(ONT) Optimizing Converged Cisco Networks

642-845

Exam B: Questions & Answers With Explanations

Study Tips

This product will provide you with questions and answers carefully compiled and written by our Expert Senior Certified Staff. Our practice questions are designed to help you learn the concepts behind the questions rather than be a strict memorization tool.

Important Note:

Please Read Carefully

Repeated readings of our Pass-Guaranteed.com Practice Exam will increase your comprehension. We constantly add to and update our Practice Exams with new questions, answers and explanations, so check that you have the latest version of this Practice Exam before you take your exam.

Update Notifications (Latest Version)

We are constantly reviewing our products. New material is added and old material is revised. Free Updates are available for 180 days after purchase. If you purchased a bundle, you will have Free Updates for 1 YEAR!

You can signup to our newsletter for instant notification whenever an update is released by becoming a Pass-Guaranteed.com member at: http://www.pass-guaranteed.com/log.htm

By becoming a Pass-Guaranteed.com member, you also get a chance to win a FREE Practice Exam of your choosing. We give away 3 Pass-Guaranteed.com Practice Exams every week to 3 lucky winners.

Pass-Guaranteed.com Product Specials

Pass-Guaranteed.com Custom Bundle Requests, cover all Pass-Guaranteed.com Products!!! You can visit our Special Bundle Discounts from Pass-Guaranteed.com or make your own Custom Bundle Request with Pass-Guaranteed.com here: http://www.pass-guaranteed.com/bundles.htm

***Pass-Guaranteed.com Custom Bundle Request Form let's you create your own Bundle Of Products!!!** You can select and group any of our products for your Custom Bundle and we will give you up to a **50% Discount** on your Custom Bundle Package. This includes our Practice Test Questions, Online Course Tutorials, Study Guides, Lab Scenarios and our Certified Online Instructor service.*

Please visit: http://www.pass-guaranteed.com/custom-request.htm If you would like to purchase a Custom Bundle from Pass-Guaranteed.com.

QUESTION: 1
You need to implement the proper IOS tools to ensure that VOIP works over the Pass network.

Which queuing and compression mechanisms are needed to effectively use the available bandwidth for voice traffic? (Select TWO)

A. Priority Queuing (PQ) or Custom Queuing (CQ)
B. Real-Time Transport Protocol (RTP) header compression
C. Low Latency Queuing (LLQ)
D. Class-Based Weighted Fair Queuing (CBWFQ)
E. TCP header compression
F. UDP header compression

Answer: D, E

Explanation:
1. Class-based weighted fair queuing (CBWFQ) extends the standard WFQ functionality to provide support for user-defined traffic classes. By using CBWFQ, network managers can define traffic classes based on several match criteria, including protocols, access control lists (ACLs), and input interfaces. A FIFO queue is reserved for each class, and traffic belonging to a class is directed to the queue for that class. More than one IP flow, or "conversation", can belong to a class.

Once a class has been defined according to its match criteria, the characteristics can be assigned to the class. To characterize a class, assign the bandwidth and maximum packet limit. The bandwidth assigned to a class is the guaranteed bandwidth given to the class during congestion.

CBWFQ assigns a weight to each configured class instead of each flow. This weight is proportional to the bandwidth configured for each class. Weight is equal to the interface bandwidth divided by the class bandwidth. Therefore, a class with a higher bandwidth value will have a lower weight.

By default, the total amount of bandwidth allocated for all classes must not exceed 75 percent of the available bandwidth on the interface. The other 25 percent is used for control and routing traffic.

The queue limit must also be specified for the class. The specification is the maximum number of packets allowed to accumulate in the queue for the class. Packets belonging to a class are subject to the bandwidth and queue limits that are configured for the class.

2. TCP/IP header compression subscribes to the Van Jacobson Algorithm defined in RFC 1144. TCP/IP header compression lowers the overhead generated by the disproportionately large TCP/IP headers as they are transmitted across the WAN. TCP/IP header compression is protocol-specific and only compresses the /IP header. The Layer 2 header is still intact and a packet with a compressed TCP/IP header can still travel across a WAN link.

TCP/IP header compression is beneficial on small packets with few bytes of data such as Telnet. Cisco's header compression supports Frame Relay and dial-on-demand WAN link protocols. Because of processing overhead, header compression is generally used at lower speeds, such as 64 kbps links.

QUESTION: 2
While planning the new Pass VOIP network, you need to determine the size of the WAN links to use. To do this, you need to calculate the bandwidth required by each call.

Which three pieces of information are used to calculate the total bandwidth of a VoIP call? (Select THREE)

A. The serialization of the interface
B. The quantization
C. The TCP overhead
D. The packetization size
E. The UDP overhead
F. The packet rate

Answer: D, E, F

Explanation:
Packet rate: Packet rate specifies the number of packets sent in a certain time interval. The packet rate is usually specified in packets per second (pps). Packet rate is the multiplicative inverse of the packetization period. The packetization period is the amount of voice (time) that will be encapsulated per packet, and is usually specified in milliseconds.

Packetization size: Packetization size specifies the number of bytes that are needed to represent the voice information that will be encapsulated per packet. Packetization size depends on the packetization period and the bandwidth of the codec used.

IP overhead: IP overhead specifies the number of bytes added to the voice information during IP encapsulation. When voice is encapsulated into Real-Time Transport Protocol (RTP), User Datagram Protocol (UDP), and IP, the IP overhead is the sum of all these headers.

Data link overhead: Data-link overhead specifies the number of bytes added during data-link encapsulation. The data-link overhead depends on the used data-link protocol, which can be different per link.

Tunneling overhead: Tunneling overhead specifies the number of bytes added by any security or tunneling protocol, such as 802.1Q tunneling, IPsec, Generic Route Encapsulation (GRE), or Multiprotocol Label Switching (MPLS). This overhead must be considered on all links between the tunnel source and the tunnel destination.

QUESTION: 3
Analog interfaces are being utilized in a number of the Pass VOIP gateways. Which two voice gateway analog-interface statements are true? (Select TWO)

A. An analog fax machine can connect to a Foreign Exchange Office (FXO) interface.
B. A router can use a Foreign Exchange Office (FXO) interface to connect to a PSTN.
C. A router can use a Foreign Exchange Station (FXS) interface to connect to a PBX.
D. An analog telephone can connect to a Foreign Exchange Station (FXS) interface.

Answer: B, D

Explanation:
Gateways use different types of interfaces to connect to analog devices, such as phones, fax machines, or PBX or public switched telephone network (PSTN) switches. Analog interfaces used at the gateways include these three types:

FXS: The FXS interface connects to analog end systems, such as analog phones or analog faxes, which on their side use the FXO interface. The router FXS interface behaves like a PSTN or a PBX, serving phones, answering machines, or fax machines with line power, ring voltage, and dial tones. If a PBX uses an FXO interface, it can also connect to a router FXS interface. In this case, the PBX acts like a phone.

FXO: The FXO interface connects to analog systems, such as a PSTN or a PBX, which on their side use the FXS interface. The router FXO interface behaves like a phone, getting line power, ring voltage, and dial tones from the other side. As mentioned, a PBX can also use an FXO interface toward the router (which will then use an FXS interface), if the PBX takes the role of the phone.

QUESTION: 4
DRAG DROP
Drag each term next to the appropriate definition. Upon completion, there will be one term left unused.

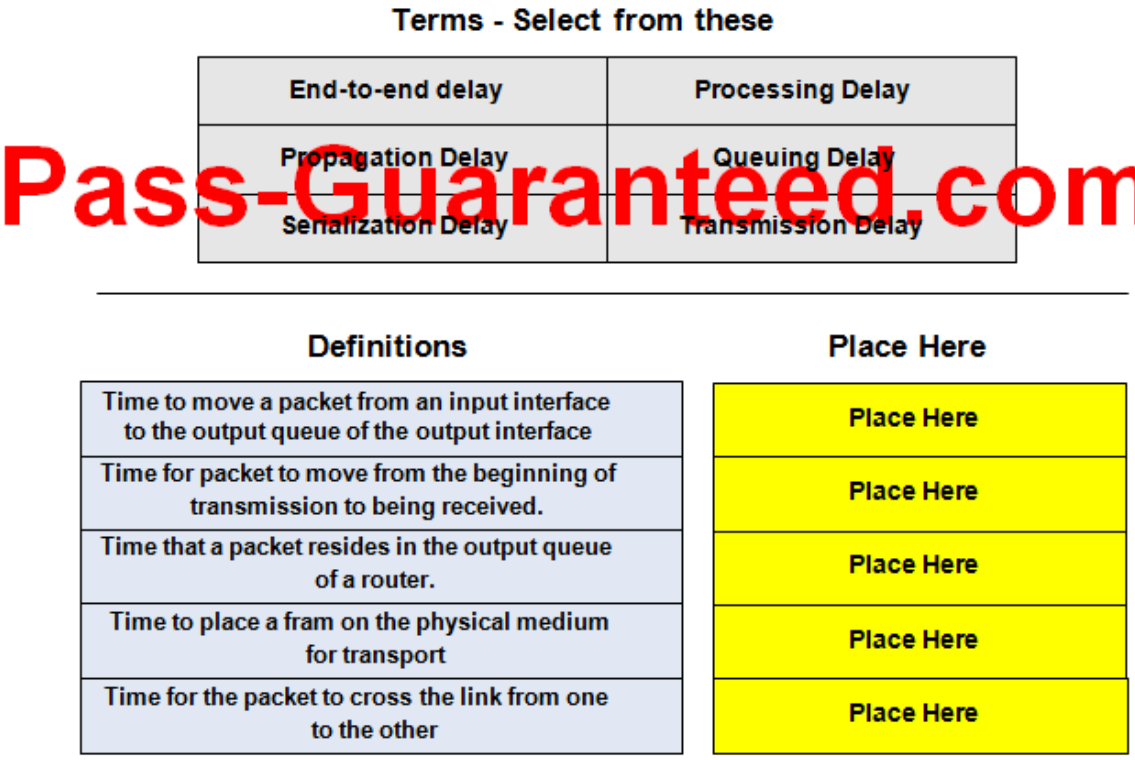

Answer:

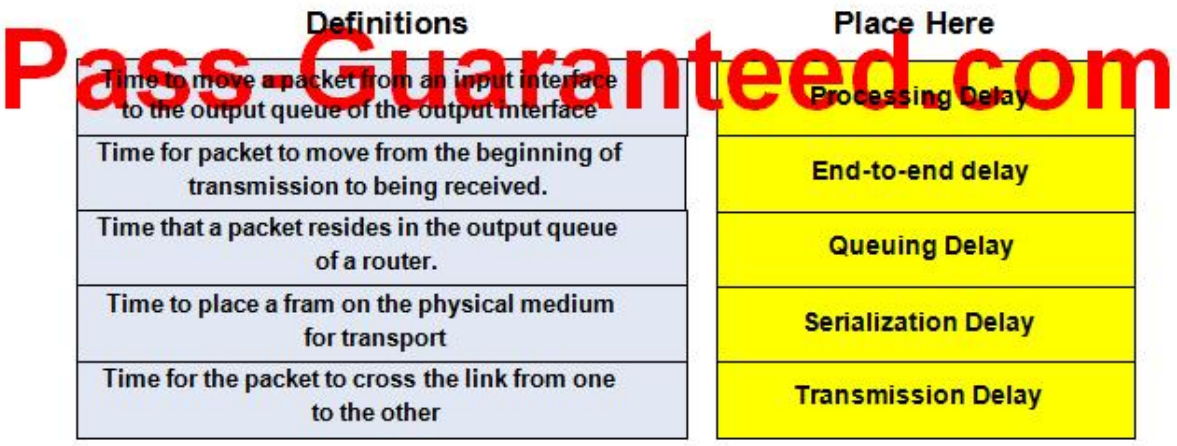

QUESTION: 5
You need to classify different packets within the Pass network so that they can be marked.

What are three traffic descriptors typically used to categorize traffic into different classes? (Select THREE)

A. DSCP
B. DLCI
C. Media type
D. IP precedence
E. Incoming interface
F. Outgoing interface

Answer: A, D, E

Explanation:
Classification is the process of identifying traffic and categorizing that traffic into classes. Classification uses a traffic descriptor to categorize a packet within a specific group to define that packet. Typically used traffic descriptors include these:

1. Incoming interface
2. IP precedence
3. differentiated services code point (DSCP)
4. Source or destination address
5. Application

After the packet has been classified or identified, the packet is then accessible for quality of service (QoS) handling on the network. Using classification, network administrators can partition network traffic into multiple classes of service (CoSs). When traffic descriptors are used to classify traffic, the source implicitly agrees to adhere to the contracted terms and the network promises QoS. Various QoS mechanisms, such as traffic policing, traffic shaping, and queuing techniques, use the traffic descriptor of the packet (that is, the classification of the packet) to ensure adherence to that agreement.

Classification should take place at the network edge, typically in the wiring closet, within IP phones, or at network endpoints. It is recommended that classification occur as close to the source of the traffic as possible.

QUESTION: 6
You need to implement QoS on the Pass network.

Which two queuing methods will allow a percentage of the available bandwidth to be allocated to each queue? (Select TWO)

A. Weighted Fair Queuing (WFQ)
B. Priority Queuing (PQ)
C. Class-based WFQ (CBWFQ)
D. Custom Queuing (CQ)
E. Low Latency Queuing (LLQ)
F. First-In, First-Out Queuing (FIFO)

Answer: C, E

Explanation:
1. CBWFQ:
Class-based weighted fair queuing (CBWFQ) extends the standard WFQ functionality to provide support for user-defined traffic classes. By using CBWFQ, network managers can define traffic classes based on several match criteria, including protocols, access control lists (ACLs), and input interfaces. A FIFO queue is reserved for each class, and traffic belonging to a class is directed to the queue for that class. More than one IP flow, or "conversation", can belong to a class.

Once a class has been defined according to its match criteria, the characteristics can be assigned to the class. To characterize a class, assign the bandwidth and maximum packet limit. The bandwidth assigned to a class is the guaranteed bandwidth given to the class during congestion.

CBWFQ assigns a weight to each configured class instead of each flow. This weight is proportional to the bandwidth configured for each class. Weight is equal to the interface bandwidth divided by the class bandwidth. Therefore, a class with a higher bandwidth value will have a lower weight.

By default, the total amount of bandwidth allocated for all classes must not exceed 75 percent of the available bandwidth on the interface. The other 25 percent is used for control and routing traffic.

The queue limit must also be specified for the class. The specification is the maximum number of packets allowed to accumulate in the queue for the class. Packets belonging to a class are subject to the bandwidth and queue limits that are configured for the class.

2. LLQ
The Low Latency Queuing (LLQ) feature provides strict priority queuing for class-based weighted fair queuing (CBWFQ), reducing jitter in voice conversations. Configured by the priority command, strict priority queuing gives delay-sensitive data, such as voice, preferential treatment over other traffic. With this feature, delay-sensitive data is sent first, before packets in other queues are treated. LLQ is also referred to as priority queuing/class-based weighted fair queuing (PQ/CBWFQ) because it is a combination of the two techniques.

For CBWFQ, the weight for a packet belonging to a specific class is derived from the bandwidth assigned to the class during configuration. Therefore, the bandwidth assigned to the packets of a class determines the order in which packets are sent. All packets are serviced equally, based on weight. No class of packets may be granted strict priority. This scheme poses problems for voice and video traffic that is largely intolerant of

delay, especially variation in delay. For voice traffic, variations in delay introduce irregularities of transmission, which manifest as jitter in the conversation.

To enqueue a class of traffic to the strict priority queue, configure the priority command for the class after specifying the class within a policy map. Classes to which the priority command is applied are considered priority classes. Within a policy map, give one or more classes priority status. When multiple classes within a single policy map are configured as priority classes, all traffic from these classes is enqueued to the same, single, strict priority queue and they will contend with each other for bandwidth

QUESTION: 7
DRAG & DROP
Using the fewest commands possible, drag the commands on the left to the blanks on the right to configure and apply a QoS policy that guarantees that voice packets receive 20 percent of the bandwidth on the S0/1/0 interface.

Answer:

Explanation:
Complete the following steps to implement the QoS.

Step 1: Configure traffic classification by using the class-map command.

A class map is created using the class-map global configuration command. Class maps are identified by case-sensitive names. Each class map contains one or more conditions that determine whether the packet belongs to the class. There are two ways of processing conditions when there is more than one condition in a class map:

Match all: All conditions have to be met to bind a packet to the class.

Match any: At least one condition has to be met to bind the packet to the class. The default match strategy of class maps is match all.

Step 2: Configure traffic policy by associating the traffic class with one or more QoS features using the policy-map command.

The name of a traffic policy is specified in the policy-map command (for example, issuing the policy-map class1 command would create a traffic policy named class1). After you issue the policy-map command, you enter policy-map configuration mode. You can then enter the name of a traffic class. Here is where you enter QoS features to apply to the traffic that matches this class.

Step 3: Attach the traffic policy to inbound or outbound traffic on interfaces, subinterfaces, or virtual circuits by using the service-policy command.

Using the service-policy command, you can assign a single policy map to multiple interfaces or assign multiple policy maps to a single interface (a maximum of one in each direction, inbound and outbound). A service policy can be applied for inbound or outbound packets.

QUESTION: 8
You want to implement a congestion avoidance mechanism within the Pass network.

Which QoS tool is used to reduce the level of congestion in the queues by selectively dropping packets?

A. Weighted Random Early Detection (WRED)
B. Low Latency Queuing (LLQ)
C. Class-based Weighted Fair Queuing (CBWFQ)
D. Modified Deficit Round Robin (MDRR)
E. None of the above

Answer: A

Explanation:
Weighted random early detection (WRED) combines RED with IP precedence or DSCP and performs packet dropping based on IP precedence or DSCP markings. As with RED, WRED monitors the average queue length in the router and determines when to begin discarding packets based on the length of the interface queue. When the average queue length exceeds the user-specified minimum threshold, WRED begins to randomly drop packets with a certain probability. If the average length of the queue continues to increase so that it becomes larger than the user-specified maximum threshold, WRED reverts to a tail-drop packet-discard strategy, in which all incoming packets are dropped. The idea behind using WRED is to maintain the queue length at a level somewhere below the maximum threshold and to implement different drop policies for different classes of traffic. WRED can selectively discard lower-priority traffic when the interface becomes congested and can provide differentiated performance characteristics for different classes of service. WRED can also be configured to produce nonweighted RED behavior.

QUESTION: 9
Four distinct packet queues in a Pass router are displayed below:

Study the exhibit carefully. Packet-based WRR (not byte-count WRR) is being used to control the output on an interface with four queues (A, B, C, D) configured. Each queue has an assigned weight of A=4, B=2, C=1, and D=1. If the queuing algorithm begins with Queue A and with packets placed into the four queues as shown in the exhibit, in which order will packets be selected from the queues for transmission?

A. 1, 3, 8, 2, 9, 4, 5, 10, 6, 7
B. 1, 2, 4, 5, 3, 9, 6, 7, 8 10
C. 1, 3, 8, 2, 4, 5, 9, 6, 7, 10
D. 1, 3, 8, 2, 9, 10, 4, 6, 5, 7
E. 1, 3, 2, 9, 4, 6, 5, 7, 8, 10
F. None of the above

Answer: A

Explanation:
In WRR, packets are accessed round-robin style, but queues can be given priorities called "weights." For example, in a single round, four packets from a high-priority class might be dispatched, followed by two from a middle-priority class, and then one from a low-priority class.

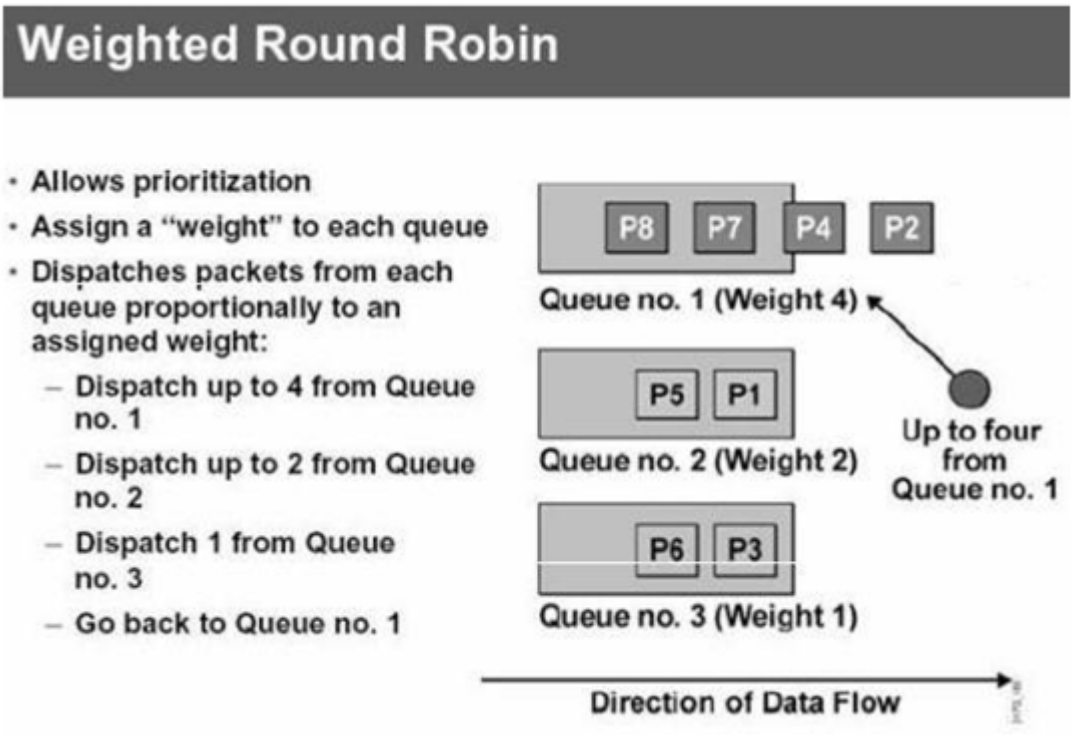

Some implementations of the WRR algorithm provide prioritization by dispatching a configurable number of bytes each round rather than a number of packets. The Cisco custom queuing (CQ) mechanism is an example of this implementation.

QUESTION: 10
You need to consider the advantages and disadvantages of using traffic shaping versus traffic policing within your network.

Which statement about traffic policing and which statement about traffic shaping are true? (Select TWO)

A. Traffic policing drops excess traffic in order to control traffic flow within specified rate limits.

B. Traffic shaping buffers excess traffic so that the traffic stays within the desired rate.

C. Traffic policing can cause UDP retransmissions when traffic in excess of specified limits is dropped.

D. A need for traffic shaping occurs when a service provider must rate-limit the customer traffic to T1 speed on an OC-3 connection.

E. Traffic policing and traffic conditioning are mechanisms that are used in an edge network to guarantee QoS.

Answer: A, B

Explanation:

Policing can be applied to either the inbound or outbound direction, while shaping can be applied only in the outbound direction. Policing drops nonconforming traffic instead of queuing the traffic like shaping. Policing also supports marking of traffic. Traffic policing is more efficient in terms of memory utilization than traffic shaping because no additional queuing of packets is needed. Both traffic policing and shaping ensure that traffic does not exceed a bandwidth limit, but each mechanism has different impacts on the traffic:

Policing drops packets more often, generally causing more retransmissions of connection-oriented protocols, such as TCP.

Shaping adds variable delay to traffic, possibly causing jitter. Shaping queues excess traffic by holding packets in a shaping queue. Traffic shaping is used to shape the outbound traffic flow when the outbound traffic rate is higher than a configured rate. Traffic shaping smoothes traffic by storing traffic above the configured rate in a shaping queue. Therefore, shaping increases buffer utilization on a router and causes unpredictable packet delays. Traffic shaping can also interact with a Frame Relay network, adapting to indications of Layer 2 congestion in the WAN. For example, if the backward explicit congestion notification (BECN) bit is received, the router can lower the rate limit to help reduce congestion in the Frame Relay network.

QUESTION: 11
Voice activity detection (VAD) suppresses the transmission of silence patterns which can mean more bandwidth will become available over a given link.

On average, and assuming that a link carries at least 24 calls, what percentage of total bandwidth could VAD save?

A. 15
B. 45
C. 55
D. 5
E. 35
F. 25
G. None of the above

Answer: E

Explanation:
In a circuit-switched telephone network, because of the nature of the network, the bandwidth of a call is permanently available and dedicated to that call. There is no way to take advantage of speech pauses, one-way audio transmission, or similar instances when a link is not being utilized. In a packet network, however, VAD can take advantage of the fact that one-third of the average voice call consists of silence. VAD detects silence, for instance, caused by speech pauses or by one-way audio transmission while a caller is listening to music on hold (MoH) when being transferred. VAD suppresses the transmission of silence and, therefore, saves bandwidth. The amount of bandwidth that can be saved by VAD depends on several factors:

1. Type of audio: During a human conversation, the two parties do not generally talk at the same time. When MoH is played, the call usually turns into a one-way call. Because of the constantly playing music, no bandwidth can be saved in this direction of the call. However, the caller listening to the music does not send any audio and no packets have to be transmitted while the call is on hold.

2. Level of background noise: VAD needs to detect silence to be able to perform silence suppression. If the background noise is too high, VAD cannot detect silence and continues the transmission.

3. Others: Differences in the language and character of speakers have an impact to the amount of silence in a call. Some calls, such as conferences or broadcasts where only one or a few participants are speaking and most of the participants are listening, allow higher bandwidth savings than other calls.

On average, the use of VAD can save about 35 percent of bandwidth. Because of the factors mentioned, there is considerable deviation per individual call. Therefore, the average of 35 percent assumes a certain statistical distribution of call types, which is usually achieved only if a link carries at least 24 calls. If you are calculating bandwidth for fewer calls, you should not take VAD into account.

QUESTION: 12
QoS pre-classification is being used on the Pass VPN.

Which three characteristics of the traffic flow are taken into consideration when the QoS-for-VPNs feature (QoS pre-classify) provides packet classification and applies appropriate QoS service on tunnel interfaces? (Select THREE)

A. IP precedence bits
B. Destination IP address
C. DE bits
D. DSCP bits
E. Source IP address
F. Original port numbers

Answer: B, E, F

Explanation:
The qos pre-classify command mechanism allows Cisco routers to make a copy of the inner IP header and to run a QoS classification before encryption, based on fields in the inner IP header. If the classification policy matches on the ToS byte, it is not necessary to use the qos pre-classify command, because the ToS value is copied to the outer header by default. In addition, a simple QoS policy that sorts traffic into classes based on IP precedence can be created. However, differentiating traffic within a class and separating it into multiple flow-based queues requires the qos pre-classify command.

Alternatively, traffic may need to be classified based on values other than IP precedence or DSCP. For example, packets may need to be classified based on IP flow or Layer 3 information, such as source and destination IP address. To do so, use the QoS for VPNs feature enabled with the qos pre-classify command.