Some New QoS Questions (Answers and Explanations)

1. CB-WRED is configured using the **random-detect** command. Which two of the following statements are true concerning the **random-detect**command? (Choose 2) A. The **random-detect** command cannot be issued for the class-default class. B. The **random-detect** command cannot be issued for the priority class(es). C. The **random-detect** command must be issued in conjunction with the bandwidth command (with the exception of the class-default class). D. The **random-detect** command should be issued in conjunction with the priority command. **Answer:** B, C **Explanation:** Weighted Random Early Detection (WRED) is effective for TCP flows, because WRED can cause some TCP flows to enter TCP slow start. When configuring class-based WRED (i.e. CB-WRED), the **random-detect** command is issued in policy-map-class configuration mode. While the **random-detect** command can be used with the **class-default** class, **random-detect** cannot be issued in policy-map-class configuration mode for a class configured with the **priority** keyword. Also, with the exception of the class-default class, the **random-detect** command must be issued along with the **bandwidth** command. **2.** Based on the following configuration, what traffic will be policed? class-map C_MUSIC

  match protocol kazaa2

  match protocol napster

!

class-map match-any C_WEB

  match protocol http

  match class-map C_MUSIC

!

policy-map P_WEB

  class C_WEB

    police 64000

!

interface serial 0/0

  service-policy output P_WEB

A. All Kazaa version 2 traffic is policed

B. All Napster traffic is policed

C. All web traffic is policed

D. All Kazaa version 2, Napster, and web traffic is policed

E. No traffic is policed

Answer:

C

Explanation:

The C_MUSIC class-map does not specify the **match-any** or **match-all** option. The default is **match-all**. Therefore, for traffic to be classified in the C_MUSIC class-map, a packet would simultaneously have to be a Kazaa version 2 packet and a Napster packet, which isn't possible.

The C_WEB class-map uses the **match-any** option, meaning that traffic will be classified in this class-map if it is HTTP traffic or if it is traffic that was classified in the C_MUSIC class-map. Since, no traffic will be classified in the C_MUSIC class-map, as described above, the only traffic that will be classified by the C_WEB class-map is HTTP traffic.

The policy-map P_WEB is configured to police (i.e. rate limit) traffic classified by the C_WEB class-map to a bandwidth of 64 kbps. (NOTE: The default conform-action is transmit, and the default exceed-action is drop.) Since only HTTP (i.e. web) traffic is matched by the C_WEB class-map, web traffic is the only traffic that is policed.

3. You are configuring a Cisco Catalyst 3560 switch port to trust CoS markings if, and only if, the marking originated from a Cisco IP Phone. In an attempt to perform this configuration, you enter the **mls qos trust device cisco-phone** command. However, your configuration does not seem to be working properly. Why is the switch not trusting CoS markings coming from an attached Cisco IP Phone?

A. A Cisco Catalyst 2950 switch supports the **mls qos trust device cisco-phone** command, but the Cisco Catalyst 3560 does not support this command

B. The **mls qos trust cos** command is missing

C. The **mls qos trust extend** command is missing

D. The **mls qos cos 5** command is missing

E. The PC attached to the phone is overriding the CoS markings

Answer:

B

Explanation:

A Cisco Catalyst 2950 switch port can be configured to trust Class of Service (CoS) markings, Differentiated Services Code Point (DSCP), or CoS markings originating from a Cisco IP Phone. The switch port can detect that a CoS marking is coming from a Cisco IP Phone via the Cisco Discovery Protocol (CDP). The **mls qos trust device cisco-phone** command does indeed tell the switch to trust a marking if, and only if, the marking comes from a Cisco IP Phone. However, the **mls qos trust device cisco-phone** command by itself does not tell the switch port which marking (i.e. CoS or DSCP) coming from the Cisco IP Phone to trust. Therefore, the **mls qos trust cos** command is also required.

4. You administer a network that transports both voice and interactive video traffic. Since these traffic types are both latency-sensitive, you decide to implement the following configuration. Which statement is true regarding the configuration?

class-map C_VOICE

  match protocol rtp audio

!

class-map C_VIDEO

  match protocol rtp video

!

policy-map P_HIGH_PRIORITY

  class C_VOICE

    priority percent 15

  class C_VIDEO

    priority percent 35

  class class-default

    fair-queue

!

interface serial 0/0

  service-policy output P_HIGH_PRIORITY

A. The configuration results in three queues, one for the C_VOICE class, one for the C_VIDEO class, and one queue for the class-default class

B. The configuration results in two queues, one priority queue and one queue for the class-default class

C. The class-default class uses FIFO as its queuing mechanism for traffic flows within its queue

D. The two priority queues use WFQ for queuing traffic within those queues

Answer:

B

Explanation:

While priority treatment (i.e. LLQ treatment) can be assigned to more than one class-map, an interface only has one priority queue. Therefore, in the above configuration, traffic classified in the C_VOICE and C_VIDEO class-maps shares the same priority queue. A second queue contains traffic classified in the class-default class-map. Therefore, the configuration only results in two queues, one shared priority queue and one queue for the class-default class. On most models of routers, only the class-default queue can be configured to use WFQ queuing for flows within the queue, while other queues use FIFO queuing for traffic within those queues.

5. Consider the following configuration:

class-map TRANSACTIONAL

  match protocol http

!

policy-map CBPOLICING

  class TRANSACTIONAL

    police 128000 conform-action set-dscp-transmit af11 exceed-action set-dscp-transmit af13 violate-action drop

!
interface serial 0/1
 service-policy input CBPOLICING

What type of class-based policing configuration is represented by this configuration?

A. Single rate, single bucket

B. Single rate, dual bucket

C. Dual rate, single bucket

D. Dual rate, dual bucket

Answer:

B

Explanation:

Cisco IOS supports single rate, single bucket; single rate, dual bucket; and dual rate, dual bucket policers. With a single rate policer, only a committed information rate (CIR) is specified, as in this question. With a dual rate policer, both a CIR and a peak information rate (PIR) are specified. Also, a single rate policer is a single bucket policer, unless the **violate** action is specified. If the violate action is specified, as it is in this question, the single rate policer uses two buckets, a Bc bucket and a Be bucket. However, a dual rate policer always uses two buckets, one bucket to transmit traffic at the CIR and one bucket to transmit traffic at the PIR.

6. You configure CB-Shaping by issuing the command **shape peak 8000 2000 2000**. This configuration shapes to what peak rate?

A. 4000 bps

B. 8000 bps

C. 16000 bps

D. 32000 bps

Answer:

C

Explanation:

In the syntax, the **8000** represents the Committed Information Rate (CIR). The first **2000** is the Committed Burst (Bc), and the second **2000** is the Excess Burst (Be). When configuring CB-Shaping, you can either shape to ?average? or shape to ?peak.? When shaping to average, traffic rates don't exceed the CIR. However, when shaping to peak, traffic rates can burst above the CIR, while some of that excess traffic could be dropped by the service provider. When shaping to peak, the peak shaping rate is calculated by the formula:

peak_rate = CIR * (1 + Be/Bc)

In this example: peak_rate = 8000 * (1 + 2000/2000) = 16,000 bps. Note that if the Bc and Be values are calculated by IOS rather than being statically configured, Bc will always equal Be, which means that the peak rate will be twice the CIR.

7. You are configuring Multilink PPP (MLP) as your Link Fragmentation and Interleaving (LFI) mechanism for a WAN link. Identify the correct statements regarding the configuration of MLP. (Choose 2)

A. The configuration of Multilink PPP requires at least two physical links (e.g. two serial interfaces)

B. The IP address is removed from any serial interface that makes up the MLP bundle

C. Any policy-map that was previously assigned to a physical interface should be reassigned to the multilink interface, that the physical interface is associated with, in order for the policy to take effect

D. The virtual multilink interface does not use an IP address. Rather, it uses the IP unnumbered feature which allows the multilink interface to share an IP address with the multilink bundle member that has the highest IP address

Answer:

B, C

Explanation:

Multilink PPP (MLP) is a Link Fragmentation and Interleaving (LFI) mechanism for PPP links. Interestingly, even though the term ?multilink? is in the title of this mechanism, MLP can be configured on a single link. Specifically, a virtual multilink interface is created. Then, one or more physical interfaces are added as members of a multilink bundle, all of which act as the single multilink interface. As a result, the virtual multilink interface is assigned an IP address, while the one or more physical interface member(s) do not have an IP address. Additionally, since the packets are logically transmitted over the virtual multilink interface, in order to apply a policy-map to the traffic using the virtual interface, the **service-policy** command should be applied to the virtual multilink interface, as opposed to the member interfaces.