address 3.1.3.2 255.255.255.0
 encapsulation frame-relay
 frame-relay map ip 3.1.3.3 150
 frame-relay interface-dlci 150
```

Reference:

[http://www.cisco.com/en/US/tech/CK713/CK237/technologies\\_tech\\_note09186a008014f8a7.shtml](http://www.cisco.com/en/US/tech/CK713/CK237/technologies_tech_note09186a008014f8a7.shtml)

---

**QUESTION 287**

Split Horizon is often used with Poison Reverse to prevent routing loops. Of the following choices, which statement is FALSE regarding the rule of Split Horizon?

- A. It can cause problems on certain Frame-Relay Hub-and Spoke configurations.
- B. It is enabled by default on multipoint Frame-Relay subinterfaces.
- C. It can be disabled for IP/RIP and IPX/RIP.
- D. It aids in preventing routing loops.
- E. None of the above.

Answer: C

Explanation:

For both point to point and point to multipoint sub-interfaces, split horizon is disabled by default. For physical serial frame relay multipoint interfaces, it is enabled by default. Frame Relay subinterfaces provide a mechanism for supporting partially meshed Frame Relay networks. Most protocols assume transitivity on a logical network; that is, if station A can talk to station B, and station B can talk to station C, then station A should be able to talk to station C directly. This is true on LANs, but is not true on Frame Relay networks unless A is directly connected to C.

Additionally, certain protocols such as AppleTalk and transparent bridging could not be supported on partially meshed networks because they require "split horizon," in which a packet received on an interface cannot be transmitted out the same interface even if the packet is received and transmitted on different virtual circuits.

By configuring Frame Relay subinterfaces, a single physical interface is treated as multiple virtual interfaces. This allows us to overcome split horizon rules. Packets received on one virtual interface can now be forwarded out another virtual interface, even if they are configured on the same physical interface.

Subinterfaces address these limitations by providing a way to subdivide a partially meshed Frame Relay network into a number of smaller, fully meshed (or point-to-point) subnetworks. Each subnetwork is assigned its own network number and appears to the protocols as if it is reachable through a separate interface. (Note that point-to-point subinterfaces can be unnumbered for use with IP, reducing the addressing burden that might otherwise result.)

IP RIP can indeed have split horizon disabled. This can be accomplished via the use of sub-interfaces, or with the "no ip split-horizon" interface command. This will disable split horizons for IP traffic, including RIP. However, IPX RIP traffic can not be disabled so this statement is false.

Incorrect Answers:

A: This statement is true. For networks using distance vector routing protocols, spoke site to spoke site connectivity can be affected due to the split horizon rule.

B: Cisco serial interfaces are multipoint interfaces by default unless specified as a point-to-point subinterface. Though less common than point-to-point subinterfaces, it is possible to divide the interface into separate virtual multipoint subinterfaces.

Multipoint interfaces/subinterfaces are still subject to the split-horizon limitations as discussed above. All nodes attached to a multipoint subinterface belong to the same network number. Typically, multipoint subinterfaces are used in conjunction with point-to-point interfaces in cases where an existing multipoint frame relay cloud is migrating to a subinterfaced point-to-point network design. A multipoint subinterface is used to keep remote sites on a single network number while slowly migrating remote sites to their own point-to-point subinterface network.

D: Routing loop prevention is the reason why split horizon was created.

References:

[http://www.cisco.com/en/US/products/sw/iosswrel/ps1818/products\\_configuration\\_guide\\_chapter09186a00800878c7.html](http://www.cisco.com/en/US/products/sw/iosswrel/ps1818/products_configuration_guide_chapter09186a00800878c7.html)

[http://www.alliancedatacom.com/manufacturers/ciscosystems/framerelay\\_design/subinterfaces.asp](http://www.alliancedatacom.com/manufacturers/ciscosystems/framerelay_design/subinterfaces.asp)

---



**QUESTION 288**

In an ATM network, what is Traffic control, traffic shaping, and traffic policing used for?

- A. They are parameters of PNNI set during PNNI configuration.
- B. They are all Quality of Service features of an ATM network.
- C. They are types of SVCs.
- D. They are types of PVCs.
- E. They are all components needed for LANE

Answer: B

Explanation:

Traffic Shaping, Policing, and Traffic Control are all forms of QoS.

---

**QUESTION 289**

Which of the following are requirements for a LANE network? (Choose all that apply)

- A. LECS
- B. SSRP
- C. BUS
- D. AAL5
- E. ILM
- F. LES

Answer: A, C, F

Explanation:

LECS, LES, and BUS are all servers that are needed for LAN Emulation (LANE).

The various LANE servers and their requirements are presented below:

The LAN Emulation Configuration Server (LECS)

LANE requires that each LAN Emulation Client (LEC) establish a virtual circuit (VC) to the LAN Emulation Configuration Server (LECS) when it goes up. The LEC then requests the ATM address of its corresponding LAN Emulation Server (LES). Once the LEC has its ATM LES address, the VC between the LEC and the LECS is removed, and the LEC no longer tries to communicate with the LECS. When the environment is stable and all LECs are up and operational, the LECS is idle. When the LECs join the emulated LAN (ELAN), they each contact the LECS individually. However, when the LANE network undergoes a disaster (for example, when the primary LECS fails), all clients go down.

The LAN Emulation Server (LES)

Each LEC will maintain a bi-directional VC to the LES of the ELAN (it may be more than one ELAN if FSSRP is used). In a typical highly loaded environment, many LAN Emulation Address Resolution Protocol (LE\_ARP) requests will be sent to the LES. The implementation of the LES on Cisco devices is quite straightforward. All incoming

LE\_ARP frames will be forwarded to the control distribute virtual channel connection (VCC).

Remember that in Cisco's implementation, LES and Broadcast and Unknown Server (BUS) processes are combined (that is, you can't put the LES for ELAN-1 on one device, and the BUS for ELAN-1 on another device).

The Broadcast and Unknown Server (BUS)

The job of the BUS is quite similar to the job of the LES. Each LEC is required to have one multicast send to the BUS. The LEC sends all its multicast, broadcast or unknown traffic to it. The BUS has a point-to-multipoint VCC to all LECs in the ELAN. Frames do not have to be examined in detail by the BUS. In other words, each incoming frame on the multicast send can be blindly forwarded to the multicast forward.

---

**QUESTION 290**

What ATM service category uses a closed-loop feedback mechanism to determine the rate at which traffic can be sent?

- A. ABR
- B. CBR
- C. UBR
- D. VBR-rt
- E. VBR-nrt

Answer: A

Explanation:

When you assign an ATM virtual circuit to the ABR service category it configures a router to transmit at a rate that varies with the amount of bandwidth available in the network or along the end-to-end transmission path. When the network is congested and other source devices are transmitting, there is little available or leftover bandwidth. However, when the network is not congested, bandwidth is available for use by other active devices. ABR allows end-system devices like routers to take advantage of this extra bandwidth and increase their transmission rates. Therefore, ABR uses mechanisms that allow ABR VCs to make use of any bandwidth available in the network at any point in time.

An ABR VC binds a source router to a contract with the ATM switch network. As part of this contract, a source router agrees to examine information that indicates whether or not the network is congested and, in turn, adapt the source transmission rate if required. In return, the ATM switch network agrees to drop no more than a maximum number of cells when congestion occurs. The ratio of dropped cells to transmitted cells is known as the cell loss ratio (CLR).

In addition, an ABR VC uses a closed-loop model. With a closed loop, a source router sends data cells or special cells (called forward resource management [RM] cells) into the ATM network. The switches in the ATM network mark or set bits in these cells as they flow along the end-to-end path. The destination router turns these cells around as backward RM cells. By setting certain bits or fields, the ATM network and destination

router provide feedback used to control the source rate in response to bandwidth changes in the network or at the destination.

---

**QUESTION 291**

What are the necessary components for an ATM SVC? (Choose all that apply)

- A. QoS type set to CBR+
- B. Path Signaling
- C. LECS
- D. ILMI PVC's
- E. LES

Answer: B, D

Explanation:

Signaling along the path, and an ILMI PVC are required for a Switched Virtual Circuit.

Incorrect Answers:

- A. The QoS configuration is not a requirement for an ATM SVC.
  - C. LECS is the LAN Emulation Configuration Server. This is a LANE requirement, not an SVC requirement.
  - E. LES is the LAN Emulation Server, which is also a LANE component.
- 

**QUESTION 292**

Which of the following statements are true regarding PNNI in an ATM network? (Choose all that apply).

- A. PNNI has QoS awareness and has layer 2 reachability support.
- B. PNNI ignores QoS.
- C. PNNI supports layer 2 and layer 3 routing.
- D. PNNI comes up automatically when a switch is brought up.
- E. PNNI stands for Private Network-Node Interface.

Answer: A, E

Explanation:

Private Network-Node Interface (PNNI) protocol provides mechanisms to support scalable, QoS-based ATM routing and switch-to-switch switched virtual connection (SVC) interoperability.

IISP provides a static routing solution that is not easily scalable and has no support for quality of service (QoS). PNNI provides a highly scalable routing solution with dynamically determined routing paths and support for QoS requirements.

Incorrect Answers:

- B. PNNI differs in IISP in that it provides QoS based routing within the ATM network.
- C. PNNI does not have the capability to support layer 3, IP based routing.
- D. PNNI must be explicitly configured throughout the network in order for it to function.

Reference:

[http://www.cisco.com/univercd/cc/td/doc/product/atm/c8540/12\\_0/13\\_19/atg/pnni.htm](http://www.cisco.com/univercd/cc/td/doc/product/atm/c8540/12_0/13_19/atg/pnni.htm)

---

**QUESTION 293**

Which of the following are ATM Reference Model Layers? (Choose all that apply)

- A. ATM layer
- B. ATM adaptation layer (AAL)
- C. Generic Flow Control (GFC) layer
- D. Session
- E. Physical layer
- F. None of the above

Answer: A, B, and E

Explanation:

The physical layer, ATM layer, and AAL make up the three layers of the ATM reference model.

Incorrect Answers:

- C. GFC is not a layer of the ATM model.
  - D. The session layer is an OSI model layer but is not part of the ATM model.
- 

**QUESTION 294**

You wish to connect your ATM core switch to an already existing switch in another ATM network. Which type of connection would you use to accomplish this?

- A. DLCI
- B. FDDI
- C. NNI
- D. UNI
- E. LMI
- F. PNNI

Answer: C

Explanation:

Network to Network Interface (NNI) is the term used for the connection between two ATM switches.

Incorrect Answers:

- A. DLCI (Data Link Circuit Identifier) is used for frame relay connections.
- B. FDDI is the Fiber Distributed Data Interface, which is a legacy high speed LAN technology.
- D. UNI is the User to Network Interface. This would describe the connection between an end user and an ATM switch.
- E. LMI is the Link Management Interface, used as a keepalive function over frame relay

networks.

F. Private Network-Node Interface (PNNI) protocol provides mechanisms to support scalable, QoS-based ATM routing and switch-to-switch switched virtual connection (SVC) interoperability.

---

**QUESTION 295**

What ATM Adaptation Layer is used to support bursty traffic requiring cells to be delivered in order?

- A. AAL1
- B. AAL2
- C. AAL3
- D. AAL4
- E. AAL5

Answer: A

Explanation:

AAL1 is designed to process Class A, constant bit rate (CBR) traffic from a higher layer application or protocol and to deliver that traffic to its destination at the same rate and at equal intervals. A prime example of this type of service is a connection that supports voice traffic or telephony circuits.

In this type of service, misordering of cells is considered more problematical than losing cells. Hence, a 3-bit sequence number (SN) is added when forming the basic unit of information transfer (the SAR-PDU) for AAL1 processing (see Figure 1-11). The sequence number embodied therein assists in detecting and correcting lost or misinserted cells.

In AAL1 traffic processing, user data is transferred between communicating peers at a constant bit rate after an appropriate connection has been established.

---

**QUESTION 296**

While troubleshooting an ATM link on one of the Certkiller routers, the "show interface atm" command was issues as shown below:

```
Router# show interfaces atm 4/0
ATM4/0 is up, line protocol is up
Hardware is cxBus ATM
Internet address is 131.108.97.165, subnet mask is 255.255.255.0
MTU 4470 bytes, BW 100000 Kbit, DLY 100 usec, rely 255/255, load 1/255
ATM E164 Auto Conversion Interface
Encapsulation ATM, loopback not set, keepalive set (10 sec)
Encapsulation(s): AAL5, PVC mode
256 TX buffers, 256 RX buffers, 1024 Maximum VCs, 1 Current VCs
Signalling vc = 1, vpi = 0, vci = 5
ATM NSAP address: BC.CDEF.01.234567.890A.BCDE.F012.3456.7890.1234.13
Last input 0:00:05, output 0:00:05, output hang never
Last clearing of "show interface" counters never
Output queue 0/40, 0 drops; input queue 0/75, 0 drops
Five minute input rate 0 bits/sec, 0 packets/sec
Five minute output rate 0 bits/sec, 0 packets/sec
 144 packets input, 31480 bytes, 0 no buffer
  Received 0 broadcasts, 0 runts, 0 giants
 13 input errors, 12 CRC, 0 frame, 0 overrun, 1 ignored, 0 abort
 154 packets output, 4228 bytes, 0 underruns
   0 output errors, 0 collisions, 1 interface resets, 0 restarts
```

Based on this information, what causes the CRC counter in the output to increment?

- A. An ATM cell was corrupted and arrived at the interface with an invalid protocol value.
- B. An ATM cell was corrupted and arrived at the interface with an invalid HEC.
- C. An ATM cell was dropped or corrupted resulting in an invalid CRC for an AAL5 frame.
- D. An OAM cell was dropped or corrupted in the WAN.
- E. None of the above

Answer: C

Explanation:

RFC 2515 defines the CRC value as:

The number of AAL5 CPCS PDUs received with CRC-32 errors on this AAL5 VCC at the interface associated with an AAL5 entity

The output of show interfaces commands on Cisco devices includes numerous counters.

One such counter is CRC, which counts the number of times (that is, for how many packets) the checksum generated by the originating station, or far end device, does not match the checksum calculated from the data received. By doing this, CRC detects changes to a protocol data unit (PDU) during transmission. It is important that we retain the true value of this PDU because we want to ensure that the destination correctly interprets the data that we're communicating.

Which CRC Are We Checking?

ATM supports five ATM adaptation layers (AALs). AAL5 appends an eight-byte trailer to the common part convergence sublayer protocol data unit (CPCS-PDU), which consists of the original layer-3 packet (for instance, an IP packet) before it segments into 53-byte cells. When you configure a permanent virtual circuit (PVC) with the encapsulation aal5snap command, you are telling it to use this AAL5 trailer. You also are specifying a Logical Link Control (LLC) or Subnetwork Access Protocol (SNAP) header, which is similarly used with Ethernet.

Note: On Cisco routers, the terms "frame", "AAL5 frames" and "CPCS-PDU" all refer to the same concept when we talk about ATM interfaces.

The following are some potential reasons for ATM CRC errors:

- Dropped cells due to traffic policing in the ATM cloud on one or more VCs attached to the ATM interface.
- Noise, gain hits, or other transmission problems on the data-link equipment.
- A faulty or failing ATM interface.

Reference:

[http://www.cisco.com/en/US/tech/CK39/CK48/technologies\\_tech\\_note09186a00800c93ef.shtml](http://www.cisco.com/en/US/tech/CK39/CK48/technologies_tech_note09186a00800c93ef.shtml)

---

**QUESTION 297**

You are seeing a large number of clocking problems on the serial interface of one of your routers. Which of the following would NOT cause this? (Choose all that apply.)

- A. Several cables connected together in a row.
- B. Impedance mismatching.
- C. Improper DSU configuration.
- D. Mismatching encapsulations on each end.
- E. Improper CSU configuration.

Answer: B, D

Explanation:

Impedance problems would cause errors on the line but not clocking problems. Although the encapsulation for any serial interface should match on each end for proper connectivity, this would also not cause clocking problems.

Incorrect Answers:

A, C, E are all possible causes for clocking problems.

---

**QUESTION 298**

A serial interface on a Cisco router is being connected to an external CSU/DSU. The CSU/DSU has an RS-232 interface with a DB-25 connection. Which cables would be used to connect the router to the external CSU/DSU?

- A. DB-60 female to DB-25 male (DTE)
- B. DB-60 male to DB-25 female (DTE)
- C. DB-60 male to DB-25 female (DCE)
- D. DB-60 female to DB-25 female (DTE)
- E. None of the above

Answer: A

Explanation:

Devices that communicate over a serial interface are divided into two classes: DTE and

DCE. The most important difference between these types of devices is that the DCE device supplies the clock signal that paces the communications on the bus. The following chart is a guideline for choosing the correct cable.

	DTE	DCE	Selectable DTE or DCE*
Device	Terminals, Data Service Unit/Channel Service Unit (DSU/CSU), Multiplexors	Modems	Hubs, Routers
Gender	Male	Female	Either

\* Selectable devices usually have a jumper, switch, or software command used to select DTE or DCE.

Incorrect Answers:

B, D. The DB-25 connection should be female, not male.

C. As shown by the chart above, the cable should be DTE, since it is connecting to a CSU/DSU.

---

**QUESTION 299**

Which of the following statement is true regarding clocking for a Cisco T1 interface?

A. The clock source command selects a source for the interface to clock received data.

By default, it is clock source loop-timed (specifies that the T1/E1 interface takes the clock from the Tx (line) and uses it for Rx).

B. Routers are DTEs and NEVER supply clocking to T1/E1 line.

C. The clock source command specifies the location of the NTP server for timing.

D. The clock source selects a source for the interface to clock outgoing data.

The default is clock source line -Specifies that the T1/E1 link uses the recovered clock from the line.

E. The clock source identifies the stratum level associated with the router T1/E1. The default is Stratum 1.

Answer: D

Explanation:

Clocking can either be internal, looped, or line. The default is line, meaning that the router is receiving clocking from the carrier network line.

Incorrect Answers:

C, E. These answers relate to NTP services, which are used for providing time stamping information to the router and does not relate to clocking. Stratum levels provide a hierarchy to the NTP source, with the highest level as 1.

---

**QUESTION 300**

You are troubleshooting connection problems from router CK1 . In doing so, you



350-001

issue the "show interface serial 0" command and see: "serial 0 is up, line protocol down (disabled)." What can you conclude from this?

- A. The Serial0 interface is operating properly.
- B. The Serial0 interface needs to be enabled with the no shut down command.
- C. The Serial0 interface is not working properly due to telco service problems.
- D. The Serial0 interface is using the wrong protocol.
- E. A loop exists in the circuit.

Answer: C

Explanation:

The line: Serial 0 is up, line protocol is down (disabled) indicates a telephone company service problem or a CSU/DSU hardware problem.

Incorrect Answers:

- A. A properly working serial interface would show "serial 0 is up, line protocol is up"
- B. If the interface was manually shut down, it would read: "serial 0 is administratively down, line protocol is down"
- D. The protocols that ride over the serial interface do not affect the interface state. If the wrong encapsulation was configured, then the interface would most likely be down.
- E. If a loop exists in the interface, then it would show up as looped, such as "serial 0 is up, line protocol is up (looped)."

Reference:

[http://www.cisco.com/en/US/products/hw/voiceapp/ps967/products\\_administration\\_guide\\_chapter09186a00801946a3.html#1057773](http://www.cisco.com/en/US/products/hw/voiceapp/ps967/products_administration_guide_chapter09186a00801946a3.html#1057773)