PassQueen
a.k.a dumpsqueen.com

Nokia Quality of Service

Alcatel-Lucent 4A0-107

Version Demo

Total Demo Questions: 26

Total Premium Questions: 260

Buy Premium PDF

https://passqueen.com

support@passqueen.com

QUESTION NO: 1

Which of the following statements regarding buffer pool allocation are FALSE? (Choose two)

A. Each network egress port receives its own buffer pool.

B. Each SAP-egress port receives its own buffer pool.

C. All SAP-ingress ports on the same MDA share a buffer pool.

D. All the network ingress ports on the same MDA share a buffer pool.

E. Each network ingress port on the same MDA receives its own buffer pool.

ANSWER: C E

Explanation:

The false statements are “All SAP-ingress ports on the same MDA share a buffer pool” and “Each network ingress port on the same MDA receives its own buffer pool.” In the SR OS QoS buffer-allocation model, SAP ingress is allocated on a per-SAP basis rather than by placing every SAP ingress entity on an MDA into one common buffer pool. This provides SAP-level buffer resources for ingress traffic handling and congestion management. Consequently, it is not accurate to describe SAP-ingress buffer allocation as a single shared pool for all SAPs on the MDA.

Conversely, network-ingress buffering is an MDA-level resource shared among the network ingress ports associated with that MDA. A network ingress port therefore does not receive an individually dedicated buffer pool simply because it is a separate port. The shared MDA resource model is important when considering the effects of aggregate ingress traffic and congestion across network ports. Egress buffer allocation differs from this ingress model: network egress and SAP egress use port-specific buffer-pool allocation. These allocation distinctions are part of the platform QoS architecture and should be considered when designing queue policies and assessing congestion behavior. See the Nokia 7750 SR Quality of Service Guide and the SR OS QoS overview.

QUESTION NO: 2

Which of the following statements regarding scheduling are TRUE? (Choose two)

A. Scheduler policies cannot be children of an egress port-scheduler.

B. HQoS provides a static method of allocating scheduling resources to each queue and/or child.

C. Scheduler policies can be applied to the MDA to schedule network ingress traffic.

D. Every time a scheduler-policy is assigned to a SAP, a new hardware scheduler is created to service the SAP.

E. A scheduler applied to multiple SAPs within the same service will result in a single hardware scheduler operating across all SAPs within the service.

ANSWER: C D

Explanation:

Scheduler policies can be applied to the MDA to schedule network ingress traffic. In SR OS, an ingress network scheduler policy is associated with an MDA and provides the hierarchy used to distribute the available ingress scheduling resources among network queues and child scheduler entities. This allows ingress traffic scheduling to be implemented at the hardware context where network ingress queues are instantiated, rather than only at an individual service access point. The scheduler-policy configuration supplies the scheduling hierarchy and operational parameters used by the forwarding hardware.

Every time a scheduler-policy is assigned to a SAP, a new hardware scheduler is created to service the SAP. A scheduler policy is a reusable configuration template, but its assignment to a SAP creates a separate scheduler instance for that SAP.

Consequently, each SAP receives its own hardware scheduling context and can independently enforce the configured hierarchy, rates, priorities, and queue treatment. This per-SAP instantiation is fundamental to hierarchical QoS because it preserves isolation while allowing the same policy definition to be consistently reused across many SAPs. Nokia's SR OS QoS documentation describes scheduler-policy use for ingress network scheduling and SAP-based hierarchical scheduling; see the Nokia SR OS documentation portal and the 7750 SR OS QoS Guide.

QUESTION NO: 3

Which of the following statements defines the committed information rate?

A. It is the rate at which a traffic stream is offered to a queue.

B. It is the rate at which packets are forwarded out of a queue.

C. It is the minimum rate at which packets may be forwarded out of a queue.

D. It is the maximum rate at which packets may be forwarded out of a queue.

ANSWER: C

Explanation:

"It is the minimum rate at which packets may be forwarded out of a queue." correctly describes the committed information rate in the context of an Alcatel-Lucent/Nokia SR OS QoS queue. The CIR is the committed portion of queue bandwidth: when there is congestion and scheduler resources are being shared, the queue is eligible to receive service sufficient to meet its configured committed rate, subject to the applicable parent scheduler and port resources being available. This gives committed traffic a predictable baseline forwarding service and allows the QoS scheduler to distinguish guaranteed bandwidth from excess bandwidth that may be used when capacity is not otherwise needed.

In SR OS queueing and scheduling, CIR is configured as a rate associated with the queue and is used by the scheduler's bandwidth-allocation behavior. The associated PIR represents the upper forwarding limit, while CIR represents the committed minimum service objective. This queue-oriented use is consistent with the broader IETF definition of committed information rate as the configured committed traffic rate used for traffic conditioning and QoS behavior. Nokia's SR OS QoS documentation is available through the Nokia Documentation Center; see also the IETF definition in RFC 2697.

QUESTION NO: 4

Click the exhibit button below.

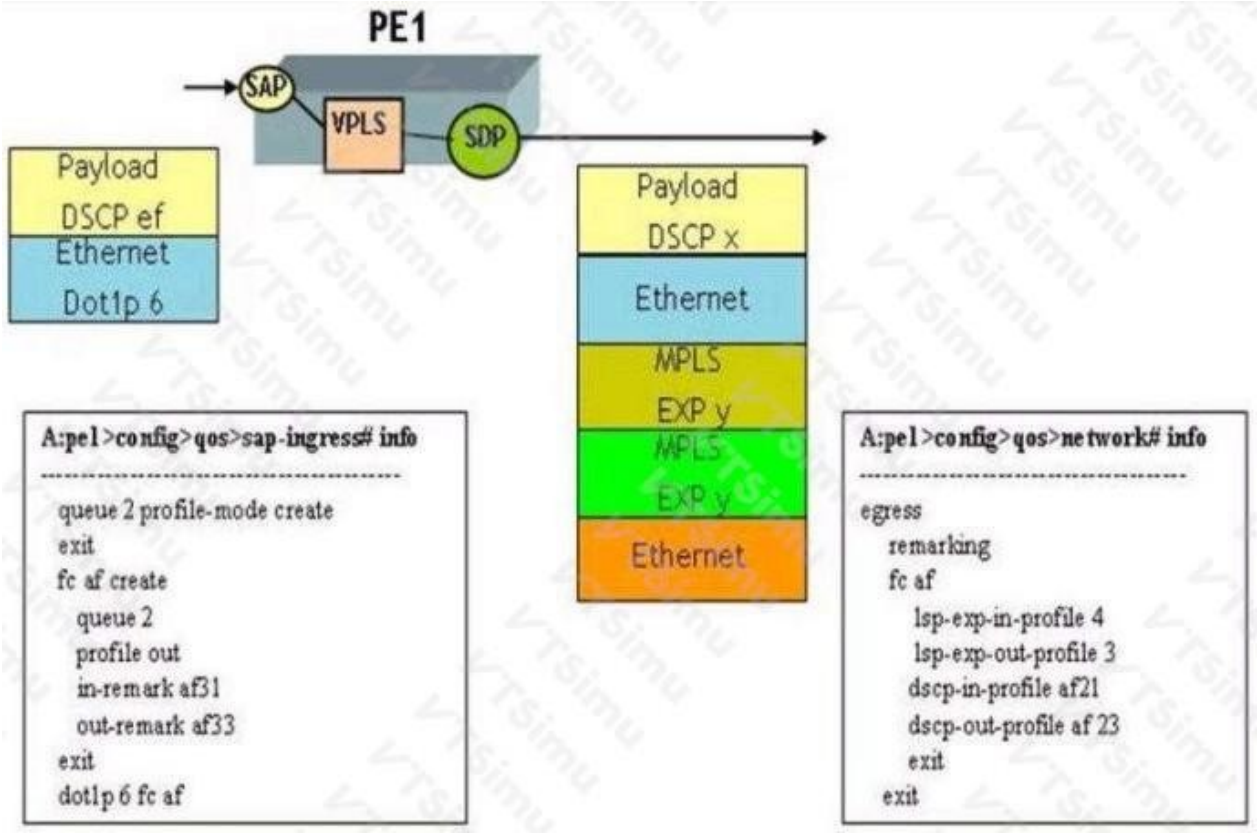

At router PE1, customer packets are arriving tagged with a dotlp value of 6 and marked with a DSCP value of EF.

Based on the configuration shown below for a VPLS service, what will be the DSCP and EXP marking for the packet egressing the network port on this PE router? (Choose two)

A. The DSCP value will be set to af21.

B. The DSCP value will be set to af23.

C. The DSCP value will remain as ef.

D. The EXP value, in the MPLS headers, will be set to 4.

E. The EXP value, in the MPLS headers, will be set to 3.

ANSWER: C E

Explanation:

The DSCP value will remain as ef, and the EXP value, in the MPLS headers, will be set to 3. In a VPLS service, the customer Ethernet frame and its IP payload are transported across the MPLS network. The configured QoS handling does not rewrite the customer IP header's DSCP field in this case, so the EF marking carried in the payload remains intact. DSCP EF is the expedited-forwarding per-hop behavior, defined for low-loss, low-latency traffic in RFC 3246.

For transport across the provider MPLS network, the router applies the configured network QoS marking to the MPLS label stack. The ingress dot1p value of 6 is classified according to the service QoS policy, and the resulting forwarding treatment is mapped by the network policy to MPLS EXP value 3. This MPLS EXP/Traffic Class value is what provider routers use to select the applicable forwarding behavior while the VPLS frame is label-switched. MPLS support for carrying differentiated-services information in the label-stack traffic-class field is specified in RFC 3270. Thus, the payload retains EF while the newly imposed MPLS headers carry EXP 3.

QUESTION NO: 5

Forwarding sub-classes are configured in the context of which QoS policy type?

A. Network

B. Scheduler

C. Slope

D. SAP-ingress

E. Network queue

ANSWER: D

Explanation:

Forwarding sub-classes are configured in a SAP-ingress QoS policy. On SR OS, a SAP-ingress policy performs classification for traffic entering a service access point and associates the classified traffic with QoS treatment information. This classification includes assigning a forwarding class and, where finer differentiation is required within that forwarding class, a forwarding subclass. The forwarding subclass provides an additional classification level that can be carried forward for subsequent QoS handling, allowing traffic that shares a forwarding class to retain distinct treatment characteristics as defined by the service ingress policy. This fits the SAP-ingress policy's role as the service-side ingress classification and marking point, before packets are handled by service queues and later forwarding stages. Nokia's SR OS QoS documentation describes SAP ingress QoS policies as the policies used for ingress traffic classification and QoS parameter assignment at SAPs. See the Nokia SR OS QoS overview and the SAP ingress QoS documentation for the policy model and classification behavior.

QUESTION NO: 6

Which of the following statements regarding Random Early Detection (RED) are TRUE? (Choose three)

A. RED can discard packets before a queue is full.

B. RED can use probabilistic packet discard.

C. RED can help prevent TCP global synchronization.

D. RED guarantees that a queue will never experience tail drop.

E. RED determines the transmission order of egress queues.

ANSWER: A B C

Explanation:

RED is an active queue-management mechanism that can begin dropping packets before a queue is full. By dropping packets probabilistically as queue occupancy rises, RED provides congestion signals before tail drop occurs. For TCP traffic, this can help avoid global synchronization, where many TCP flows reduce their transmission rates at the same time after a common tail-drop event.

RED does not guarantee that packets will never be dropped when a queue becomes full. It is also not a scheduling algorithm; packet scheduling determines the order in which queues receive service. The IETF describes the role of active queue management in RFC 2309.

QUESTION NO: 7

Which of the following are entities of the Nokia 7750 SR on which you can NOT apply a networkqueue policy? (Choose three)

A. Router interfaces

B. Ingress MDA

C. Ingress Ethernet network ports

D. Egress Ethernet network ports

E. Service SAPs

ANSWER: A C E

Explanation:

Network queue policies are specifically intended for network-facing hardware queueing contexts on the 7750 SR. Their supported attachment points include an ingress MDA, where ingress network queue resources are controlled, and an egress Ethernet network port, where the policy defines the treatment and resource allocation of egress network queues. Consequently, **Router interfaces**, **Ingress Ethernet network ports**, and **Service SAPs** are the three entities on which a network queue policy cannot be applied. Router-interface QoS behavior is configured through the appropriate network QoS policy mechanisms rather than a network queue policy attachment. An ingress Ethernet port is not the egress port queue context to which a network queue policy is attached. Service SAPs use service/SAP QoS policy constructs, preserving the separation between service traffic QoS and network port/MDA queue resources. This distinction is important because the policy type must match both the traffic domain and the hardware point at which queue scheduling and buffering occur. Nokia's SR OS QoS documentation describes the supported network queue policy application contexts and their role in network QoS configuration. See the Nokia SR OS documentation portal and the 7750 SR documentation set.

QUESTION NO: 8

When deploying a Nokia MPLS network with QoS, how are the EXP bits set?

A. The first 3 bits of the DSCP value are copied directly into the EXP field on the first network egress.

B. The EXP bits are set on the first network egress, based on a combination of the SAP-ingress classification and the network policy applied.

C. The 3 priority bits are copied directly from the dot1p header into the EXP field of the first MPLS header.

D. At SAP ingress, the default DSCP map classifies traffic into an internal forwarding class and profile; the network QoS policy then maps those to EXP bits at first network egress.

ANSWER: B D

Explanation:

In Nokia SR OS QoS processing, classification at SAP ingress first assigns the packet an internal forwarding class and profile. For IP traffic, this classification can use the default DSCP-to-forwarding-class map when no explicit classifier entry overrides it. These internal QoS attributes, rather than a direct bit-for-bit copy of the received IP or Ethernet priority marking, are carried through the service and used for subsequent network treatment.

When the packet is sent toward the MPLS core at its first network egress, the applied network QoS policy maps the forwarding class and profile to the MPLS Traffic Class field, historically called the EXP field. This makes "The EXP bits are set on the first network egress, based on a combination of the SAP-ingress classification and the network policy applied" correct. The statement describing default DSCP mapping at SAP ingress is also correct when completed to show that the resulting forwarding class and profile are then mapped by the network QoS policy to EXP bits. This separation of ingress classification from network MPLS marking is fundamental to SR OS QoS behavior. See the Nokia SR OS Quality of Service Guide and the MPLS Traffic Class definition in RFC 5462.

QUESTION NO: 9

Which of the following are commonly included as QoS objectives in a service-level agreement? (Choose three)

A. Packet delay

B. Packet delay variation

C. Packet loss

D. The number of forwarding classes configured on a router

E. The DSCP value used by every customer application

ANSWER: A B C

Explanation:

Packet delay, **packet delay variation**, and **packet loss** are common QoS objectives. They describe measurable outcomes experienced by traffic delivered through the network. Delay and delay variation are particularly important for real-time applications, while packet-loss objectives help define the expected reliability of the service.

A forwarding class is an internal traffic-treatment construct, and a DSCP value is a packet marking used to select treatment. Both may be referenced in a service design, but neither alone represents a measured service-performance objective. The IP performance-metrics framework is described in RFC 2330.

QUESTION NO: 10

Which of the following statements regarding the MPLS Traffic Class field are TRUE? (Choose three)

A. The MPLS Traffic Class field is 3 bits long.

B. The MPLS Traffic Class field is present in each label-stack entry.

C. The MPLS Traffic Class field was historically called the EXP field.

D. The MPLS Traffic Class field is 6 bits long.

E. The MPLS Traffic Class field is carried only in the bottom label.

F. The MPLS Traffic Class field is part of the IP header.

ANSWER: A B C

Explanation:

The MPLS Traffic Class field is three bits long, is present in each MPLS label-stack entry, and was historically called the EXP field. A network device can use the Traffic Class value to classify MPLS packets and select the appropriate forwarding treatment. Since each label-stack entry includes its own Traffic Class field, QoS information can be carried through label operations such as push, swap, and pop when the applicable QoS policy is configured.

The MPLS label-stack encoding is defined in RFC 3032, and the renaming of the EXP field to the Traffic Class field is described in RFC 5462.

QUESTION NO: 11

It is necessary for Tier 1 schedulers to obtain bandwidth from a higher tier scheduler.

A. TRUE

B. FALSE

ANSWER: B

Explanation:

FALSE is correct. In the SR OS hierarchical scheduling model, a Tier 1 scheduler is the top-level scheduler in a scheduler hierarchy. It can be associated with the applicable scheduling context, such as a port or other top-level attachment point, and

is not required to be a child of a scheduler in a higher tier. Its available rate is determined by its configured scheduling and shaping parameters together with the capacity available at that attachment point.

Higher scheduler tiers are used only when an operator builds additional levels of hierarchy beneath the Tier 1 level to divide available bandwidth among groups of traffic, services, customers, or child schedulers. Such a hierarchy is optional and is configured when more granular bandwidth allocation or contention management is needed. A Tier 1 scheduler can therefore operate as the root of the configured scheduler tree; it does not need to obtain bandwidth from another scheduler tier. This distinction is important when designing QoS policies, because the parent-child bandwidth relationship begins only when subordinate scheduler tiers are explicitly configured below the top-level scheduler.

See the Nokia SR OS documentation portal and its Quality of Service documentation for hierarchical scheduler concepts and configuration details.

QUESTION NO: 12

Which of the following CANNOT be used as external marking bits?

A. IP Precedence

B. Dot1p

C. Forwarding class

D. EXP

ANSWER: C

Explanation:

Forwarding class is correct because it is an internal QoS classification value maintained by the Nokia SR OS router, rather than a set of bits carried in an incoming or outgoing packet header. The router assigns a forwarding class after classifying traffic, using configured match criteria such as packet markings, interfaces, or other policy rules. It then uses that internal class to select QoS treatment, including queue mapping, priority handling, scheduling behavior, and remarking actions. Because forwarding class is local to the router's QoS policy and is not a standardized field encoded on the wire, it cannot serve as external marking bits. External markings are packet-visible header fields that devices can read and write as traffic traverses the network; forwarding class can be derived from such a marking or used to determine a marking to apply, but it is not itself transmitted with the packet. Nokia's SR OS QoS documentation describes forwarding classes as internal classification constructs and distinguishes them from packet-header QoS markings. See the Nokia SR OS QoS overview and the Nokia SR OS forwarding-class documentation.

QUESTION NO: 13

Which of the following statements about behavior aggregate classification is TRUE?

A. It sets both the forwarding class and the packetâ€™s dot1p, lsp-exp, precedence, or dscp field to the same value.

B. It groups all packets together into a single aggregate stream, and gives all packets the same treatment.

C. It copies the three most significant bits in the packetâ€™s dscp field into the forwarding class.

D. It assigns a forwarding class based on the packetâ€™s dot1p, lsp-exp, precedence, or dscp value.

ANSWER: D

Explanation:

“It assigns a forwarding class based on the packet's dot1p, lsp-exp, precedence, or dscp value.” is correct because behavior aggregate (BA) classification bases forwarding treatment on a packet's existing QoS marking. Depending on the packet encapsulation and the applicable QoS policy, SR OS can use an IEEE 802.1p priority value, MPLS label EXP/Traffic Class value, IP precedence value, or DSCP value to identify the behavior aggregate. The policy's classification mapping then associates that marked aggregate with a forwarding class, and may also determine the packet's profile, such as in-profile or

out-of-profile. This permits packets bearing the same relevant marking to receive a consistent per-hop treatment without inspecting transport-layer flows or applying a multi-field match. The forwarding class is therefore selected from the configured QoS interpretation of the received marking, rather than being inferred from a complete flow definition. Nokia's SR OS QoS documentation describes forwarding-class assignment and QoS classification behavior in its SR OS documentation library. The meaning and use of the DSCP field for differentiated forwarding behavior are standardized in RFC 2474.

QUESTION NO: 14

When allocating scheduling resources, which of the following statements BEST describes how the egress port scheduler allocates bandwidth?

A. The egress port scheduler takes the ratio of requested bandwidth to total bandwidth request from all children. It allocates bandwidth fairly, using these percentages and not exceeding the egress interface speed.

B. The egress port scheduler allocates bandwidth based on a WRR distribution method that takes into account the CIR and PIR requested at each level and the max speed, either default or configured, on egress.

C. The egress port scheduler allocates bandwidth on a first come, first served distribution method that takes into account the total CIR and PIR requested at each level and the max PIR requested from a single level.

D. The egress port scheduler will limit the bandwidth offered to the egress port based on the oversubscription factor, provisioned under the router interface, while using a WRR scheduler for the remaining bandwidth.

E. The egress port scheduler will use the Russian Dolls Method to allocate bandwidth to each level of requested bandwidth. HPO-Only is automatically set to 100%, allowing only high priority traffic to be buffered during congestion.

ANSWER: B

Explanation:

The egress port scheduler allocates bandwidth based on a WRR distribution method that takes into account the CIR and PIR requested at each level and the max speed, either default or configured, on egress. This describes the hierarchical scheduling behavior used by SR OS QoS. The egress port scheduler is the top-level scheduling point for traffic leaving a port, and it makes bandwidth available to its child schedulers according to configured service commitments and peak limits. CIR establishes the committed bandwidth consideration, while PIR permits eligible traffic to use additional capacity when it is available. The scheduler must also respect the effective egress port rate, whether that rate is derived from the physical port speed or is administratively configured. WRR provides a fair, weighted distribution of available capacity among contending children while preserving the hierarchy of scheduler levels and their associated CIR/PIR parameters. As a result, allocation is not merely a calculation of each child's instantaneous requested percentage; it is a QoS scheduling process that accounts for configured rates, priority treatment, and the capacity actually available at the egress interface. Nokia's SR OS QoS documentation describes the hierarchical scheduler model and the bandwidth allocation behavior of egress scheduling entities. See the Nokia SR OS Quality of Service Guide and the Nokia SR OS egress QoS documentation.

QUESTION NO: 15

When using an egress port scheduler, which scheduling loop priority are orphaned queues serviced on?

A. 1st

B. 8th

C. 9th

D. 17th

ANSWER: D

Explanation:

Orphaned queues are serviced on the 17th scheduling loop priority. In SR OS egress port scheduling, a queue is considered orphaned when it is not associated with a parent scheduler within the configured egress scheduler hierarchy. The port scheduler must still provide these queues an opportunity for service, so SR OS handles them through the dedicated orphan-queue scheduling loop. This is loop 17, following the scheduler loops used to process the configured scheduler hierarchy. This behavior ensures that queues which are not explicitly attached to a scheduler are not excluded from egress transmission; they remain eligible for service according to the egress port scheduler's defined orphan-queue handling. The distinction is important when designing hierarchical QoS policies, because attaching queues to the intended scheduler levels gives administrators explicit control over their scheduling treatment, while unattached queues are processed through the orphaned-queue mechanism. Nokia's SR OS Quality of Service documentation describes egress-port scheduler policies, scheduler hierarchy, and orphan queue operation in the SR OS QoS Guide: Egress Port Scheduler Policy. Additional SR OS QoS documentation is available through the Nokia SR OS documentation portal.

QUESTION NO: 16

What are the two least significant bits of the IP Differentiated Services field used for?

A. Explicit Congestion Notification

B. IP precedence

C. VLAN priority

D. MPLS Traffic Class

E. Packet profile marking

ANSWER: A

Explanation:

Explicit Congestion Notification is correct. The Differentiated Services field is eight bits wide. Its six most significant bits contain the DSCP value, while its two least significant bits are used for ECN where ECN is supported. ECN allows network equipment to signal congestion without necessarily discarding a packet.

The ECN bits are not part of the DSCP and do not increase the number of DSCP values available for QoS classification. The DS field structure is defined in RFC 2474, and ECN operation is specified in RFC 3168.

QUESTION NO: 17

Packets arrive at the VPLS SAP on PE1 with DSCP marking of af11 and Dot1p of 2. According to the QoS policies applied at routers PE1, P, and PE2, what is the DSCP value of packets egressing router PE2? (All the SDPs are MPLS-encapsulated, and all interfaces are using their default trust states).

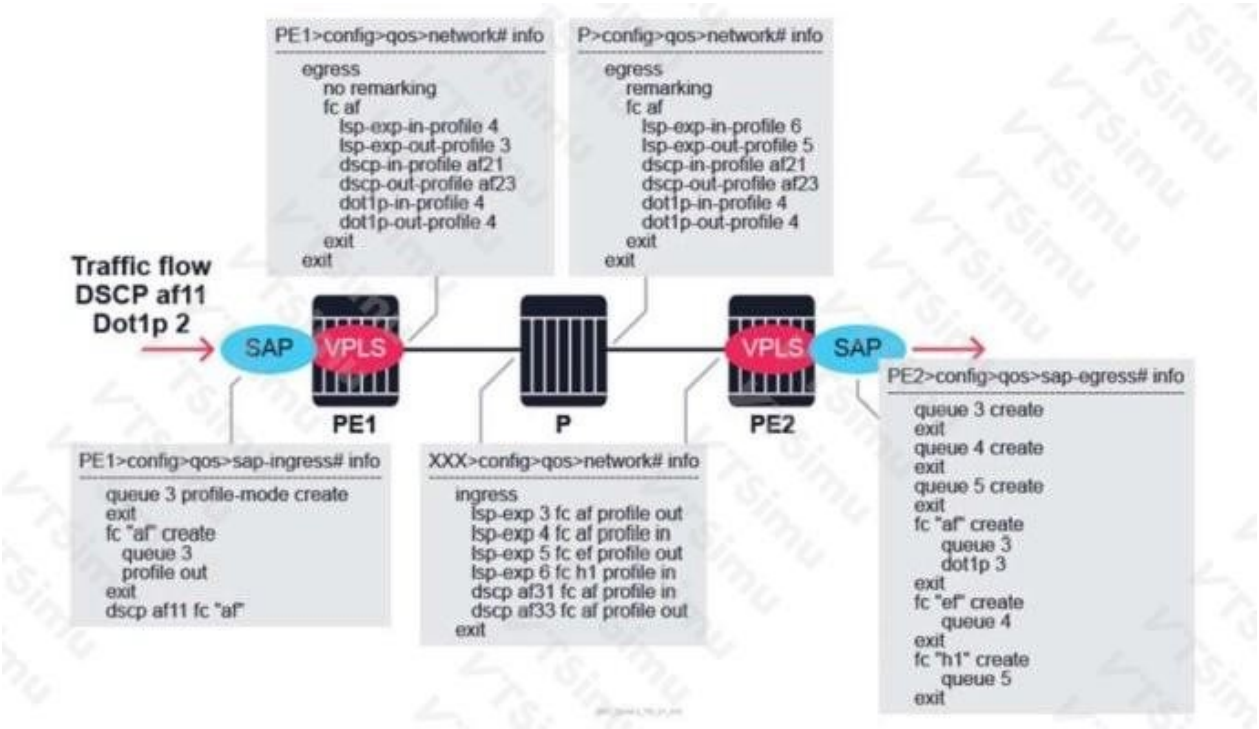

A. af11

B. af21

C. af23

D. af31

ANSWER: A

Explanation:

The DSCP value egressing PE2 is **af11**. The customer packet enters the VPLS service carrying its original IP header, including the af11 DSCP marking. QoS classification at the service access point uses the configured classification and the applicable default trust behavior to select the forwarding class and profile used within the service. For transport across the MPLS-enabled SDPs, the forwarding treatment is represented by the MPLS transport encapsulation and its traffic-class information. This allows PE1, P, and PE2 to apply the required per-hop QoS behavior without requiring the original customer IP DSCP field to be changed.

At the far-end PE2, removal of the MPLS encapsulation exposes the original customer Ethernet and IP packet. Because the applied egress treatment does not perform a DSCP remarking operation for this traffic, the packet retains the DSCP value with which it entered the VPLS service: af11. This distinction between transport QoS markings and preservation of the customer packet's IP header is fundamental to MPLS service QoS operation. Nokia's QoS documentation describes classification, forwarding-class handling, and marking behavior in the 7750 SR OS QoS overview and the interaction of QoS with services in the service QoS documentation.

QUESTION NO: 18

On the Nokia 7750 SR, what is the maximum number of ingress queues that a single network queue policy can support in a VPLS?

A. 8

B. 16

C. 32

D. 48

ANSWER: C

Explanation:

The correct maximum is **32**. On the Nokia 7750 SR, a network queue policy defines the queueing treatment used for traffic at the network ingress. In a VPLS service, the ingress network queue policy can be configured with queue entries numbered through queue 32, so one policy supports up to 32 ingress queues. These queues provide the per-class buffering and scheduling framework used when packets enter the router from the network side before they are processed and forwarded within the VPLS context. The queue count is a platform QoS capability and is not determined by the number of customer SAPs or VPLS service instances using the policy. A policy may use fewer queues when the service design has fewer traffic classes, but 32 is the supported upper limit for a single ingress network queue policy. Nokia's QoS documentation describes network queue policies and their queue configuration model in the 7750 SR OS QoS guide. See the Nokia 7750 SR OS Quality of Service Guide and the services QoS documentation for the relevant policy and queue configuration context.

QUESTION NO: 19

What is an SLA? (Choose two)

A. An SLA is used to provide automated, real-time testing and alarming for throughput, latency, and jitter across a provider's network.

B. An SLA is an agreement between a customer and a provider that dictates the treatment of customer traffic across the provider's network.

C. An SLA allows customers to control all traffic within the service provider's network by prioritizing their traffic over others as desired.

D. An SLA allows a customer to pre-mark traffic and ensure that traffic is treated as per the agreement within the provider's network.

E. An SLA is a standard set of network QoS policies that a provider shares to all its customers, allowing them to better understand the treatment of traffic within the provider's network.

ANSWER: B D

Explanation:

A service-level agreement is a contractual agreement between a service provider and its customer that defines the service being delivered and the performance and handling commitments applicable to that service. In a QoS context, it establishes the agreed treatment of the customer's traffic across the provider network, such as the traffic classes supported and the service objectives associated with those classes. This makes "An SLA is an agreement between a customer and a provider that dictates the treatment of customer traffic across the provider's network." correct.

An SLA can also specify the customer/provider traffic-marking arrangement. A customer may pre-mark frames or packets to identify their intended class of service, while the provider applies traffic treatment consistent with the agreed service definition, subject to its ingress classification and marking policy. Therefore, "An SLA allows a customer to pre-mark traffic and ensure that traffic is treated as per the agreement within the provider's network." correctly describes a practical QoS function enabled by an SLA. Nokia's SR OS QoS documentation describes the use of classification, marking, and forwarding treatment to implement differentiated services: Nokia SR OS QoS overview. See also the ITU-T definition of a service-level agreement: ITU-T Recommendation E.860.

QUESTION NO: 20

Click the exhibit button below.

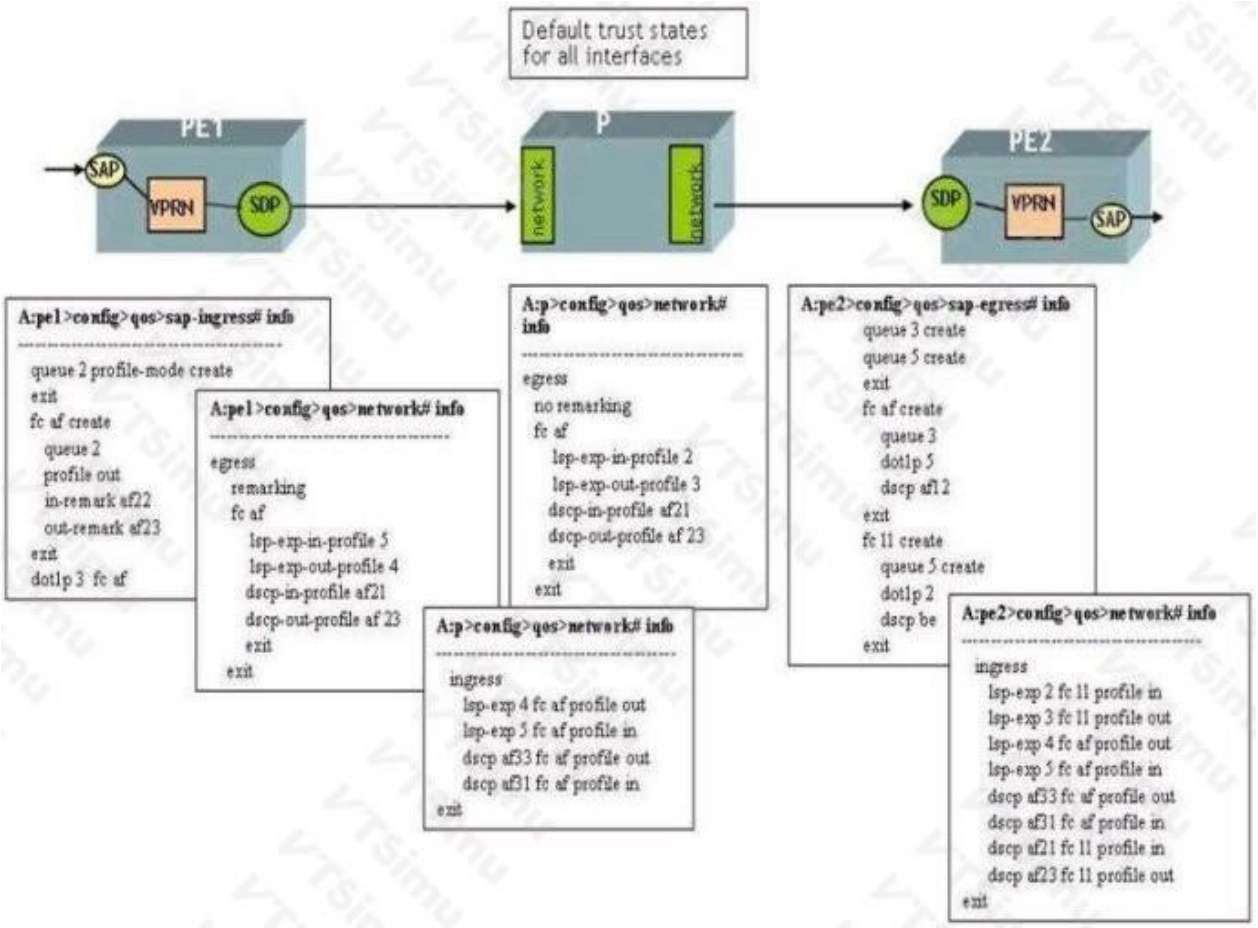

All interfaces are using their default trust states and MPLS is used as the transport tunnel. The SAP-ingress, SAP-egress, and network QoS policies have been configured as shown below.

Assume that the default network-queue policy is used on each router.

At router PE 1, customer traffic is arriving marked with DSCP BE and tagged with a dot1p value of 3.

Based on the configuration shown below for the VPRN service, what will be the DSCP and dot1p marking for the packet egressing at router PE 2? (Choose two)

A. The DSCP value will be set to af21.

B. The DSCP value will be set to be.

C. TheDSCPvaluewillbesettoaf23.

D. The DSCP value will be set to af 12.

E. Thedot1pvaluewillbesetto5.

F. Thedot1pvaluewillbesetto2.

G. The dot1p value will be set to 3.

ANSWER: D E

Explanation:

The correct egress markings are *af 12* for DSCP and *5* for dot1p. At PE 1, classification at the ingress SAP uses the applicable SAP-ingress QoS policy and the interface's default trust behavior to derive the internal forwarding class and profile from the received customer frame. The forwarding class is then carried across the MPLS transport by the network QoS treatment, including the MPLS EXP value selected for the packet. The default network-queue policy does not alter the service-policy classification and rewrite result established by the configured QoS mappings.

At PE 2, the network QoS policy maps the MPLS transport marking back to the internal service forwarding treatment. The SAP-egress policy then performs the customer-facing remarking configured for that forwarding class and profile. Consequently, the outgoing IP packet receives DSCP af12, while the outgoing Ethernet VLAN priority field receives dot1p 5. These are independent egress rewrite fields: the DSCP value is written into the IP header and the dot1p value is written into the 802.1Q tag, so both selected markings apply to the same packet. Nokia's SR OS QoS documentation describes classification, forwarding-class handling, and ingress/egress marking behavior in the SR OS QoS overview. DSCP codepoint semantics are defined by RFC 2474.

QUESTION NO: 21

Which of the following are the minimum requirements needed to classify traffic marked with DSCP AF11, AF12, and AF13 into the AF forwarding class and provide fair treatment to the drop precedence settings? (Choose three)

A. Profile mode queues

B. Priority mode queues

C. 1 queue

D. 3 queues

E. 3 forwarding classes

F. 1 forwarding class, with 2 subclasses

G. 1 forwarding class, with 3 subclasses

ANSWER: A C F

Explanation:

Profile mode queues, **1 queue**, and **1 forwarding class, with 2 subclasses** are the minimum requirements. AF11, AF12, and AF13 belong to the same Assured Forwarding class: the first AF digit identifies the forwarding class, while the second digit identifies increasing drop precedence within that class. Therefore, these packets require a common forwarding-class treatment and can share a single queue rather than needing separately scheduled queues.

On SR OS, a profile-mode queue is required because it supports differentiated congestion behavior for traffic associated with the forwarding class. The two subclasses provide the classification structure needed to distinguish the assured traffic treatment from the higher drop-precedence treatment while retaining one shared AF queue. This lets the scheduler allocate service fairly to the AF class as a whole, while congestion management can preferentially discard traffic associated with the less-preferred drop-precedence markings. Thus, the implementation preserves the central AF behavior: common forwarding resources for the class combined with differentiated discard treatment under congestion.

The IETF definition of the Assured Forwarding PHB group and its three drop-precedence levels is available in RFC 2597. SR OS QoS concepts and queue behavior are documented in the Nokia SR OS QoS Guide.

QUESTION NO: 22

Which of the following statements about a queue's rate parameters is FALSE?

A. If CIR = PIR, then all packets will be considered as in-profile.

B. CIR can be set to 0.

C. CIR can be larger than PIR.

D. CIR can be set to 'max'.

ANSWER: C

Explanation:

“CIR can be larger than PIR.” is the false statement. A queue’s committed information rate (CIR) defines the rate that is committed to the queue, while its peak information rate (PIR) establishes the upper rate limit at which the queue may be served. Because the committed rate is necessarily bounded by the peak rate, the configured CIR must not exceed PIR. This relationship ensures that the QoS scheduler can first allocate committed bandwidth and then, where capacity and scheduling policy permit, allow traffic to use bandwidth up to the configured peak. The rate values may be expressed as explicit rates or, where supported by the relevant queue context, as *max*; nevertheless, the fundamental CIR-to-PIR constraint remains applicable. Nokia SR OS QoS documentation describes queue rate configuration and the use of CIR and PIR as the committed and maximum scheduling parameters. See the Nokia SR OS documentation portal and the SR OS QoS overview.

QUESTION NO: 23

When deploying a Nokia MPLS network with QoS, how are the EXP bits set?

A. The first 3 bits of the DSCP value are copied directly into the EXP field on the first network egress.

B. The EXP bits are set on the first network egress, based on a combination of the SAP-ingress classification and the network policy applied.

C. The 3 priority bits are copied directly from the dot1p header into the EXP field of the first MPLS header.

D. The EXP bits are set based on a default map that translates DSCP codes into internal forwarding classes on SAP-ingress.

E. At the first network egress, an FC-to-EXP mapping in the network QoS policy maps the packet’s internally assigned forwarding class to EXP.

ANSWER: B E

Explanation:

The EXP bits are set on the first network egress, based on a combination of the SAP-ingress classification and the network policy applied. At service-access ingress, SR OS classifies the packet and assigns an internal forwarding class and profile using the applicable SAP ingress QoS policy. When the packet is sent into the MPLS network, the network QoS policy uses that internal forwarding class/profile context to select the MPLS Traffic Class value, historically called EXP, for the imposed label stack. This separates customer-facing classification from the provider network’s MPLS marking behavior and allows the operator to control MPLS treatment consistently through network QoS policy mappings.

At the first network egress, an FC-to-EXP mapping in the network QoS policy maps the packet’s internally assigned forwarding class to EXP. This is the specific mapping operation that produces the three-bit MPLS marking. The mapping can be tailored by policy, allowing forwarding classes to receive the MPLS traffic-class values required by the provider’s transport QoS design. MPLS Traffic Class/EXP marking is defined as a three-bit field used for differentiated forwarding treatment; see RFC 3270. Nokia SR OS QoS documentation and configuration guidance are available through the Nokia SR OS documentation portal.

QUESTION NO: 24

Click the exhibit button below.

```
qos
 network 30 create
 ingress
      dot1p 2 fc be profile out
      dot1p 3 fc h2 profile in
      dot1p 4 fc h1 profile out
      dot1p 5 fc ef profile in

      lsp-exp 0 fc be profile out
      lsp-exp 1 fc l2 profile in
      lsp-exp 2 fc af profile out
      lsp-exp 3 fc af profile in
      lsp-exp 4 fc h1 profile in
      lsp-exp 5 fc ef profile in
      lsp-exp 6 fc h1 profile in
      lsp-exp 7 fc nc profile in

      dscp be fc be profile out
      dscp ef fc ef profile in
      dscp cs1 fc l2 profile in
      dscp nc1 fc h1 profile in
      dscp nc2 fc nc profile in
      default-action fc l1 profile out
      exit
```

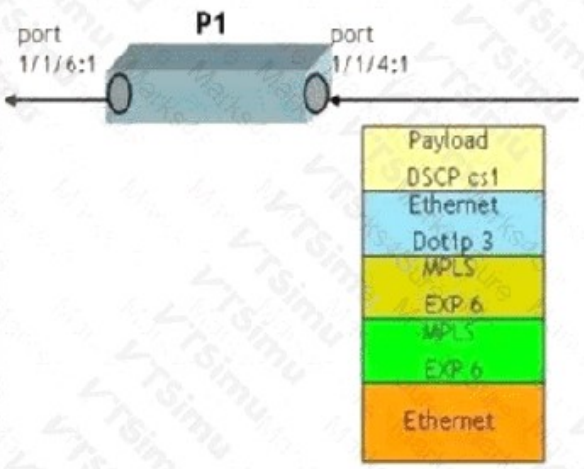


```
A:P1>config>router# info
#------------------------------------------
echo "IP Configuration"
#------------------------------------------
      interface "toPE2"
            address 10.1.4.1/27
            port 1/1/4
            qos 30
      exit
```

Based on the configuration of the network policy (below), what will be the forwarding class associated with a MPLS encapsulated customer packet that arrives on a dotlQ encapsulated network port 1/1/4:1 on P1 with the following characteristics:

EXP value = 6

DSCP value = cs1

Dot1pvalue = 3

A. EF

B. H1

C. L1

D. H2

E. L2

ANSWER: D

Explanation:

H2 is the forwarding class selected by the configured network policy for an MPLS-encapsulated packet with an EXP value of 6. On a network-facing SAP or port, the router performs classification according to the received encapsulation. Because this packet is MPLS encapsulated, the MPLS traffic-class field, historically called EXP, is the relevant marking used by the MPLS classification rule in the network policy. The policy mapping associates the received value of 6 with forwarding class H2, so the packet enters the H2 forwarding-class queues and receives the corresponding in-profile or out-of-profile treatment defined by that policy.

The customer packet's inner IP DSCP marking of cs1 and the Dot1p value of 3 do not replace the MPLS classification result in this case. Those markings may be used for traffic arriving without MPLS encapsulation, or can be copied or remarked by explicitly configured QoS behaviors, but the received MPLS label-stack traffic-class value governs this network-policy lookup.

This behavior supports consistent QoS handling across an MPLS transport domain, where the outer MPLS marking conveys the service treatment selected at the network boundary. Nokia SR OS QoS documentation is available from the Nokia Service Router documentation portal; MPLS traffic-class-field terminology is also defined in RFC 5462.

QUESTION NO: 25

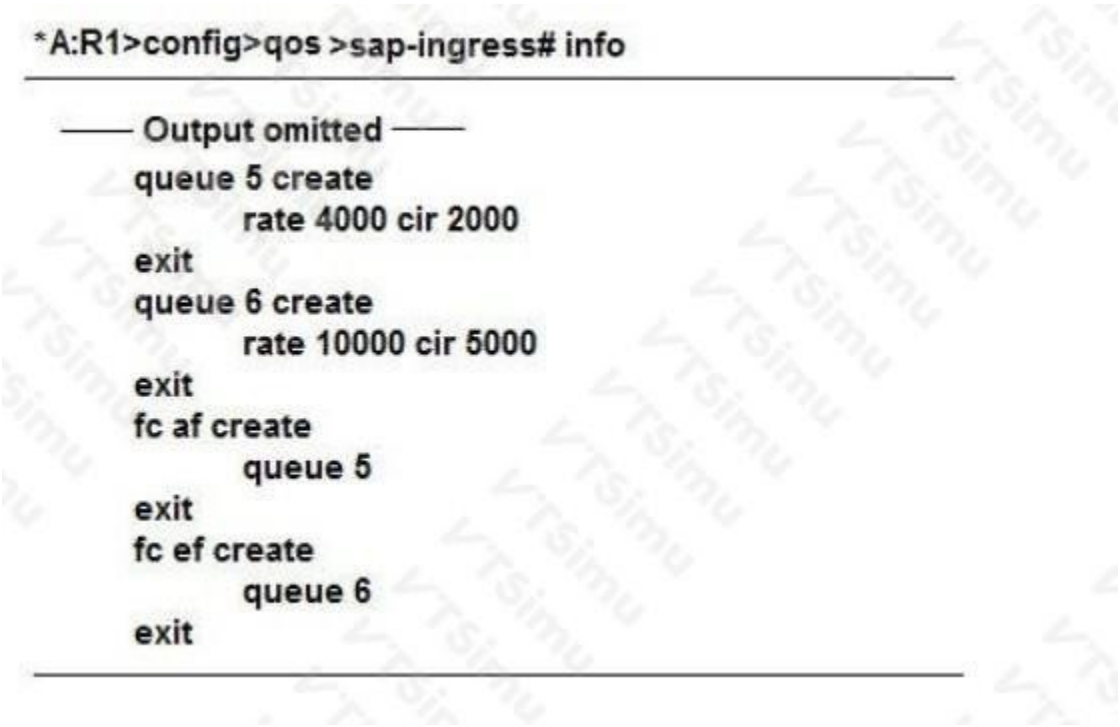

Based on the image shown above, and assuming that the default scheduler of the Nokia 7750 SR is on an IOM2 card, in which order will packets be scheduled?

A. Queue 6 within-CIR packets, then queue 5 within-CIR packets, then queue 5 and queue 6 above-CIR packets in round-robin.

B. Queue 6 within-CIR packets, then queue 5 within-CIR packets, then queue 5 and queue 6 above-CIR packets in biased round-robin.

C. Queue 6 within-CIR packets, then queue 5 within-CIR packets, then queue 6 above-CIR packets, and finally queue 5 above-CIR packets.

D. Queue 6 and queue 5 within-CIR packets in biased round-robin, then queue 5 and queue 6 above-CIR packets in round-robin.

ANSWER: B

Explanation:

Queue 6 within-CIR packets, then queue 5 within-CIR packets, then queue 5 and queue 6 above-CIR packets in biased round-robin is correct. On the IOM2 default scheduling implementation, traffic that is within a queue's committed information rate is handled before excess, above-CIR traffic. The priority relationship shown in the scheduler configuration is therefore applied first to the within-CIR portions: queue 6 is served before queue 5. This protects the committed service associated with each queue and honors the configured priority treatment for in-profile traffic.

After the committed portions have been serviced, packets exceeding CIR remain eligible for transmission up to their peak limits. The IOM2 default scheduler does not drain all above-CIR traffic from one of these queues before considering the other. Instead, it uses biased round-robin behavior for the above-CIR contention stage. That mechanism shares available excess bandwidth between the eligible queues according to the scheduler's bias or weighting treatment while maintaining service opportunities for both queues. This two-stage result—priority handling of within-CIR traffic followed by biased round-robin handling of above-CIR traffic—is the required packet-scheduling order. Nokia's SR OS QoS documentation is available

through the Nokia SR OS documentation portal; platform context is provided by the Nokia 7750 Service Router product page.

QUESTION NO: 26

Which of the following are possible criteria for classifying packets at the SAP-ingress? (Choose three)

A. The EXP bits in the MPLS header.

B. The packet's source and destination IP addresses.

C. The service ID of the service to which the SAP belongs.

D. The dot1p bits in the frame header.

E. The packet's Ethertype.

F. The packet's DSCP bits in the IP header.

ANSWER: A D F

Explanation:

At SAP ingress, SR OS assigns a forwarding class by examining QoS markings carried in the packet or frame headers. The EXP bits in the MPLS header are a valid classification key for MPLS-encapsulated traffic. MPLS Traffic Class/EXP information is specifically intended to carry per-hop QoS behavior, as described in RFC 5462. The dot1p bits in the frame header are likewise a valid SAP-ingress criterion for Ethernet traffic, allowing an ingress QoS policy to map IEEE 802.1p priority values to forwarding classes.

The packet's DSCP bits in the IP header are the other standard SAP-ingress classification criterion. DSCP is the differentiated-services field used to identify the traffic treatment that should be applied to an IP packet; its standardized semantics are defined by RFC 2474. In practical SR OS QoS configuration, these MPLS EXP, Ethernet dot1p, and IP DSCP markings are mapped into forwarding classes and profiles by the SAP-ingress QoS policy. The question’s original set omits DSCP, so it must be added to provide the required three valid packet-classification criteria.