PassQueen
a.k.a dumpsqueen.com

Nokia SRA Composite Exam

Nokia 4A0-C02

Version Demo

Total Demo Questions: 87

Total Premium Questions: 878

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, IP Routing	252
Topic 2, MPLS	40
Topic 3, Services	586
Total	878

QUESTION NO: 1

Click the exhibit button below. Given the slope-policy (below), which of the following statements are TRUE? (Choose two)

```
A:sr1a>config>qos>slope-policy# info detail
----------------------------------------------
            description "Slope policy"
            high-slope
                start-avg 50
                max-avg 70
                max-prob 60
                no shutdown
            exit
            low-slope
                start-avg 10
                max-avg 60
                max-prob 80
                no shutdown
            exit
            time-average-factor 1
----------------------------------------------
```

A. All out-of-profile traffic will be dropped before that of any in-profile traffic.

B. The discard probability slope of an in-profile traffic is steeper than that of an out-of-profile traffic.

C. When the shared buffer utilization reaches 55%, both in-profile and out-of-profile packets may be dropped.

D. The maximum probability with which an out-of-profile packet can be dropped is higher than that of an in-profile packet.

E. When the shared buffer utilization is at 61%, only in-profile packets are currently in the shared buffer pool.

ANSWER: C D

Explanation:

When the shared buffer utilization reaches 55%, both in-profile and out-of-profile packets may be dropped because the slope-policy graph shows that the discard curves for both traffic profiles are active at that utilization level. A slope policy applies probabilistic packet discard according to shared-buffer occupancy: once a profile's configured start-average threshold has been reached, packets of that profile can be selected for discard at a probability that increases along its configured slope. Therefore, overlap between the active portions of the two curves means that packets from either profile may be discarded.

The maximum probability with which an out-of-profile packet can be dropped is higher than that of an in-profile packet because the graph assigns the out-of-profile curve a higher maximum discard probability. This preserves preferential treatment for in-profile traffic while congestion develops: out-of-profile traffic is exposed to the more severe discard behavior, whereas in-profile traffic retains the lower configured maximum probability. Nokia SR OS slope policies use shared-buffer utilization and profile-specific slope parameters to implement this congestion-management behavior. See the Nokia Quality of Service Guide: Slope Policies and the Nokia SR OS documentation portal.

QUESTION NO: 2

Which of the following statements are TRUE regarding the PE device in a QoS enabled network? (Choose two)

A. A device that provides no differentiation between customer traffic flows.

B. A device that classifies and marks customer traffic.

C. A device that uses MPLS EXP bits to differentiate between traffic flows.

D. A device that creates macroflows towards the core from multiple customer sites.

E. A device that passes macroflows while optionally changing QoS markings.

ANSWER: B D

Explanation:

In a QoS-enabled MPLS service network, the provider edge is the boundary at which customer traffic enters and leaves the provider domain. “A device that classifies and marks customer traffic.” is correct because the provider edge can classify packets using customer-facing criteria, such as ingress interface, service context, or packet header fields. It then associates the traffic with the appropriate forwarding class and applies QoS markings used by the provider network. This ingress classification and marking establishes the service treatment that traffic is intended to receive as it traverses the network.

“A device that creates macroflows towards the core from multiple customer sites.” is also correct. The provider edge aggregates customer traffic into provider-facing forwarding aggregates, commonly called macroflows, for transport across the MPLS core. This aggregation allows the core to scale by carrying a manageable set of QoS treatment classes rather than maintaining customer-specific microflow state throughout the network. Nokia’s QoS architecture describes classification, forwarding classes, marking, policing, and scheduling as core QoS functions at service and network boundaries. MPLS DiffServ behavior also relies on aggregating packets into behavior aggregates for consistent per-hop treatment. See the Nokia SR OS documentation portal and RFC 3270, MPLS Support of Differentiated Services.

QUESTION NO: 3

What are the limitations of a traditional layer 2 switch? (Choose 3}

A. Isolates collision domains.

B. Requires layer 2 loop prevention.

C. Slow response to topology changes.

D. Limited to 4094 vlans per switch.

E. Increased flooding.

ANSWER: B C D

Explanation:

Traditional Ethernet Layer 2 networks must use a loop-prevention mechanism because redundant physical links can create bridging loops. Such loops can cause broadcast and unknown-unicast traffic to circulate indefinitely, so protocols such as Spanning Tree Protocol block selected redundant paths to maintain a loop-free active topology. This introduces a design and operational limitation: available links may be placed in a blocking state rather than used for forwarding.

Slow response to topology changes is also a traditional Layer 2 limitation. When a link or bridge state changes, the network must reconverge and MAC forwarding information may need to be relearned. Depending on the protocol mode and timers, this can temporarily affect traffic forwarding. Modern rapid variants improve convergence, but the underlying topology-control and relearning behavior remains an important consideration in conventional bridged networks.

Limited to 4094 vlans per switch is correct because IEEE 802.1Q identifies VLANs with a 12-bit VLAN ID field. Values 1 through 4094 are usable VLAN identifiers; the remaining values are reserved. This fixed identifier space constrains large Layer 2 segmentation designs and is one reason service-provider networks use additional encapsulation and virtualization technologies when greater scale is needed. See the IEEE 802.1Q standard overview and Nokia’s Layer 2 services overview.

QUESTION NO: 4

An I-VPLS can be associated with many B-VPLS services.

A. TRUE

B. FALSE

ANSWER: B

Explanation:

FALSE is correct. In Nokia Provider Backbone Bridging VPLS, the I-VPLS is the customer-facing service instance and the B-VPLS is the backbone transport service used to carry encapsulated customer traffic across the provider backbone. The service association is designed so that an I-VPLS is bound to one B-VPLS, which provides a single backbone context for that I-VPLS. This one-to-one mapping from the I-VPLS perspective preserves the service model and determines the backbone identifier and forwarding context used when customer frames are encapsulated for transport over the PBB network.

A B-VPLS can provide shared backbone connectivity for multiple I-VPLS instances, allowing service multiplexing and backbone MAC-address scaling. However, that does not permit one I-VPLS to be simultaneously associated with many B-VPLS services. The association is configured at the I-VPLS service level and identifies the particular B-VPLS that carries its traffic. Nokia's Layer 2 Services documentation describes PBB-VPLS architecture, including the relationship between customer-facing I-VPLS services and backbone-facing B-VPLS services. See the Nokia SR OS PBB-VPLS documentation and the Nokia Layer 2 Services guide.

QUESTION NO: 5

Click on the exhibit below.

```
A:PE5# show service service-using
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm    Opr      CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
10           VPLS      Up     Up       1                 09/12/2006 18:03:34
20           VPLS      Up     Up       1                 09/11/2006 11:14:45
100          VPLS      Up     Up       1                 09/12/2006 17:51:07
-------------------------------------------------------------------------------
Matching Services : 3
-------------------------------------------------------------------------------
```

Management VPLS 100 has been created for redundant spoke-SDP connections. User VPLSs 10 and 20 are using the same spoke-SDPs that are being managed by VPLS 100. Unexpected behavior has been reported. Based on the following output, what is the most likely reason for the unexpected behavior?

A. VPLS 100 was not actually configured as a management VPLS.

B. VPLS 10 and VPLS 20 are not actually configured as a user VPLS.

C. VPLS 10 and VPLS 20 are not using the same spoke-sdps that the management VPLS is using.

D. Based on the output there is now way of telling if there is a problem with the configuration.

ANSWER: A

Explanation:

VPLS 100 was not actually configured as a management VPLS. In SR OS spoke-SDP redundancy, merely creating a VPLS that contains the relevant redundant spoke-SDP bindings does not cause it to manage the bindings for other VPLS services. The service must be explicitly designated with the management-VPLS configuration role. That role makes the VPLS responsible for determining the active and standby state of the redundant spoke-SDPs and for propagating the resulting forwarding state to the associated user VPLSs.

The output indicates that VPLS 100 is operating as an ordinary VPLS rather than with the required management role. Consequently, VPLSs 10 and 20 cannot rely on VPLS 100 to control which common spoke-SDP is active. This can produce inconsistent forwarding behavior, such as traffic using an unintended path or failing to switch as expected during a spoke-SDP failure. The configuration must explicitly identify VPLS 100 as the management VPLS before its status can govern the shared spoke-SDP redundancy behavior used by the user VPLSs. Nokia's SR OS service documentation describes VPLS and spoke-SDP behavior in the VPLS overview and the associated services overview.

QUESTION NO: 6

Which of the following are categories of forwarding classes on the Alcatel-Lucent 7750 SR? (Choose two)

A. Real-time

B. Assured

C. Non-conforming

D. High priority

E. In-profile

ANSWER: A D

Explanation:

Real-time and High priority are forwarding-class categories used by SR OS QoS to distinguish traffic requiring preferential treatment. A forwarding class provides an internal traffic classification that QoS policies use when packets are mapped to queues and scheduled for transmission. The real-time category identifies delay- and jitter-sensitive traffic, such as voice or interactive media, for handling consistent with its service requirements. The high-priority category identifies traffic that should receive preferential queue treatment relative to lower-priority traffic when resources are contended. These classifications allow an operator to apply coherent QoS behavior across packet classification, queue assignment, scheduling, and congestion-management functions rather than treating every packet flow independently. The 7750 SR supports a set of predefined forwarding classes, and QoS policy configuration maps packet marking and other classification criteria into those classes. Nokia's SR OS QoS documentation describes forwarding classes and their use in queueing and scheduling; see the Nokia SR OS QoS Guide: Forwarding Classes and the Nokia SR OS QoS overview.

QUESTION NO: 7

Which of the following are trusted boundaries, by default? (Choose two)

A. An IES SAP

B. A VPLS SAP

C. A network port

D. AVPRN SAP

E. An access port

ANSWER: C D

Explanation:

By default, a network port and a VPRN SAP are trusted boundaries in SR OS CPU-protection processing. A trusted boundary identifies an ingress location where control-plane traffic is treated as arriving from a trusted part of the provider network or from a routed VPN context. This default classification is significant because CPU protection uses the ingress boundary to distinguish traffic sources before applying policy and queue treatment for packets directed to the router's control plane. Network ports normally connect to the routing infrastructure, making them trusted by default. Similarly, a VPRN SAP is associated with a Layer 3 VPN service and is treated as a trusted service boundary under the default CPU-protection model. Operators can alter defaults through configuration when their security design requires a different trust model, but the question asks specifically for the out-of-the-box classifications. Nokia's SR OS documentation describes CPU protection and its use of trusted and untrusted boundaries for control-plane traffic handling; see the Nokia SR OS documentation portal and the Nokia 7750 SR OS Infocenter.

QUESTION NO: 8

The policies shown below from the configuration of an Alcatel-Lucent 7750 SR have been created to implement an Extranet VPRN between the main sites of two separate VPRNs. How should these policies be applied?

A. ample Policy should be applied as an export policy on the PE connected to the main site of Customer 1

B. ample Policy 3 should be applied as an export policy on the PE connected to the main site of Customer 2

C. ample Policy 3 should be applied as an export policy on both PEs connected to the main site of Customer 1 and the main site of Customer 2

D. ample Policy should be applied as an export policy on both PEs connected to the main site of Customer 1 and the main site of Customer 2

E. ample Policy should be applied as an export policy on all PEs connected to sites of Customer 1

ANSWER: A

Explanation:

“ample Policy should be applied as an export policy on the PE connected to the main site of Customer 1” is correct. An extranet VPRN selectively shares VPN-IP routes between otherwise separate customer VPNs. On SR OS, this is accomplished through VPRN route-target and VPN route-policy processing: routes eligible for sharing are accepted and advertised from the originating VPRN with the attributes that identify them as extranet routes. Applying the policy as an export policy on the PE serving Customer 1’s main site controls the routes Customer 1 contributes to the inter-VPRN exchange at the point where they enter the provider VPN control plane. This maintains the intended hub or main-site sharing model while allowing the receiving VPRN to import only the explicitly authorized routes. Export policy evaluation is therefore the appropriate place to select and tag Customer 1’s advertised VPN routes for the extranet relationship. Nokia’s SR OS documentation describes VPRN route-policy use and the import/export processing applied to VPN routes in its VPRN configuration material: Nokia SR OS VPRN documentation and Nokia SR OS route-policy documentation.

QUESTION NO: 9

Click on the exhibit below.

How many LSPs are required in total between Metro A and Metro B if 100 hierarchical VPLS services are deployed across the two metro networks?

A. 100

B. 200

C. 1

D. 2

ANSWER: D

Explanation:

2 is correct because the hierarchical VPLS design can aggregate the traffic for all 100 customer VPLS instances onto shared MPLS transport between Metro A and Metro B. The service count does not require a separate inter-metro LSP for every VPLS instance. Instead, the VPLS service configuration uses the available service-distribution path and transport tunnel infrastructure to carry the individual service traffic across the provider network.

An MPLS LSP is unidirectional. Therefore, bidirectional connectivity between the two metro networks requires one LSP from Metro A toward Metro B and another LSP from Metro B toward Metro A. Together, these two directional LSPs provide the shared inter-metro transport required by the hierarchical VPLS deployment, irrespective of the 100 logical VPLS services using that transport. This separation of service instances from underlying packet-switched transport is a core scalability benefit of hierarchical VPLS architectures. RFC 4665 describes the hierarchical VPLS model and its use of provider-network transport for interconnected service domains: RFC 4665. MPLS LSP directionality is also defined in the MPLS architecture: RFC 3031.

QUESTION NO: 10

Which of the following regarding the TTL of a BGP message is TRUE?

A. TTL is a well-known mandatory BGP attribute.

B. The default TTL value of an eBGP message is 1, and 64 for an iBGP message.

C. The default TTL value of an iBGP message can be modified using the multihop command.

D. The default TTL value of an eBGP message can be modified using the "ttl-security" command.

ANSWER: D

Explanation:

The default TTL value of an eBGP message can be modified using the "ttl-security" command is correct. TTL is carried in the IP header of the TCP packets that transport BGP, and it provides a useful way to constrain where a BGP session can originate. For a directly connected eBGP peer, the usual transmitted TTL is one. When TTL security is enabled, the router instead sends BGP packets with a TTL of 255 and validates the received TTL against the configured number of permitted hops. This implements the Generalized TTL Security Mechanism: a legitimate peer close to the router will deliver packets whose TTL remains near 255, whereas packets injected from a more distant source will not meet the required threshold. On Nokia SR OS, the `ttl-security` configuration is intended for eBGP peers and changes the TTL behavior needed for this protection; it is therefore a valid mechanism for changing the normal directly connected eBGP TTL behavior. The underlying mechanism and its use of an initial TTL of 255 are specified by the IETF in RFC 5082. Nokia SR OS BGP configuration and operational documentation is available through the Nokia Service Router documentation portal.

QUESTION NO: 11

Two management VPLSs have been created. mVPLS 1 and mVPLS 2. Two user VPLS services exist. uVPLS 10 and uVPLS 20. What command can be used to determine which management VPLS is managing uVPLS 10?

A. show stp 10

B. show service vpls 1 stpC, show service id 10 stp

C. show service id 1 stp or show service id 2 stp

ANSWER: C

Explanation:

`show service id 1 stp` or `show service id 2 stp` is correct because management VPLS operation is determined from the STP information maintained by each management VPLS. A management VPLS provides the shared STP control plane for the user VPLS instances assigned to it. Displaying STP status for management VPLS 1 and/or management VPLS 2 lets the operator inspect the managed VPLS associations and identify the management VPLS responsible for uVPLS 10. In practice, the output from the applicable management VPLS STP display identifies the user VPLS services under that management domain, along with the relevant STP state and topology information. This is particularly useful when multiple management VPLSs exist, because the user VPLS service ID alone does not establish which management instance owns its STP operation. Nokia SR OS documentation describes VPLS and its management/control-plane behavior in the Layer 2 Services documentation, available from the Nokia SR OS documentation portal. The platform command syntax and operational show commands are documented in the associated SR OS command reference.

QUESTION NO: 12

Click the exhibit.

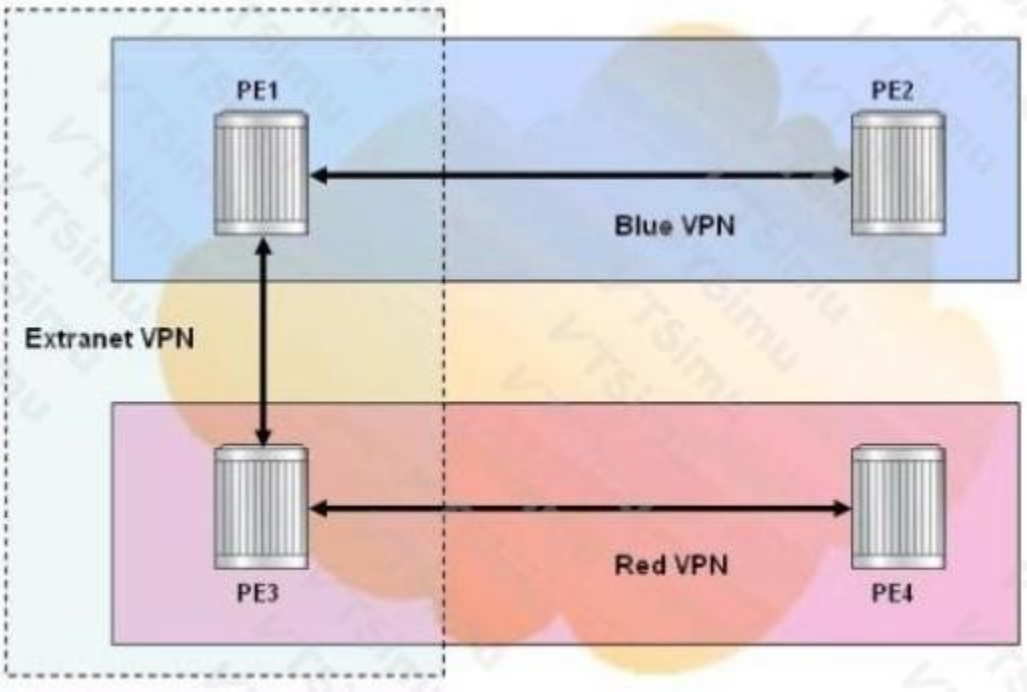

How many route targets and route distinguishers are required to implement the extranet VPRN?

A. Three route targets and three route distinguishers

B. Five route targets and three route distinguishers

C. Three route targets and two route distinguishers

D. Five route targets and two route distinguishers

ANSWER: C

Explanation:

Three route targets and two route distinguishers is correct. A route distinguisher is required for each separate VPRN routing domain represented in the design, so the two distinct customer VPN route tables require two route distinguishers. The route distinguisher makes otherwise potentially overlapping IPv4 prefixes unique when they are advertised as VPN-IPv4 routes through MP-BGP. It identifies the VPN route context; it does not itself control which VPRNs receive the route.

The extranet connectivity is controlled by route targets. Three route-target communities provide the necessary import/export policy: route membership can remain associated with each of the two customer VPRNs while a dedicated route target identifies the routes intended for controlled exchange through the extranet. Import policies on the participating VPRNs then select the shared extranet routes without removing the separation of the underlying customer routing domains. This is the normal RFC 4364 model used by SR OS VPRN services: RDs provide uniqueness, while RTs express route distribution and membership policy. See RFC 4364, BGP/MPLS IP Virtual Private Networks and RFC 4360, BGP Extended Communities.

QUESTION NO: 13

Click on the exhibit below.

The SAP on PE-A has been configured with dot1q encapsulation. The SAP on PE-B has been configured with qinq encapsulation. CE-A sends frames with a single tag of 100. The VLAN tag from CE-A must be passed transparently. CE-B is expecting a frame with a top tag of 200 and bottom tag of 100. Which of the following is the correct SAP id for PE-B?

A. sap 1/1/1:200.100

B. sap 1/1/1:200*

C. sap 1/1/1:*

D. sap 1/1/1:0*

ANSWER: B

Explanation:

`sap 1/1/1:200*` is the correct SAP identifier because a QinQ SAP identifies the outer VLAN tag and can use a wildcard for the customer, or inner, VLAN tag. Here, the outer tag required at CE-B is 200, so 200 is specified as the first tag value. The wildcard permits the incoming customer VLAN ID to remain unchanged through the service. Consequently, the single customer tag of 100 received from CE-A is retained as the inner tag, while PE-B presents the required QinQ frame with an outer tag of 200 and an inner tag of 100 toward CE-B. This is the needed tag manipulation behavior when the service must transparently carry the original customer VLAN while adding a service-facing outer VLAN tag at the QinQ attachment point. Nokia SR OS documentation describes SAP encapsulation and VLAN matching as part of Layer 2 service configuration, including the use of VLAN identifiers and wildcard matching for service access points. See the Nokia SR OS Layer 2 Services documentation and the Nokia SR OS Layer 2 Services configuration guide.

QUESTION NO: 14

Click the exhibit button below.

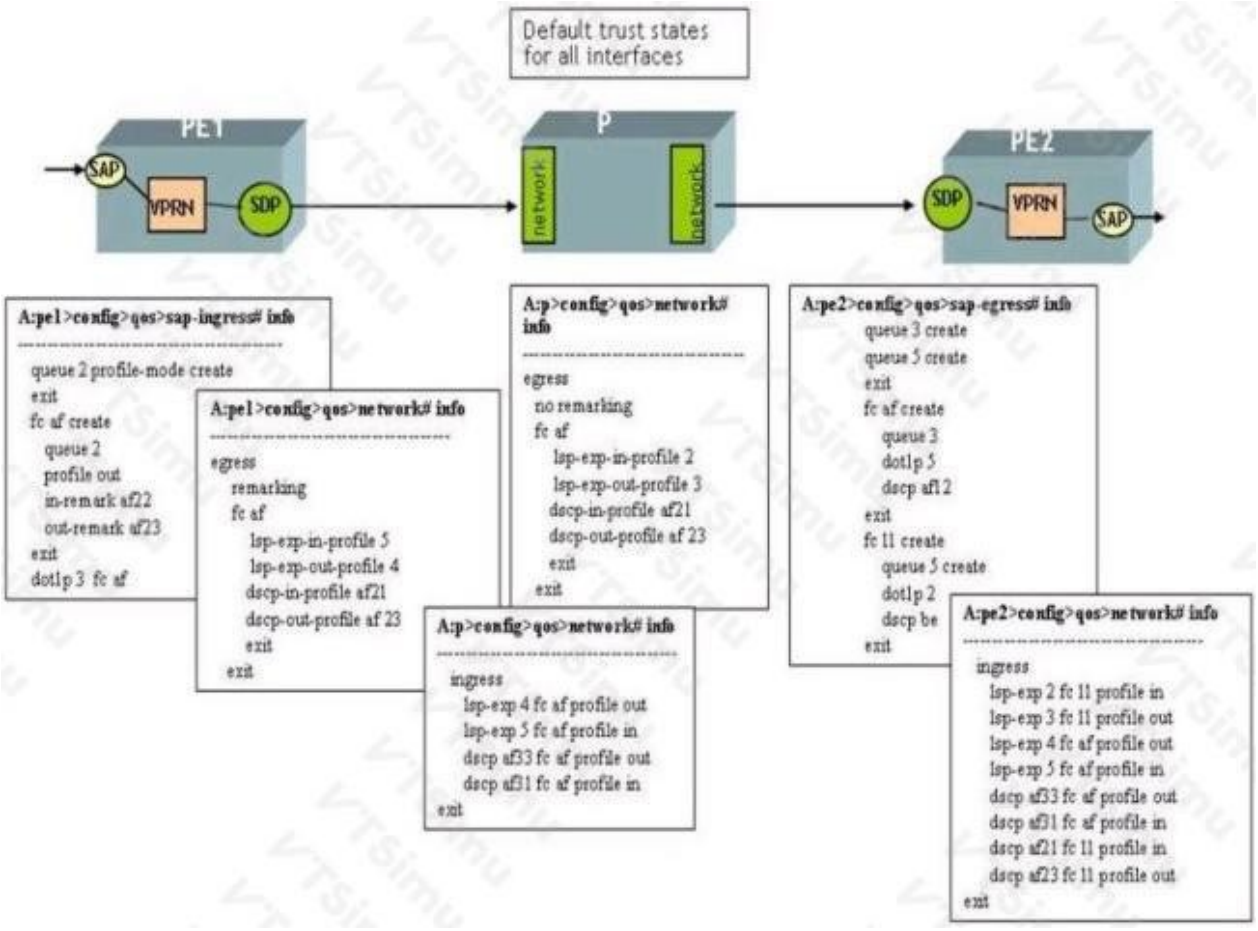

All interfaces are using their default trust states and MPLS is used as the transport tunnel. The SAP-ingress, SAP-egress, and network QoS policies have been configured as shown below. Assume that the default network-queue policy is used on each router.

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. The DSCP value will be set to af23.

D. The DSCP value will be set to af12.

E. The dot1p value will be set to 5.

F. The dot1p value will be set to 2.

G. The dot1p value will be set to 3.

ANSWER: D E

Explanation:

The DSCP value will be set to af12, and the dot1p value will be set to 5. The ingress SAP QoS policy first classifies the customer frame according to the applicable ingress marking map and assigns its forwarding-class and profile treatment. Across the MPLS core, the packet's service QoS treatment is carried through the MPLS transport using the network QoS handling and the MPLS traffic-class/EXP encoding associated with that treatment. The default trust behavior does not cause the original customer DSCP and VLAN priority values simply to be preserved end to end; the service QoS policies and MPLS transport mapping determine the forwarding treatment.

At PE 2, the packet is mapped back into the appropriate service forwarding class. The configured SAP egress QoS policy then performs egress remarking from that forwarding-class and profile result. Its DSCP remarking map selects AF12, while its IEEE 802.1p remarking map selects priority 5 for the outgoing customer Ethernet frame. Thus, the markings observed on the packet leaving the egress SAP are AF12 in the IP header and dot1p 5 in the VLAN tag. Nokia's QoS documentation describes SAP ingress classification, service forwarding classes, and SAP egress packet marking in the SR OS QoS overview and the SAP ingress QoS documentation.

QUESTION NO: 15

A customer has requested a VPLS service. The customer is using dot1q encapsulation and they have requested that all vlan tags be transparently passed. Which two SAP IDs will accept VLAN tags and pass them transparently? (Choose 2)

A. sap 1/1/1

B. sap 1/1/1:0*

C. sap 1/1/1:0

D. sap 1/1/1:*

ANSWER: A D

Explanation:

`sap 1/1/1` is a null-encapsulation SAP identifier. It does not classify traffic using a VLAN encapsulation value, so frames arriving at the SAP, including frames carrying customer VLAN tags, are delivered to the VPLS unchanged. This is an appropriate approach when the service must be transparent to customer VLAN tagging rather than terminating an individual VLAN at the SAP.

`sap 1/1/1:*` is the dot1q wildcard SAP form. The wildcard represents the complete set of VLAN encapsulation values on the port, allowing the SAP to accept VLAN-tagged frames regardless of their customer VLAN ID. The customer's original dot1q header remains part of the customer frame as it is switched through the VPLS, which provides VLAN-transparent transport across the Layer 2 service.

These two SAP forms therefore support a VPLS design in which tagged customer traffic is accepted without restricting the service to one particular VLAN value. Nokia's SR OS service documentation describes SAP identifiers, encapsulation matching, and wildcard use in the VPLS service model. See the Nokia SR OS VPLS documentation and the Nokia SR OS SAP documentation.

QUESTION NO: 16

On the Alcatel-Lucent 7750 SR, which of the following QoS-related policies can be configured on a SAP? (Choose three)

A. The slope-policy

B. The SAP-ingress policy

C. The SAP-egress policy

D. The network-queue policy

E. The scheduler-policy

ANSWER: B C E

Explanation:

The SAP-ingress policy, the SAP-egress policy, and the scheduler-policy are QoS policy types that can be associated with a Service Access Point on a 7750 SR. A SAP-ingress policy defines how traffic received from a customer-facing SAP is classified, policed, marked, and queued before it enters the service. This permits service-specific enforcement of ingress bandwidth and forwarding treatment. A SAP-egress policy provides the corresponding egress QoS treatment for traffic leaving through the SAP, including classification, queue assignment, queue parameters, and remarking behavior where configured.

A scheduler-policy can also be applied in the SAP egress context to control the scheduling hierarchy and allocation of available egress bandwidth among queues or child schedulers. This makes it possible to apply differentiated scheduling behavior per SAP rather than relying solely on interface-level behavior. Together, these policy attachments provide the principal ingress, egress, and hierarchical scheduling controls for customer service traffic at a SAP. Nokia's QoS documentation describes SAP ingress and egress QoS policy application and the use of scheduler policies in service egress scheduling: Nokia 7750 SR OS Quality of Service Guide and Nokia Scheduler Policies documentation.

QUESTION NO: 17

Click the exhibit.

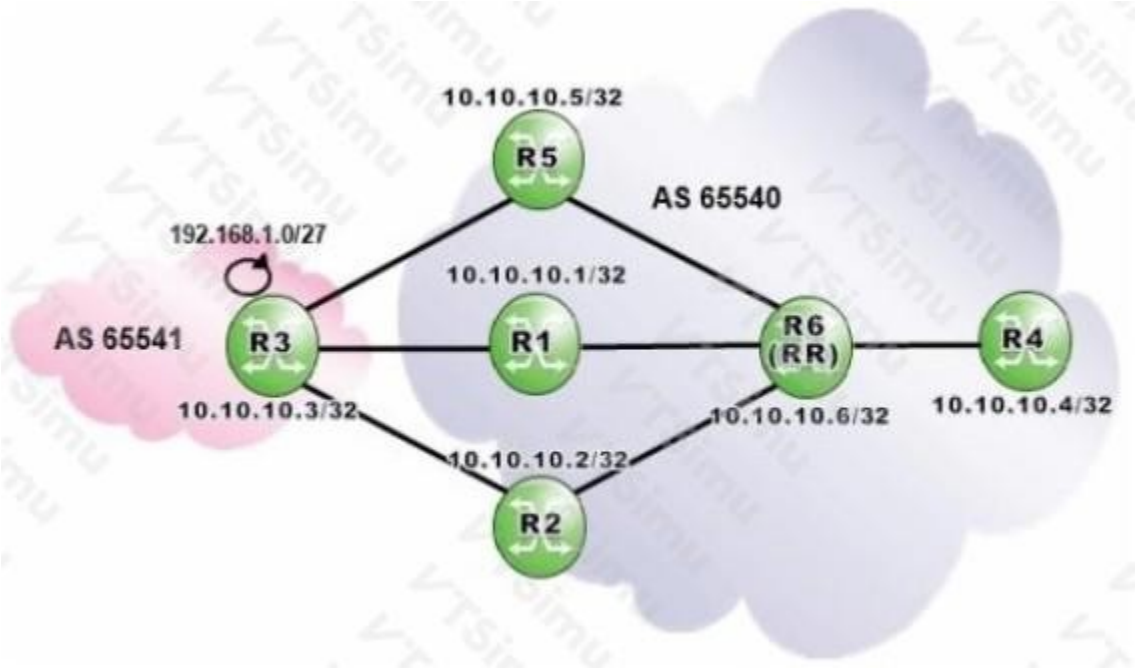

Router R6 is a route reflector with clients R1, R2, R4 and R5. R4 receives three routes from router R6 and is configured with ECMP 3. Given the following BGP configuration on router R4, how many primary and backup paths will be present in router R4's BGP routing table?

```
R4>config>router>bgp# info
----------------------------------------
        multipath 3
        add-paths
            ipv4 send 3 receive
        exit
        backup-path ipv4
        group "IBGP"
            peer-as 65540
            neighbor 10.10.10.6
            exit
        exit
```

A. Two primary paths and one backup path.

B. One primary path and two backup paths.

C. Three primary path and one backup path,

D. Three primary paths.

ANSWER: D

Explanation:

Three primary paths. is correct. Router R4 receives three eligible BGP paths for the same destination from its route reflector. The configured ECMP value of 3 permits up to three equal BGP paths to be used concurrently. When the BGP multipath criteria in the displayed configuration are met, all three received paths are retained as equal active, or primary, paths rather than choosing only one best path. The ECMP limit is significant because it caps the number of simultaneously active equal-cost next hops; here, the number of eligible equal paths exactly matches that limit.

A backup path is a path retained as an inactive alternate after the active path set has been selected. In this case, no additional eligible route remains after the three received routes have been installed as the three primary paths. Therefore, R4's BGP routing table contains three primary paths and no backup path. Nokia SR OS BGP behavior, including best-path selection and multipath operation, is documented in the Nokia SR OS documentation portal. The general BGP route-selection framework is also defined by the IETF in RFC 4271.

QUESTION NO: 18

Click the exhibit.

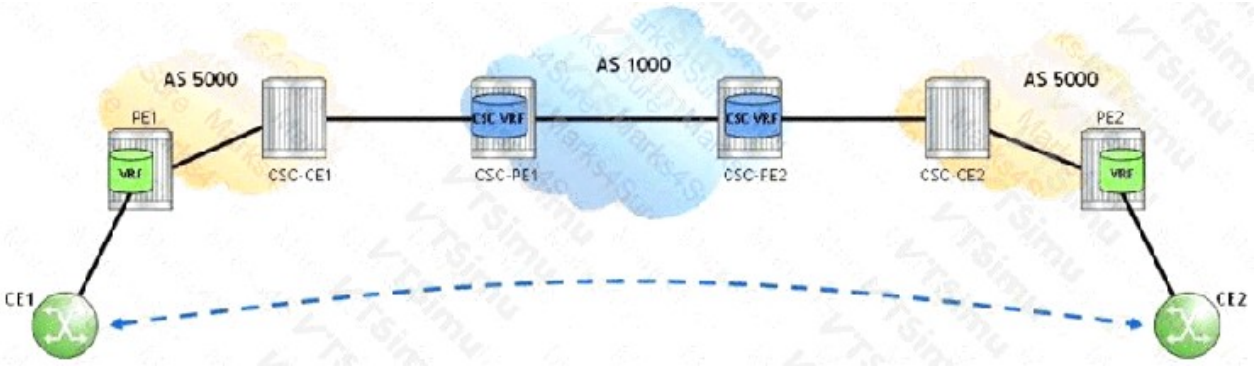

A Carrier Supporting Carrier (CSC) VPRN is configured on an Alcatel-Lucent 7750 SR for a customer carrier offering VPRN services. Which of the following configuration steps is NOT required?

A. Creating an export policy on CSC-CE1 and CSC-CE2 to advertise local PE addresses to the network provider.

B. Configuring eBGP between CSC-CE and CSC-PE routers; setting the address family to IPv4; and enabling the advertisement of labels for IPv4 routes.

C. Configuring iBGP between PE and CSC-CE within each customer carrier site; setting the address family to IPv4; and enabling the advertisement of labels for IPv4 routes.

D. Configuring MP-iBGP between PE1 and PE2; setting the address family to IPv4; and enabling the advertisement of labels for IPv4 routes.

ANSWER: D

Explanation:

Configuring MP-iBGP between PE1 and PE2; setting the address family to IPv4; and enabling the advertisement of labels for IPv4 routes is not required for the CSC VPRN arrangement shown. In a Carrier Supporting Carrier design, the customer carrier's PE-loopback reachability is transported across the network-provider VPRN. The customer carrier exchanges labeled IPv4 routes with the network-provider edge through the CSC attachment, allowing the customer carrier to retain MPLS/BGP reachability for its own downstream VPN services. The CSC-facing BGP sessions therefore require labeled-unicast IPv4 route advertisement, and the relevant customer-carrier PE addresses must be exported into the provider service. The network-provider VPRN supplies the connectivity between the provider edge routers; it does not require a separate labeled-IPv4 MP-iBGP session directly between PE1 and PE2 for this purpose. Nokia's CSC implementation is based on exchanging IPv4 routes with BGP labels at the CSC interface, while the VPRN service provides the provider-core transport context. See Nokia's CSC VPRN documentation and the IETF's BGP/MPLS IP VPN architecture.

QUESTION NO: 19

On the Alcatel-Lucent 7750 SR, where can network-queue policies be applied? (Choose two)

A. On the network egress port.

B. On the SDP.

C. On the access ingress port.

D. On the egress MDA.

E. On the SAP.

F. On the ingress MDA.

ANSWER: A F

Explanation:

On the 7750 SR, a network-queue policy defines the queueing behavior used for network-facing traffic, including queue parameters and the forwarding-class-to-queue association. Its attachment point depends on the forwarding direction and the hardware resource being configured. For egress network traffic, the policy is applied at the network egress port, where the router schedules and transmits traffic toward the connected network. This allows the egress port's network queues to enforce the configured QoS treatment before frames leave the router.

For ingress network traffic, the policy is applied at the ingress MDA. The ingress MDA owns the relevant receive-side queueing resources, so applying the policy there establishes how network traffic entering ports associated with that MDA is classified into and handled by ingress network queues. These locations reflect the SR OS separation between ingress queue resources at the MDA level and egress queue resources at the physical port level. Nokia's SR OS QoS documentation describes network queue policies and their ingress and egress attachment contexts in the SR OS documentation library; platform and operating-system information is also available from Nokia's SR OS product page.

QUESTION NO: 20

Which of the following should be performed if a service provider does not want packets to make use of the shared buffer space when queued?

A. Set the CBS to 0.

B. Set the MBS equal to the CBS.

C. Set the high-priority-only equal to the MBS.

D. Set the MBS to 0.

E. Set the high-priority-only to 0.

ANSWER: B

Explanation:

Set the MBS equal to the CBS. In Nokia SR OS QoS queue management, the Committed Buffer Size (CBS) defines the queue buffer space that is committed or reserved for the queue, while the Maximum Buffer Size (MBS) defines the maximum amount of buffer space the queue may consume. The portion between CBS and MBS represents the queue's potential use of shared buffer resources. Packets may use this shared space when the queue grows beyond its committed allocation, subject to the platform's buffer-management behavior and available shared resources. Configuring MBS to exactly match CBS removes that additional range: the queue's maximum permitted buffer occupancy is then limited to its committed allocation. Consequently, queued packets cannot consume shared buffer space. This setting is appropriate where strict buffer isolation and predictable per-queue resource limits are required, although the selected CBS must still be sized carefully for the expected burst characteristics of the traffic. Nokia's SR OS documentation describes queue buffer management and the relationship of committed and maximum buffer thresholds in its QoS material; see the Nokia SR OS documentation portal and the SR OS QoS overview.

QUESTION NO: 21

Which of the following statements is true regarding the Route Distinguisher? (Choose three)

A. It is an 8 byte value containing 3 fields

B. The Route Distinguisher must have the same value as the Route Target.

C. The Administrator field contains either an AS number or an IP number

D. To ensure uniqueness, Route Distinguishers are assigned by the IANA.

E. The Route Distinguisher is not used if the customer addresses do not overlap

F. The Assigned Number field must contain a public IP address

G. The Route Distinguisher is prepended to an IPv4 prefix to create a globally unique VPN-IPv4 address.

ANSWER: A C G

Explanation:

A Route Distinguisher is an 8-octet value used in BGP/MPLS IP VPNs to make customer IP prefixes globally unique in the VPN-IPv4 address family. It has a type field and fields that identify an administrator and an assigned number; the exact field lengths depend on the RD type. Therefore, “It is an 8 byte value containing 3 fields” accurately describes the general RD format. The administrator portion can identify the assigning service provider through an autonomous system number or an IPv4 address, so “The Administrator field contains either an AS number or an IP number” is also correct. For example, the standard RD formats include two-octet ASN-based, IPv4-address-based, and four-octet ASN-based forms.

The RD is prepended to a customer IPv4 prefix to create a 12-byte VPN-IPv4 address. This allows identical customer prefixes, such as the same private IPv4 subnet used by different VPN customers, to coexist as distinct routes in the provider’s BGP control plane. The RD itself does not define VPN membership; route targets perform import and export policy functions. Nokia’s VPN documentation and the base BGP/MPLS VPN specification describe this separation between route uniqueness and route distribution. See RFC 4364, section 4.2 and Nokia SR OS Route Distinguisher documentation.

QUESTION NO: 22

Which of the following statements are true regarding P devices in an MPLS VPRN? (Choose two.)

A. Participate in service provider core routing

B. P devices are not required to be MPLS enabled. MPLS is only required on the PE devices

C. Run a common routing protocol with the CE router

D. Must support MP-BGP

E. Do not have any connections to the CE

F. Must be aware of the VPRNs

ANSWER: A E

Explanation:

“Participate in service provider core routing” is correct because a provider (P) router is part of the service-provider transport core. It participates in the provider’s internal routing and label-switched forwarding infrastructure, allowing it to forward labeled packets between provider-edge routers. In a VPRN/MPLS Layer 3 VPN, the P router’s role is to provide scalable transit across the backbone; it uses the transport information needed to reach the next hop and swap or forward MPLS labels through the core. RFC 4364 describes this PE-to-PE VPN transport model, in which the provider backbone carries VPN traffic without requiring every transit router to maintain customer VPN routing state: RFC 4364: BGP/MPLS IP Virtual Private Networks.

“Do not have any connections to the CE” is also correct. Customer-edge devices attach to provider-edge routers, where customer-facing interfaces, VPRN service contexts, and customer route exchange are maintained. A P router is an interior backbone device, so it has no direct CE-facing attachment and does not need per-customer VPRN awareness. This separation is a core scalability property of MPLS VPRNs: PE routers impose and remove VPN-related labels at the service edge, while P routers forward transit traffic across the MPLS core. Nokia’s Layer 3 service documentation describes VPRN operation and the PE role in delivering routed VPN services: Nokia SR OS Layer 3 Services Guide.

QUESTION NO: 23

What is the result of configuring the following policy statement as a BGP import policy on the Alcatel-Lucent 7750 SR?

```
policy-statement "Policy-1"
    entry 10
        from
            protocol bgp
            prefix-list "List-1" "List-2"
        exit
        action reject
            exit
    exit
    default-action accept
```

A. All BGP routes are accepted.

B. All routes matching prefix lists " List-1 " and " List-2 " are rejected.

C. BGP routes matching prefix lists " List-1 " or " List-2 " are rejected.

D. BGP routes matching prefix lists " List-1 " and " List-2 " are rejected.

ANSWER: C

Explanation:

BGP routes matching prefix lists " List-1 " or " List-2 " are rejected. In SR OS routing policy evaluation, multiple prefix-list match conditions of the same type within a policy entry are evaluated as alternatives: a route matches the entry when its prefix matches either referenced prefix list. Once the entry is matched, its `reject` action is applied to that BGP route during import processing, preventing the route from being accepted into the BGP routing information base. This is a common way to construct an inbound BGP filter that blocks multiple independently defined groups of prefixes while retaining the ability to manage each group in a separate prefix list. The behavior follows SR OS policy semantics, where match expressions of different types can be combined for more specific matching, but repeated instances of the same match criterion such as prefix lists provide an OR relationship. Nokia's routing-policy documentation describes policy entries, match criteria, and route actions in the SR OS route policies guide. BGP import-policy application is also covered in Nokia's SR OS BGP documentation.

QUESTION NO: 24

Click the exhibit.

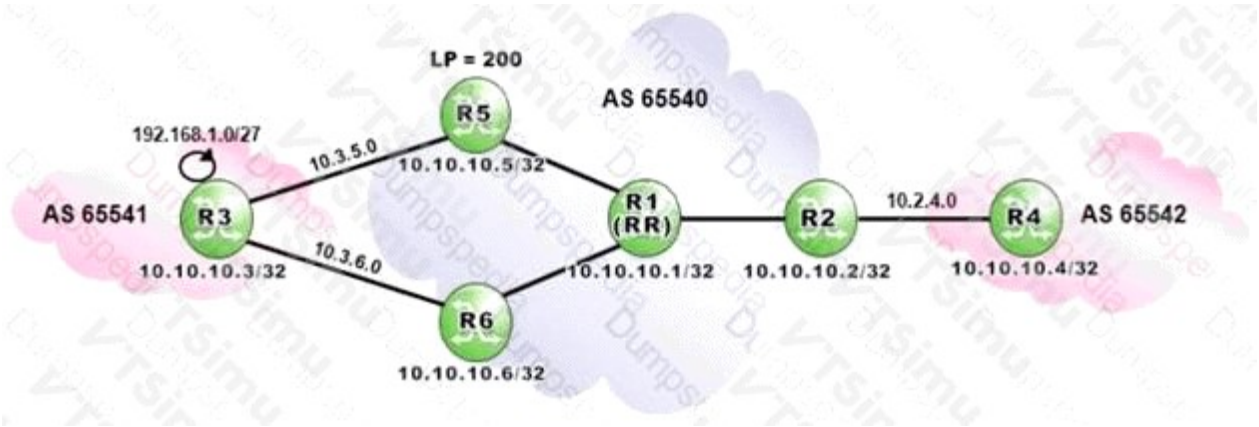

Router R1 is a route reflector with clients R2, R5 and R6. Prefixes advertised by router R5 have a local preference of 200. Router R3 advertises the prefix 192.168.1.0/27 to routers R5 and R6.

Assuming that router R6 is configured with "advertise-external", what is the expected output of "show router bgp routes" when executed on router R1?

A.

```
Flag  Network                                   LocalPref   MED
      Nexthop                                   Path-Id     VPNLabel
      As-Path
-------------------------------------------------------------------------------
u*>i  192.168.1.0/27                            200         None
      10.10.10.5                                None        -
      65541
i     192.168.1.0/27                            None        None
      10.10.10.6                                None        -
      65541
```

B.

```
Flag  Network                                   LocalPref   MED
      Nexthop                                   Path-Id     VPNLabel
      As-Path
-------------------------------------------------------------------------------
u*>i  192.168.1.0/27                            200         None
      10.3.5.3                                  None        -
      65541
i     192.168.1.0/27                            200         None
      10.10.10.5                                None        -
      65541
i     192.168.1.0/27                            100         None
      10.10.10.6                                None        -
      65541
```

C.

```
Flag  Network                                   LocalPref   MED
      Nexthop                                   Path-Id     VPNLabel
      As-Path
-------------------------------------------------------------------------------
u*>i  192.168.1.0/27                            None        None
      10.10.10.6                                None        -
      65541
i     192.168.1.0/27                            200         None
      10.10.10.5                                None        -
      65541
```

D.

```
Flag  Network                                   LocalPref   MED
      Nexthop                                   Path-Id     VPNLabel
      As-Path
-------------------------------------------------------------------------------
u*>i  192.168.1.0/27                            200         None
      10.10.10.5                                None        -
      65541
i     192.168.1.0/27                            100         None
      10.10.10.6                                None        -
      65541
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: D

Explanation:

The output represented by *Option D* is correct because R1 receives two relevant paths for 192.168.1.0/27 from its route-reflector clients. R5 advertises the path learned from R3 with a Local Preference of 200. Since Local Preference is evaluated early in the BGP best-path decision process, this is R1's preferred and active BGP route for the prefix. The selected route therefore shows R5 as the next hop/path source and retains the Local Preference value of 200.

R6 also has an eBGP-learned route from R3. Normally, after R6 receives the more-preferred reflected iBGP route via R1, it would advertise only its overall best path to iBGP peers. The `advertise-external` setting changes that advertisement behavior: it permits R6 to advertise its best external path even when that path is not its overall selected BGP route. Consequently, R1 can display R6's externally learned route as an additional, non-active path while still selecting the higher-local-preference path received from R5.

This behavior follows SR OS BGP route-reflection and best-path processing, including the significance of Local Preference and external-path advertisement. See Nokia's SR OS documentation portal and the SR OS BGP documentation.

QUESTION NO: 25

In a VPRN, which of the following are supported as PE-CE routing methods on the Alcatel-Lucent 7750 SR? (Choose four).

A. Static

B. RIP

C. OSPF

D. IS-IS

E. MP-BGP

F. BGP

ANSWER: A B C F

Explanation:

Static, RIP, OSPF, and BGP are supported PE-CE routing methods for a VPRN on the 7750 SR. A VPRN is a Layer 3 VPN service in which the provider-edge router maintains a separate customer routing context for each VPN. Static routing is available where routes are configured directly, making it suitable for small or stable customer sites. RIP and OSPF provide dynamic IPv4 route exchange between the PE and CE; RIP supports simpler deployments, while OSPF is used where a link-state IGP and faster convergence are required. BGP can also run between the PE and CE, allowing the customer and provider to exchange VPN reachability using policy-controlled path-vector routing. These methods operate in the VPRN service context and install customer reachability into the appropriate VPN routing table, while VPN route distribution across the provider MPLS backbone is handled separately by the provider network. Nokia's SR OS documentation describes VPRN as a service supporting customer route exchange through static routes and supported routing protocols, and SR OS provides the routing-service framework used for these PE-CE adjacencies. See the Nokia SR OS documentation portal and Nokia's Service Router Operating System overview.

QUESTION NO: 26

Click the exhibit button below. Given the SAP-ingress policy, which of the following statements are TRUE? (Choose three)

```
A:sr1a>config>qos>sap-ingress# info
----------------------------------------------
            queue 1 create
            exit
            queue 3 profile-mode create
                rate 50000 cir 10000
             exit
            queue 5 create
                parent "high" level 7 weight 100 cir-level 7
                rate 30000 cir 20000
             exit
            queue 6 create
                parent "high" level 8 weight 100 cir-level 8
                rate 2000 cir 2000
             exit
            queue 11 multipoint create
            exit
            fc "af" create
                queue 3
                profile in
            exit
            fc "ef" create
                queue 6
            exit
            fc "h2" create
                queue 5
            exit
            ip-criteria
                entry 10 create
                    match protocol tcp
                        dst-port eq 1234
                    exit
                    action fc "ef" priority high
                exit
                entry 15 create
                    match protocol udp
                    exit
                    action fc "h2"
                exit
                entry 20 create
                    match protocol icmp
                    exit
                    action fc "h2"
                exit
            exit
            default-fc af
```

A. All AF traffic will be marked " in-profile " at the SAP-ingress.

B. All TCP traffic will be mapped to forwarding class EF.

C. UDP traffic is placed in queue 3 while ICMP traffic is placed in queue 5.

D. EF traffic will receive higher scheduling priority than H2 traffic.

E. Traffic that does not match any of the IP criteria will be placed in queue 3.

ANSWER: A D E

Explanation:

All AF traffic will be marked “in-profile” at the SAP-ingress because the policy’s AF classification entry assigns the AF forwarding class and explicitly applies the in-profile setting. On SR OS, ingress classification establishes both the forwarding

class and the packet profile used for subsequent QoS treatment. This allows the SAP-ingress policy to identify AF-marked traffic and retain it as in-profile traffic as it enters the service.

EF traffic will receive higher scheduling priority than H2 traffic because the forwarding-class-to-queue treatment defined by the policy gives the EF queue a higher priority level than the H2 queue. Queue priority is used by the scheduler when both queues have eligible traffic, so the configured priority relationship determines which class receives preferential service.

Traffic that does not match any of the IP criteria will be placed in queue 3 because the policy's default classification behavior maps unmatched packets to the default forwarding class/queue treatment associated with queue 3. A SAP-ingress policy therefore provides a defined disposition even when no explicit IP match criterion applies. Nokia's QoS documentation describes SAP ingress classification, forwarding-class assignment, profile marking, and queue mapping in the SR OS SAP ingress QoS guide and the SR OS Quality of Service documentation.

QUESTION NO: 27

An Alcatel-Lucent 7750 SR receives a route from an eBGP peer with a MED value of 500. What

MED value is sent to iBGP peers?

A. 100

B. The IGP cost

C. 500

D. None

ANSWER: C

Explanation:

500 is correct. The Multi-Exit Discriminator (MED) is a BGP path attribute used to indicate a preferred entry point into the advertising autonomous system when multiple interconnection points exist. When the 7750 SR receives the route from its external BGP neighbor, the received MED of 500 is retained as part of that route's BGP path information. Advertising the route to internal BGP peers does not, by default, alter or remove the MED attribute, so the iBGP peers receive MED 500.

This behavior lets all BGP speakers within the local autonomous system retain the external neighbor's expressed preference and apply consistent route-selection information where MED comparison is applicable. MED is an optional non-transitive attribute: it is carried within the receiving AS via iBGP but is not normally advertised onward to a different external AS. A routing policy could explicitly modify, clear, or set a different MED, but the question specifies no such policy; therefore, the received value is propagated unchanged to iBGP peers. This handling follows the BGP MED definition in RFC 4271, section 5.1.4 and is consistent with the Nokia SR OS BGP documentation at Nokia SR OS Routing Protocols Guide—BGP.

QUESTION NO: 28

Which of the following statements BEST describes the IntServ model?

A. Has the ability to distinguish traffic in 32 different ways.

B. Uses the RSVP protocol to create microflows through a network.

C. Has the ability to manually configure foiwarding classes through a network.

D. Uses the RSVP protocol to create macroflows through a network.

ANSWER: B

Explanation:

Uses the RSVP protocol to create microflows through a network. Integrated Services (IntServ) is a QoS architecture that provides resource reservation on a per-flow basis. An application signals its traffic requirements using the Resource

Reservation Protocol (RSVP), and each RSVP-capable router along the selected path maintains soft state and performs admission control before committing resources such as bandwidth and buffer space. This establishes a distinct reservation for the individual traffic flow, commonly called a microflow, enabling IntServ to offer quantifiable service commitments for flows that require predictable performance. RSVP reservations are receiver-oriented and must be periodically refreshed, allowing the network to adapt when receivers or paths change. The per-flow state and signaling requirements are central characteristics of IntServ and are also why its deployment is generally less scalable in large core networks than aggregate-oriented QoS approaches. Nokia's QoS architecture documentation describes RSVP as supporting per-service and per-flow resource reservations, while the IETF's Integrated Services architecture defines the model's resource-reservation framework and service guarantees. See RFC 1633: Integrated Services in the Internet Architecture and RFC 2205: Resource ReSerVation Protocol (RSVP).

QUESTION NO: 29

Click the exhibit.

For the Inter-AS model B VPRN, which of the route updates for prefix 192.168.1.0/27 is incorrect?

A. The route update sent from PE1.

B. The route update sent from PE2.

C. The route update sent from ASBR1.

D. The route update sent from ASBR2.

ANSWER: B

Explanation:

The route update sent from PE2 is the incorrect update. In an Inter-AS model B VPRN, the ASBRs exchange VPN-IPv4 routes directly using MP-eBGP, including the VPN label needed to reach the egress PE. The advertising ASBR changes the BGP next hop when it advertises the VPN route to the ASBR in the neighboring autonomous system. The receiving ASBR then distributes that VPN-IPv4 route internally toward PE2 using MP-iBGP. PE2 resolves the VPN route's BGP next hop through the MPLS transport path toward the advertising ASBR; the packet is subsequently forwarded across the inter-AS boundary and onward to the egress PE using the appropriate label stack. PE2 does not originate a replacement VPN-IPv4 advertisement for the remote customer prefix as part of this control-plane propagation. When PE2 advertises the learned route to its attached CE, that CE-facing advertisement is an ordinary IPv4 route within the relevant VPRN context and does not carry VPN-IPv4 route-target or MPLS VPN-label information. This behavior follows the Inter-AS option B control-plane model defined for BGP/MPLS IP VPNs. See RFC 4364 and Nokia's VPRN documentation.

QUESTION NO: 30

In a VPRN the PE device is configured to run which of the following protocols? (Choose three.)

A. MP-BGP for exchanging customer routes with other PEs

B. A routing protocol for exchanging customer routes with the CE

C. MPLS for exchanging labels with other provider core devices

D. MPLS for exchanging labels with the CE devices

E. Targeted LDP for exchanging VPRN labels with other PE devices

F. A label signaling protocol for defining transport tunnels between PE and CE devices

ANSWER: A B C

Explanation:

A VPRN PE participates in both the provider MPLS backbone and the customer's Layer 3 VPN routing domain. **MP-BGP for exchanging customer routes with other PEs** is used to distribute VPN-IPv4 or VPN-IPv6 reachability between PE routers. These routes carry the VPN route-target information needed to import routes into the appropriate VPRN and identify the VPN context across the provider network. **A routing protocol for exchanging customer routes with the CE** is also required at the PE–CE interface. Depending on the service design, this can be static routing, BGP, OSPF, IS-IS, RIP, or another supported PE–CE routing method, operating within the VPRN context. Finally, **MPLS for exchanging labels with other provider core devices** provides the label-switched transport through the provider backbone. The PE imposes an outer transport label for the LSP across the core and an inner VPN/service label that identifies the destination VPRN at the remote PE. Together, VPN route distribution, PE–CE route exchange, and MPLS transport allow customer packets to reach remote VPRN sites while remaining logically separated from other VPNs. See Nokia's VPRN overview documentation and RFC 4364.

QUESTION NO: 31

The following question relates to this command syntax: "config>service# vpls 9000 customer 6 create". Which of the following statements are correct? (Choose 2)

A. The customer-id in this command is 9000.

B. The service-id in this command is 9000.

C. This value 9000 should be unique across the network as a best practice.

D. The vc-id will be 6 for the service by default.

E. The egress vc-label will be 9000.

ANSWER: B C

Explanation:

In the SR OS service configuration hierarchy, the syntax for creating a VPLS includes the service identifier immediately after the `vpls` keyword and the customer identifier after the `customer` keyword. Therefore, in `vpls 9000 customer 6 create`, `9000` identifies the VPLS service and `6` identifies the customer associated with that service. The statement "The service-id in this command is 9000." is consequently correct.

Using a service ID that is unique throughout the provider network is also an important operational best practice. Although service objects are configured locally on each router, a consistent network-wide service-ID plan makes provisioning, fault isolation, operational tooling, and cross-node service correlation much easier. It is especially valuable when the same VPLS is instantiated on multiple provider-edge routers, because engineers can identify the intended service directly from its common identifier. Thus, "This value 9000 should be unique across the network as a best practice." is correct. Nokia's SR OS documentation describes VPLS as a service configured using a service ID and customer ID under the service context; see the Nokia SR OS VPLS services overview and the Nokia SR OS service model documentation.

QUESTION NO: 32

Click the exhibit.

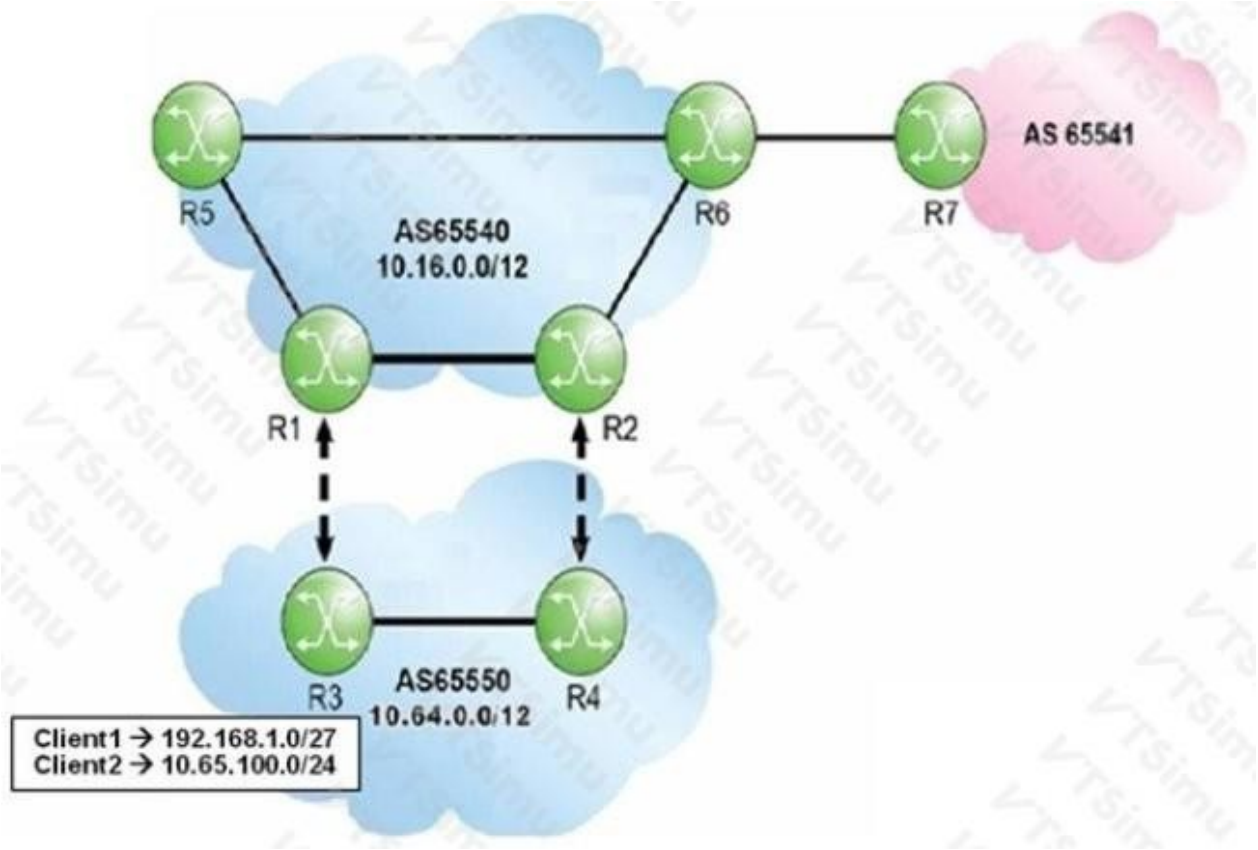

AS 65550 owns CIDR block 10.64.0.0/12. Router R3 is advertising Client1 and Client2 networks into BGP. Router R6 advertises the aggregate 10.64.0.0/12 to AS 65541 with "summary-only" enabled. Which of the following routes are received by router R7?

A. Client1, Client2 and the aggregate.

B. Only the aggregate.

C. Only Client1 and the aggregate.

D. Only Client2 and the aggregate.

ANSWER: C

Explanation:

Only Client1 and the aggregate. is correct. The `summary-only` setting on R6 controls advertisements made by R6 toward its BGP peer in AS 65541. R6 advertises the aggregate route, 10.64.0.0/12, rather than propagating the contributing client-specific routes over that advertisement path. This lets AS 65541 use a single covering route for destinations within the block while reducing the number of routes it must receive and maintain.

As shown in the exhibit, R7 also has an eligible BGP path through which it receives the Client1 prefix independently of R6's summarized advertisement. Therefore, R7 installs or receives both the independently learned Client1 route and the aggregate announced by R6. The aggregate remains useful as the covering route for destinations within 10.64.0.0/12 that do not match a more-specific route. BGP forwarding uses longest-prefix matching, so traffic for Client1 follows its more-specific prefix, while other matching destinations can use the /12 aggregate.

Nokia SR OS documents aggregate-route advertisement and suppression behavior in its BGP documentation. General BGP route-selection and longest-prefix forwarding concepts are also described in RFC 4271.

QUESTION NO: 33

Which of the following entities can scheduler policies be applied to? (Choose two)

A. Network ingress port

B. Epipe ingress SAP

C. ES egress SAP

D. Ingress MDA

E. VPRN ingress interface

F. Shared buffer space

ANSWER: B C

Explanation:

Scheduler policies provide hierarchical scheduling control within the service QoS context. They define scheduler levels, scheduling behavior, rate parameters, and the association of forwarding classes or queues with scheduler nodes. An **Epipe ingress SAP** is a supported attachment point because ingress service traffic entering an Epipe through a SAP can be subject to a scheduler-policy hierarchy. This permits the service to control how traffic received at that SAP is scheduled and rate-managed before it is handled by the service forwarding path.

An **ES egress SAP** is also a supported attachment point. On service egress, the scheduler policy establishes the hierarchy used to schedule traffic toward the SAP, allowing bandwidth and priority treatment to be applied to the service's outbound queues and forwarding classes. These SAP-level ingress and egress applications are service-QoS uses of scheduler policies and allow QoS behavior to be tailored per service attachment rather than only through physical-resource configuration.

Nokia's QoS documentation describes scheduler-policy configuration and its use in service QoS contexts, including SAP ingress and egress scheduling. See the Nokia SR OS Scheduler Policies documentation and the Nokia SR OS Service QoS documentation.

QUESTION NO: 34

Click the exhibit.

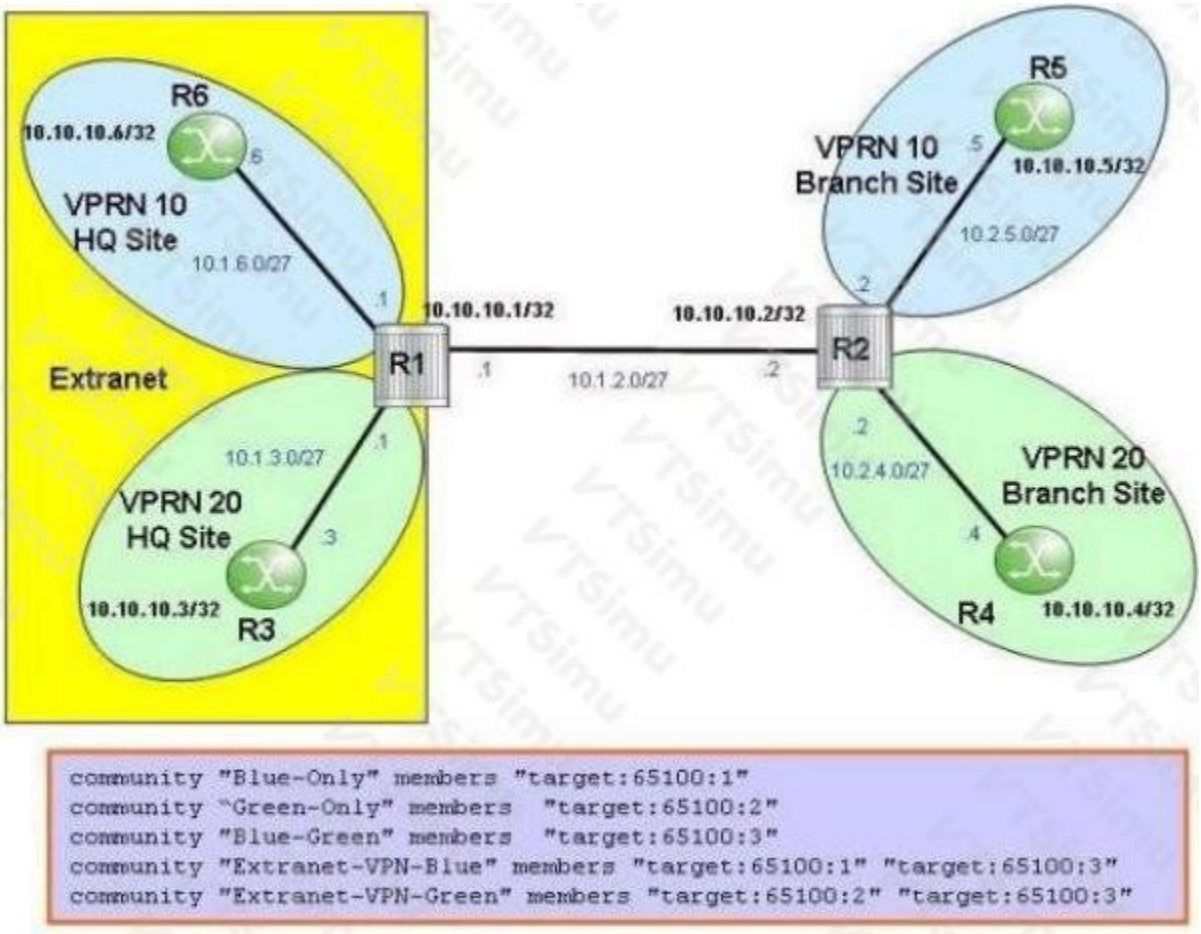

The headquarter sites of VPRN 10 and VPRN 20 are part of an extranet VPRN. Which route targets should be exported under the VPRN 20 service on router R1?

A. Target:65100:2 and target:65100:3

B. Target:65100:1 and target:65100:2

C. Target:65100:1 and target:65100:3

D. Target:65100:1, target:65100:2 and target:65100:3

ANSWER: A

Explanation:

Target:65100:2 and target:65100:3 is correct. In an MPLS Layer 3 VPN, an exported route target is attached to the VPN-IPv4 route as a BGP extended community. Other VPRN services can then import that route only when their import policy includes a matching route target. This allows the normal VPRN membership and the selective extranet connectivity to coexist.

For VPRN 20, **target:65100:2** identifies routes as belonging to the VPRN 20 route-target membership. Exporting it ensures that routes originated by the VPRN 20 headquarters can be imported by the intended VPRN 20 sites. **target:65100:3** is the shared extranet route target shown for the headquarters interconnection. Applying that target to exported VPRN 20 routes makes the selected VPRN 20 headquarters prefixes available to the other extranet participant, VPRN 10, whose extranet import policy includes the same shared target.

Using both targets therefore preserves VPRN 20 connectivity while advertising its headquarters routes into the common extranet route-target domain. This is the standard route-target mechanism for controlled VPN route distribution, as defined in RFC 4364 and implemented in Nokia SR OS L3 VPN services; see the Nokia SR OS Layer 3 VPN documentation.

QUESTION NO: 35

Which of the following statements regarding BGP route selection is TRUE?

A. Lower local preference is preferred over higher local preference.

B. Higher MED is preferred over lower MED.

C. Higher origin code is preferred over lower origin code.

D. Lower BGP router ID is preferred over higher BGP router ID.

ANSWER: D

Explanation:

“Lower BGP router ID is preferred over higher BGP router ID.” is correct. In the BGP best-path selection process, the router ID is a late tie-breaker used when candidate paths remain otherwise indistinguishable after the preceding applicable selection criteria have been evaluated. The BGP speaker compares the router IDs associated with the paths and selects the path associated with the numerically lowest router ID. Router IDs are represented as 32-bit values, conventionally displayed in dotted-decimal notation, so the comparison is numerical rather than a textual comparison of the displayed address.

This tie-breaker provides deterministic path selection: routers evaluating equivalent candidate paths can consistently choose the same preferred route based on an unambiguous identifier. In operational practice, the router ID is normally configured explicitly or derived from an address according to platform configuration rules; network designers should therefore ensure that each BGP speaker has a unique, stable router ID. The BGP decision process and its tie-breaking behavior are specified in RFC 4271, A Border Gateway Protocol 4 (BGP-4). Nokia’s routing documentation and product resources are available through the Nokia Documentation Center.

QUESTION NO: 36

What are the two primary problems the service provider must consider when providing traditional Layer 3 VPN services using only a single common routing table in the provider core? (Choose two)

A. Memory exhaustion in the provider core

B. Route leaking between the customer networks

C. CPU utilization for route processing

D. Unwanted packet forwarding between customer networks

ANSWER: B D

Explanation:

Route leaking between the customer networks and **Unwanted packet forwarding between customer networks** are the fundamental isolation issues when a provider uses one common routing table for multiple Layer 3 VPN customers. Customer VPNs require separate routing domains: each customer must see and use only its own prefixes, even where different customers use identical private IPv4 address ranges. A shared provider routing table has no inherent per-customer route context, so routes learned from one customer can become visible or be selected for another customer unless the provider applies extensive and error-prone filtering. This is route leakage at the control-plane level.

The corresponding data-plane concern is that packets can be forwarded toward another customer based on entries in that same shared table. Even if prefix advertisements are filtered carefully, maintaining complete customer separation through policy alone is operationally difficult, particularly as customers, sites, and routes change. Modern Layer 3 VPN designs address both requirements by associating routes and forwarding information with VPN-specific routing/forwarding instances and using route distinguishers and route targets to preserve separation. RFC 4364 describes how BGP/MPLS IP VPNs provide distinct customer routing contexts and prevent ambiguous overlapping routes; see RFC 4364. Nokia’s Service Router documentation also covers Layer 3 VPN routing and forwarding separation: Nokia Service Router documentation.

QUESTION NO: 37

Click the exhibit.

If router A originates a BGP route for prefix 192.168.0.1/27, what will the update contain when it reaches router B?

A. AS Path of 65200, Next Hop of router A, Origin of IGP.

B. AS Path of 65200, Next Hop of router A, Origin of incomplete.

C. Null AS Path, Next Hop of router A, Origin of IGP.

D. Null AS Path, Next Hop of router A, Origin of incomplete.

ANSWER: C

Explanation:

Null AS Path, Next Hop of router A, Origin of IGP. is correct. When router A originates the stated prefix directly into BGP, the route is locally originated in its BGP routing information base. A locally originated BGP route has an empty AS_PATH because it has not yet traversed any autonomous system. AS_PATH prepending occurs when a BGP speaker advertises a route to an external BGP peer; the originating router does not insert its own AS number into the route before making that initial external advertisement. Router A is the advertising BGP speaker, so its address is carried as the NEXT_HOP attribute for the advertisement to router B. The ORIGIN attribute is IGP for a route introduced into BGP through the BGP network-origination mechanism. This identifies the route as originating from an interior-routing source as defined by BGP path-attribute semantics. RFC 4271 defines the construction of AS_PATH for locally originated routes and specifies the ORIGIN attribute values and meaning; see RFC 4271, AS_PATH and RFC 4271, ORIGIN.

QUESTION NO: 38

Click on the exhibit below.

```
*A:PE1# show service id 61 sdp 3 detail | match Pw
Peer Pw Bits       : pwFwdingStandby
```

The MDU attached to this PE has been configured with active/standby SDP's configuration option has been configured for the PE to recognize that this SDP is in standby mode?

A. On the PE the endpoint has been configured with "supress-standby-signalling"

B. On the PE the endpoint has been configured with "no supress-standby-signalling"

C. On the MDU the endpoint has been configured with "supress-standby-signalling"

D. On the MDU the endpoint has been configured with "no supress-standby-signalling1

E. On the MDU the endpoint has been configured with "no suppress-standby-signalling"

ANSWER: E

Explanation:

On the MDU the endpoint has been configured with "no suppress-standby-signalling" is correct. In an active/standby spoke-SDP protection design, the PE needs to learn the remote endpoint's operational protection state so that it can treat the corresponding SDP as standby. This state is signalled by the remote system hosting the endpoint—in this case, the MDU—not by the receiving PE. Configuring `no suppress-standby-signalling` on the MDU endpoint permits the endpoint's standby status to be advertised across the SDP control plane. The attached PE can then receive that signalling and recognize that the associated remote SDP endpoint is in standby state. This ensures the PE forwards traffic using the active path while retaining the standby SDP as the protected alternative, allowing the intended switchover behavior if the active path becomes unavailable. Nokia's service documentation describes endpoint status signalling and spoke-SDP redundancy as part of Layer 2 VPN service operation; see the Nokia SR OS VPLS service documentation and the Nokia documentation portal.

QUESTION NO: 39

Which of the following parameters are used in a hierarchical scheduler-policy to specify the priority of the child scheduler/queue for allocating bandwidth beyond its committed rate? (Choose two)

A. CIR

B. CIR level

C. CIR weight

D. Rate

E. Level

F. Weight

ANSWER: E F

Explanation:

Level and **Weight** control the distribution of bandwidth that is available beyond a child scheduler or queue's committed information rate. Level establishes the strict-priority ordering used during excess-bandwidth allocation: children assigned a higher scheduling priority are considered before children at lower priority levels. Weight then applies among children that are eligible at the same level, providing weighted sharing of the remaining bandwidth. Together, these parameters let a hierarchical scheduler policy provide both preferential treatment and proportional sharing for traffic that exceeds committed-rate service.

This distinction is important in a hierarchy because the parent scheduler first satisfies the applicable committed-rate treatment and then uses the excess-capacity scheduling attributes to arbitrate available bandwidth among its child schedulers or queues. A level therefore expresses relative precedence, while a weight expresses the relative allocation entitlement within an equal-precedence group. Nokia's QoS and scheduler documentation describes hierarchical scheduling and the scheduler-policy controls used to manage child bandwidth allocation. See the Nokia SR OS scheduler policies documentation and the Nokia SR OS QoS documentation.

QUESTION NO: 40

How does the IEEE 802.1ag loopback function differently from the loopback feature in 802.3ah? (Choose 2)

A. Theloopbackfuntionin802.1ag is an intrusive test that will stop customer data.

B. The loopback function of 802.lag is a non-intrusive "ping" sent from a MEP to a remote MEP or MIP.

C. The loopback function in 802.3ah is an intrusive test that will loop customer data.

D. The loopback function of 802.lag is an intrusive test that will "loop" all customer data back to the original sender.

ANSWER: B C

Explanation:

The loopback function of 802.1ag is a non-intrusive "ping" sent from a MEP to a remote MEP or MIP. Ethernet Connectivity Fault Management uses a Loopback Message and Loopback Reply to verify connectivity and forwarding to a specified maintenance point. The reply is generated for the diagnostic message, while ordinary service frames continue to be forwarded normally. This makes the function suitable for fault isolation and verification without placing the customer service into a loopback state.

The loopback function in 802.3ah is an intrusive test that will loop customer data. IEEE 802.3ah Ethernet in the First Mile provides remote loopback at the Ethernet OAM sublayer. Once remote loopback is enabled, frames received by the remote OAM entity are returned toward the initiating device rather than being forwarded onward through the service. Consequently, it is a disruptive diagnostic mode and must be used with awareness that normal customer traffic is looped back during the test. Nokia's Ethernet CFM documentation describes 802.1ag loopback messages and replies, while the IEEE 802.3ah standard overview identifies Ethernet OAM as part of the amendment's functionality. See Nokia Ethernet CFM documentation and IEEE 802.3ah standard information.

QUESTION NO: 41

On the Alcatel-Lucent 7750 SR, where can network-queue policies be applied? (Choose two)

A. On the network egress port.

B. On the SDP.

C. On the access ingress port.

D. On the egress MD

E. On the SAP.

F. On the ingress MDA.

ANSWER: A F

Explanation:

On the 7750 SR, network queue policies provide QoS queue configuration for traffic handled by the network side of the router. The egress direction uses a network-queue policy at the network egress port, where the policy defines the egress network queues and their scheduling, buffering, and congestion-management behavior for traffic transmitted from that port. This association is made at the port because egress transmission resources and queue scheduling are port-based.

For ingress traffic, a network-queue policy is associated with the ingress MDA. The MDA is the hardware forwarding entity that receives network traffic, so the ingress network queues and associated ingress QoS treatment are allocated at that level. This distinction reflects the 7750 SR QoS architecture: egress network queuing is configured at the network port, while ingress network queuing is configured on the receiving MDA. Nokia's QoS documentation describes network QoS policies and their association points within the SR OS forwarding architecture. See the Nokia SR OS Network QoS documentation and the Nokia SR OS QoS overview.

QUESTION NO: 42

Click on the exhibit below.

If the SDP between R2 and R1 goes down, what devices will be reachable from R2 through the VPLS?

A. Only R4 will be reachable through the VPLS.B All devices will be reachable through the VPLS.

B. R4 and R3 will remain reachable.

C. No devices will be reachable.

ANSWER: B

Explanation:

R4 and R3 will remain reachable. In the exhibited VPLS topology, the failure of the SDP between R2 and R1 affects the pseudowire/transport path associated with R1 only; it does not remove the other active VPLS service connections from R2. R2 continues to have operational VPLS connectivity toward R3 and R4, so traffic can still be forwarded to the customer-facing devices attached through those participating provider routers. The VPLS forwarding database adapts to the loss of the failed service path, while active pseudowires continue to provide Layer 2 reachability within the same VPLS instance. Thus, the service impact is scoped to the failed R2-to-R1 SDP rather than causing a complete VPLS outage at R2. This behavior follows the SR OS VPLS model, in which a VPLS service can use multiple SDPs/pseudowires and each operational service path contributes independently to connectivity. See Nokia's VPLS documentation and the SR OS Layer 2 Services overview.

QUESTION NO: 43

Which protocol is used to signal the status of the active links between the PE and the CE in MC-LAG?

A. MSTP

B. MCLAG

C. LACP

D. DHCP

E. A proprietary protocol

ANSWER: C

Explanation:

LACP is used to signal the operational status of the active member links between a customer edge device and provider edge devices in an MC-LAG deployment. Link Aggregation Control Protocol exchanges control information across the aggregated Ethernet links, allowing the customer edge device to identify which member links are eligible to forward traffic. This is essential when the MC-LAG peers present a shared logical LAG toward the customer edge while maintaining coordination between the two provider edge systems. If a member link, peer, or associated forwarding path changes state, LACP communicates the resulting aggregation state so traffic is sent only over usable active links. In SR OS MC-LAG designs, LACP therefore provides the PE-to-CE link-status signaling function and supports resilient active/standby or active/active forwarding behavior, depending on the configured mode and network design. Nokia's SR OS documentation describes MC-LAG operation and its use with LACP, while the IEEE standard defines LACP as the control protocol for link aggregation. See the Nokia SR OS MC-LAG documentation and IEEE 802.1AX Link Aggregation standard.

QUESTION NO: 44

What is the MTU impact on the SDP path-mtu if fast reroute facilities-mode is implemented on the LSP?

A. There will be no impact on the path-mtu.

B. The path-mtu will decrease by4 bytes.

C. The path-mtu will increase by 4 bytes.

D. The path-mtu will decrease by 8 bytes.

E. The path-mtu will increase by 8 bytes.

ANSWER: B

Explanation:

The path-mtu will decrease by4 bytes. RSVP-TE fast reroute in facilities mode protects an LSP by redirecting traffic onto a pre-established bypass LSP around the protected resource. While traffic is using that bypass, the router must impose an additional MPLS label to identify and transport the packet through the bypass tunnel, in addition to the label already required for the protected LSP. An MPLS label-stack entry is 4 bytes, so the effective payload capacity available along the SDP's LSP path is reduced by exactly 4 bytes. Consequently, SR OS accounts for this potential additional label when determining the SDP path MTU, ensuring that packets which fit the advertised path MTU can still be transported if a facilities-mode fast-reroute switchover occurs. This adjustment avoids oversize frames caused by the temporary extra label-stack depth during protection. Nokia's MPLS documentation describes RSVP-TE LSPs and fast-reroute behavior, including bypass-based protection, in the 7750 SR OS MPLS Guide. Additional SR OS MPLS and RSVP-TE feature documentation is available from the Nokia SR OS documentation portal.

QUESTION NO: 45

Click the exhibit.

Router R6 is a route reflector with clients R1, R2, R4 and R5. Router R4 receives three routes from router R6 and is configured with ECMP 3. Given the following BGP configuration on routerR4, how many primary and backup paths will be present in router R4's BGP routing table?

A. Three primary paths,

B. Two primary paths and one backup path.

C. One primary path and two backup paths.

D. One primary path and one backup path.

ANSWER: B

Explanation:

Two primary paths and one backup path. is correct. Router R4 receives three eligible paths for the same destination from its route reflector. The BGP multipath settings shown in the configuration control how many equivalent paths are retained as active, or primary, paths rather than merely how many paths the forwarding plane can support through ECMP. In this case, the configured primary multipath limit is two, so two of the equal eligible BGP paths are selected as primary paths. The remaining eligible path is retained as a backup path, providing a preselected alternate if one of the primary paths is withdrawn or becomes unusable. The ECMP value of three establishes that the router can use up to three equal-cost forwarding next hops, but it does not override the BGP primary-path selection limit configured for this route family or peer group. Consequently, R4's BGP route table records two active primary paths and one backup path for the prefix. Nokia's BGP documentation describes BGP path selection and multipath operation in the unicast-routing guide: Nokia SR OS BGP documentation. Additional SR OS routing documentation is available at Nokia SR OS Unicast Routing Protocols.

QUESTION NO: 46

A service provider is using GRE for his transport tunnel on the Alcatel-Lucent 7750 SR. How can traffic be marked as it traverses the service provider's network? (Choose two)

A. Using dot1p bits of the 802.1q Ethernet frame header

B. Using DSCP bits within the ToS field of the IP packet header.

C. Using the IP precedence bits within the ToS field of the IP packet header.

D. Using the CLP bit of the ATM cell header.

E. Using the EXP bits of the MPLS transport label header.

ANSWER: A B

Explanation:

Using dot1p bits of the 802.1q Ethernet frame header and using DSCP bits within the ToS field of the IP packet header are valid marking mechanisms for GRE-transported traffic on a 7750 SR. GRE encapsulates the customer packet in an outer IP header, and the provider can apply a differentiated-services code point to that outer IP header. The DSCP is the six-bit differentiated-services field used by IP networks to identify the forwarding treatment, such as the forwarding class and drop behavior, that should be applied while the packet crosses the provider network. This makes DSCP marking directly applicable to the routed IP transport carrying GRE. On Ethernet links, an 802.1Q VLAN tag can also carry the three-bit priority code point commonly called 802.1p or dot1p. Network devices can use that Layer 2 marking to select the appropriate QoS treatment on an Ethernet hop. Thus, either the outer IP DSCP marking or Ethernet dot1p marking can communicate the intended traffic class through the provider transport. GRE's use of an outer IP delivery header is defined in RFC 2784, and DSCP semantics are specified in RFC 2474.

QUESTION NO: 47

Click the exhibit button below. Traffic arriving at router PE 1 is marked with a dot1p value of 4. Given the network policies applied on the routers in the diagram below, what is the EXP value of packets arriving at router PE 2?

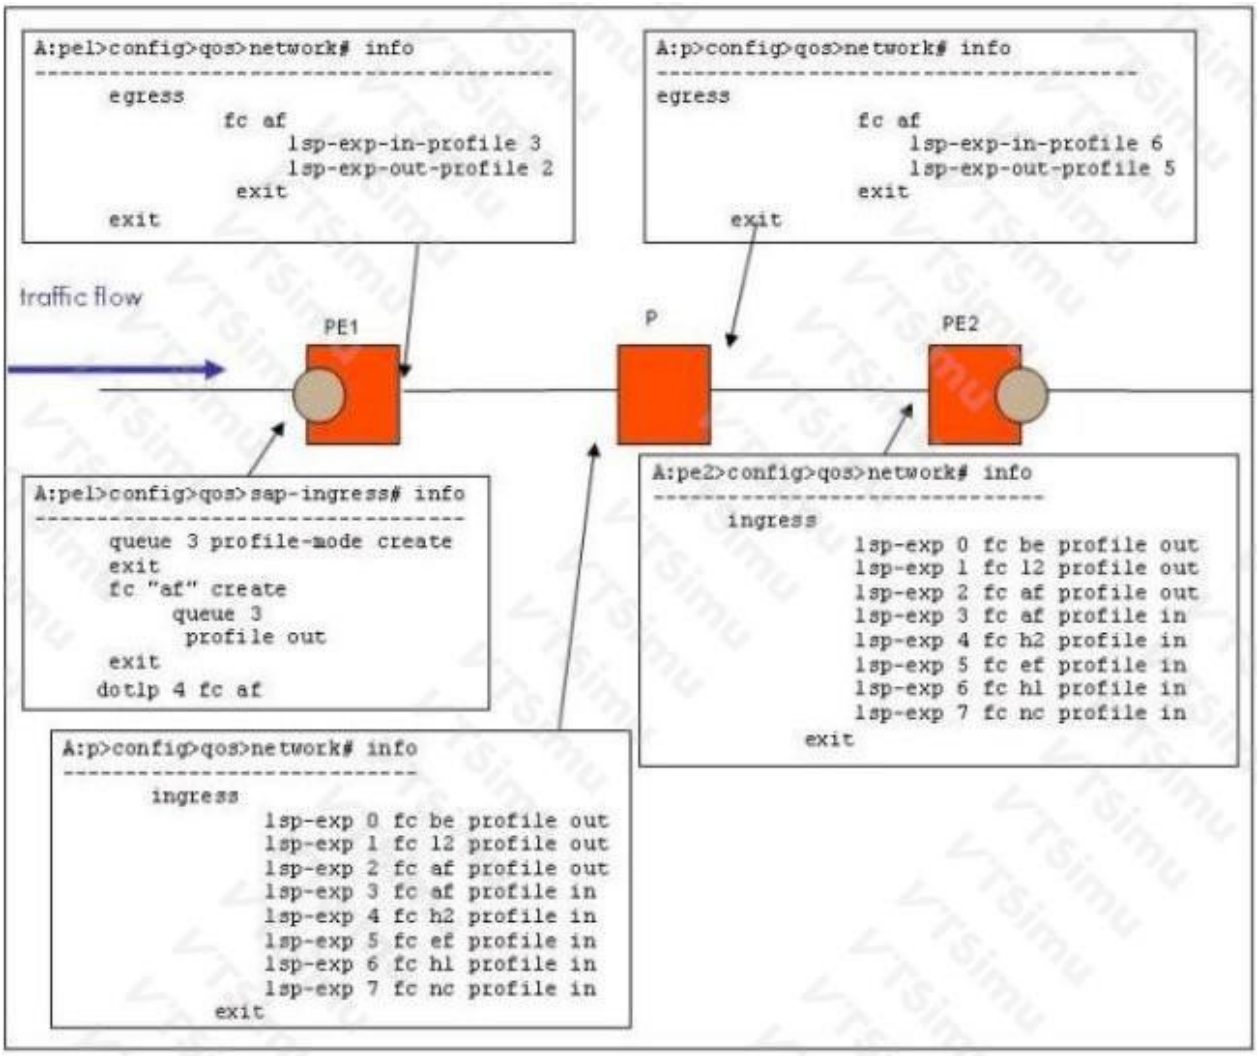

Note: Only relevant portions of the network policy are shown.

A. 2

B. 3

C. 5

D. 6

E. None of the above.

ANSWER: A

Explanation:

2 is the EXP value arriving at PE 2. The ingress classification at PE 1 uses the received IEEE 802.1p value of 4 to select the forwarding treatment defined by the applied network policy. When PE 1 pushes the MPLS transport label, the policy's forwarding-class and profile mapping determines the MPLS Traffic Class field, historically called EXP. For the classified traffic shown, that MPLS marking is set to 2.

The MPLS label carries this three-bit traffic-class value through the provider network so that core nodes can apply the intended QoS behavior without needing the original Ethernet dot1p marking. The relevant policy processing therefore occurs at ingress when the packet is classified and encapsulated; the MPLS header presented on arrival at PE 2 retains the resulting traffic-class marking. This is the expected Nokia SR OS behavior for network QoS policies and MPLS packet marking.

For Nokia's QoS concepts and the relationship between classification, forwarding classes, profiles, and packet marking, see the Nokia SR OS QoS overview and the Nokia SR OS MPLS overview.

QUESTION NO: 48

Click the exhibit button below. Given this scheduler-policy configuration, which of the following can be said about the scheduler called "high"? (Choose two)

```
A:sr1a>config>qos>scheduler-policy# info
----------------------------------------------
            tier 1
                scheduler "Root" create
                    description "This is the root scheduler"
                    rate 50000 cir 50000
                exit
            exit
            tier 2
                scheduler "high" create
                    description "High priority traffic"
                    parent "Root" level 5 cir-level 5
                    rate 30000 cir 25000
                exit
           tier 3
                scheduler "low" create
                    description "Low priority traffic"
                    parent "high" level 2 cir-level 2
                    rate 50000 cir 10000
                exit
            exit
----------------------------------------------
```

A. It is a parent scheduler.

B. It is the top-level scheduler.

C. It can allocate up to 50 Mbps of bandwidth to scheduler "low" since the bandwidth is not used by another Tier 3 scheduler.

D. It is a child scheduler.

E. It can only have queues as children.

ANSWER: A D

Explanation:

“It is a parent scheduler” and “It is a child scheduler” are correct because the configured scheduler hierarchy places *high* at an intermediate level. A scheduler is a parent when one or more lower-level schedulers are associated beneath it, allowing it to distribute its available bandwidth among those children. At the same time, *high* is a child because it is itself attached beneath another scheduler in the policy hierarchy. These are not mutually exclusive roles: intermediate schedulers commonly act as both a child of an upstream scheduling node and a parent of downstream schedulers or queues.

SR OS scheduler policies support hierarchical scheduling, in which scheduler levels establish parent-child relationships used for bandwidth control and arbitration. This hierarchy permits traffic to be shaped and prioritized at more than one level, such as an aggregate level followed by a service, customer, or queue grouping level. The scheduler named *high* therefore participates in both directions of that hierarchy: it receives scheduling treatment from its own parent and provides scheduling treatment to its descendants. Nokia’s QoS documentation describes scheduler policies and hierarchical scheduler operation in the SR OS Scheduler Policy documentation and the SR OS QoS overview.

QUESTION NO: 49

What of the following is a key benefit of Active/Standby pseudowires?

A. The MDU does not need to support MPLS.

B. Can be used in ring networks of 4 or more nodes

C. Does not require the use of a layer 2 loop prevention protocol such as RSTP

D. The MDU does not have to be aware of the active/standby configuration

ANSWER: C

Explanation:

“Does not require the use of a layer 2 loop prevention protocol such as RSTP” is a key benefit because an active/standby pseudowire redundancy design deliberately permits only one pseudowire in a redundant set to forward customer traffic at a time. The standby pseudowire is retained as a backup path and becomes active only when the currently active path fails or becomes unavailable. This controlled forwarding state prevents the duplicate Layer 2 forwarding paths that could otherwise form a bridging loop between the attachment circuits and the MPLS transport network.

As a result, redundancy and failover are provided by the pseudowire control and status mechanisms rather than by relying on a customer-facing spanning-tree topology to block one of the paths. This is especially valuable for Ethernet access and dual-homed service designs, where avoiding Layer 2 loops without introducing spanning-tree convergence behavior can improve predictability and reduce operational complexity. The IETF pseudowire redundancy framework defines active and standby pseudowire behavior, including selecting one active forwarding path and maintaining backup connectivity. Nokia service-router Layer 2 service implementations use this pseudowire-redundancy model to support resilient services.

See RFC 6718: Pseudowire Redundancy and RFC 4447: Pseudowire Setup and Maintenance Using LDP.

QUESTION NO: 50

Click on the exhibit below.

Assume a CE router is connected to the SAP of each PE. The arp cache of all CE routers is empty and the FDB for VPLS 300 is empty. The CE router connected to PE-A initiates a ping towards the CE router connected to PE-C. The ping is successful. Which of the following statements are true? (Choose 2)

A. The FDB on PE-A will contain entries for all CE routers.

B. The FDB on PE-B will contain an entry for the CE router connected to PE-A.

C. The FDB on PE-B will contain an entry for all CE routers.

D. The FDB on PE-C will contain entries for the CE router connected to PE-A and the CE router connected to PE-C.

ANSWER: B D

Explanation:

“The FDB on PE-B will contain an entry for the CE router connected to PE-A” is correct because the initial ARP request is an unknown-unicast/broadcast Ethernet frame when every CE ARP cache and VPLS forwarding database is empty. PE-A learns the source MAC address of its locally attached CE from that ingress frame and floods the request through the VPLS. PE-B receives the flooded frame and learns the CE connected to PE-A’s source MAC address against the VPLS forwarding path toward PE-A.

“The FDB on PE-C will contain entries for the CE router connected to PE-A and the CE router connected to PE-C” is also correct. PE-C learns the PE-A CE source MAC when it receives the flooded ARP request. The CE attached to PE-C then sends an ARP reply, allowing PE-C to learn that local CE’s source MAC on its SAP. The reply is forwarded toward the already learned destination MAC for the CE at PE-A, enabling the subsequent ping traffic to be unicast. This behavior follows SR OS VPLS source-MAC learning and flooding operation; see Nokia’s SR OS documentation portal and the Nokia Service Router Operating System overview.

QUESTION NO: 51

In a provider core consisting of 6 PE and 4 P routers, what is the minimum number of BGP sessions required in total for correct operation of VPRN services?

A. 0

B. 5

C. 10

D. 15

E. 45

ANSWER: D

Explanation:

15 is correct. In a traditional Nokia VPRN deployment using MP-BGP to distribute VPN-IPv4 routes, BGP peering is required between the Provider Edge routers that originate and consume customer VPN routes. The Provider routers only provide MPLS label switching across the backbone; they do not hold VPRN customer routing instances and do not need to participate in the VPN BGP mesh. Therefore, only the six PE routers are counted when determining the required BGP sessions.

Without route reflectors, the PE routers form a full iBGP mesh. The number of sessions in a full mesh is calculated as $n \times (n - 1) / 2$, where n is the number of BGP-speaking PE routers. With six PE routers, this is $6 \times 5 / 2$, which equals 15 BGP sessions. This represents the minimum for the stated topology when no route-reflector function is specified. MP-BGP distributes the VPN routes and their associated labels between PE routers, while the MPLS transport core carries the labeled packets. Nokia’s service documentation is available through the Nokia Service Router documentation portal; the underlying BGP/MPLS IP VPN control-plane model is defined in RFC 4364.

QUESTION NO: 52

Which of the following prefix-lists is the most specific match for prefixes 192.168.64.1/20 and

192.168.32.5/19?

A. Prefix 192.168.0.0/17 longer.

B. Prefix 192.168.0.0/19 longer.

C. Prefix 192.168.96.0/19 longer.

D. Prefix 192.168.0.0/18 longer.

ANSWER: A

Explanation:

Prefix 192.168.0.0/17 longer. is the correct choice because it is the most-specific listed aggregate that contains both route prefixes while also satisfying the *longer* match condition. The address range of 192.168.0.0/17 extends from 192.168.0.0 through 192.168.127.255. After applying the masks to the prefixes in the question, the relevant network routes are 192.168.64.0/20 and 192.168.32.0/19. Both networks fall within that /17 aggregate: the /20 occupies addresses beginning at 192.168.64.0, and the /19 occupies addresses beginning at 192.168.32.0. In SR OS policy terminology, *longer* requires the evaluated route to have a prefix length greater than the configured prefix length. The evaluated /20 and /19 routes are therefore both longer than /17. A prefix-list entry can match a route only when the configured prefix is an aggregate of that route, so the selected /17 correctly provides a common enclosing prefix for both routes. This follows standard CIDR containment and longest-prefix concepts described in RFC 4632 and Nokia SR OS routing-policy documentation at Policy Statements.

QUESTION NO: 53

Click the exhibit button below. Given the SAP-ingress policy, which of the following statements are FALSE? (Choose three)

```
*A:srcdev3r1>config>qos# info
----------------------------------------------
#----------------------------------------------
echo "QoS Policy Configuration"
#----------------------------------------------
        sap-ingress 66 create
            queue 1 create
            exit
            queue 3 profile-mode create
                parent "high" level 5 weight 100 cir-level 5
                rate 50000 cir 10000
            exit
            queue 5 create
                parent "high" level 7 weight 100 cir-level 7
                rate 30000 cir 20000
            exit
            queue 6 create
                parent "high" level 8 weight 100 cir-level 8
                rate 2000 cir 2000
            exit
            queue 11 multipoint create
            exit
            fc "af" create
                queue 3
                profile in
            exit
            fc "ef" create
                queue 6
            exit
            fc "h2" create
                queue 5
            exit
            ip-criteria
                entry 10 create
                    match protocol tcp
                        dst-port eq 1234
                    exit
                    action fc "ef" priority high
                exit
                entry 15 create
                    match protocol udp
                    exit
                    action fc "h2"
                exit
                entry 20 create
                    match protocol icmp
                    exit
                    action fc "h2"
                exit
            exit
            default-fc "af"
        exit
----------------------------------------------
```

A. Queue 5 is using priority mode.

B. UDP traffic can be serviced at a maximum rate of 20 Mbps.

C. Ping traffic is place in queue 5.

D. Queues 5 and 6 will receive equal portions of any remaining available bandwidth to meet their above-CIR bandwidth requirements.

E. Queue 3 is scheduled before queue 5.

ANSWER: B D E

Explanation:

The false statements are “UDP traffic can be serviced at a maximum rate of 20 Mbps,” “Queues 5 and 6 will receive equal portions of any remaining available bandwidth to meet their above-CIR bandwidth requirements,” and “Queue 3 is scheduled before queue 5.” The SAP-ingress policy must be evaluated as a combination of its forwarding-class classification, queue association, queue CIR/PIR or rate settings, and scheduler configuration. A protocol match alone does not establish the stated 20-Mbps service limit; the applicable queue and its configured rate parameters determine the traffic treatment. Similarly, excess bandwidth allocation is governed by the configured scheduling behavior and weights rather than an assumption that two queues automatically split unused bandwidth equally. Queue service order is also controlled by the scheduler’s configured priority and scheduling mode, not simply by queue numbering. Nokia SR OS QoS supports queue-specific rate control and scheduler-based distribution of available bandwidth, including differentiated treatment above committed rates. These policy elements are what determine the actual ingress service outcome. See Nokia’s SR OS documentation portal and the Nokia 7750 Service Router product documentation resources for SR OS QoS and scheduling concepts.

QUESTION NO: 54

Click the exhibit button below. Given the output of the #show pools 1/2/2 network-egress command on a GigE port, what can the service provider deduce? (Choose two)

```
A:sr1a>config# show pools 1/2/2 network-egress

===============================================================
Pool Information
===============================================================
Port           : 1/2/2
Application    : Net-Egr            Pool Name          : default
Resv CBS       : Sum
---------------------------------------------------------------
-------------------------------------------------
Utilization           State    Start-Avg Max-Avg   Max-Prob
---------------------------------------------------------------
-------------------------------------------------
High-Slope            Down       70%       90%       80%
Low-Slope             Up         10%       50%       80%

Time Avg Factor     : 7
Pool Total     : 20480 KB
Pool Shared    : 12288 KB           Pool Resv          : 8192 KB

Pool Total In Use   : 704 KB
Pool Shared In Use  : 192 KB              Pool Resv In Use   :
512 KB
WA Shared In Use    : 1 KB

Hi-Slope Drop Prob  : 0                 Lo-Slope Drop Prob : 0
---------------------------------------------------------------
FC-Maps             MBS       Depth   A.CIR       A.PIR
                    CBS               O.CIR       O.PIR
---------------------------------------------------------------
be l2               192       190     0           1000000
                    0                 0           Max
af                  10240     456     250000      1000000
                    1792              250000      Max
l1                  5120      0       250000      1000000
                    512               250000      Max
h2                  10240     0       1000000     1000000
                    1792              Max         Max
ef                  10240     0       1000000     1000000
                    1792              Max         Max
h1                  512       0       100000      10000
                    512               100000      10000
nc                  5120      0       100000      1000000
                    512               100000      Max
```

A. QoS is not configured on the router.

B. Packets in forwarding classes "be" and "l2" are being queued.

C. Traffic belonging to forwarding class "ef" is not experiencing queuing delays.

D. Traffic belonging to forwarding class "l1" will never make use of the shared buffer pool.

E. Out-of-profile traffic in the shared buffer pool is being dropped.

ANSWER: B C

Explanation:

The pool statistics show nonzero buffer utilization for the *be* and *l2* forwarding classes. A nonzero in-use or allocated-buffer value means packets associated with those forwarding classes are resident in their egress queues at the instant the command is issued. Therefore, packets in forwarding classes "be" and "l2" are being queued. Queue occupancy is a direct indication that egress buffering is being used while traffic awaits transmission on the GigE port.

The *ef* forwarding class shows no queued buffer use in the displayed network-egress pool statistics. With no packets occupying the EF egress queue, EF traffic has no queue-residence time at that observation point; consequently, traffic belonging to forwarding class "ef" is not experiencing queuing delays. This conclusion concerns current queueing delay, rather than an absolute guarantee about future congestion. Nokia SR OS QoS uses forwarding classes to associate traffic with queueing and buffer-management behavior, and the pool display provides the relevant per-forwarding-class buffer-use counters. See Nokia's SR OS QoS overview and egress queuing documentation.

QUESTION NO: 55

When a Service Provider offers VPRN services to its customers, which of the following functions are expected to be the responsibility of the Service Provider? (Choose three)

A. Distributing the customer generated labels between sites

B. Distributing the customer routing information between customer sites

C. Forwarding the customer originated data packets to the appropriate destination

D. Forwarding the provider originated data packets to the appropriate customer site

E. Providing secure layer 3 routing exchange between sites

ANSWER: B C E

Explanation:

A VPRN is a provider-managed Layer 3 VPN that uses separate customer routing and forwarding contexts on provider edge routers. Accordingly, **Distributing the customer routing information between customer sites** is a core provider function: customer routes learned at one VPRN attachment are carried across the provider network and installed only in the relevant VPN routing context. The provider also performs **Forwarding the customer originated data packets to the appropriate destination**. After ingress classification into the VPRN, the provider edge encapsulates traffic for transport across the MPLS/IP backbone, and the egress provider edge delivers it toward the correct customer-facing interface. Finally, VPRN service design inherently delivers **Providing secure layer 3 routing exchange between sites**. VPN route-target policy and distinct VRF/VPRN routing tables constrain route distribution and traffic forwarding to the intended customer VPN, preserving separation between customers sharing the same provider infrastructure. These are the essential responsibilities of a service provider delivering an RFC 4364-style BGP/MPLS IP VPN service. Nokia's service documentation describes VPRN as a Layer 3 VPN service implemented on provider edge routers and carried over the provider MPLS network. See Nokia SR OS VPRN service overview and RFC 4364: BGP/MPLS IP Virtual Private Networks.

QUESTION NO: 56

Click the exhibit.

The displayed output is from an ASBR supporting an inter-AS model A VPRN. Which of the following statements about this output is TRUE?

A. The neighbor shown is an ASBR router.

B. The address family between the peers should be VPN-IPv4 instead of IPv4.

C. The neighbor address cannot be the interface address 10.3.4.4.

D. The omitted AS number under "BGP summary" must be equal to 64496.

ANSWER: A

Explanation:

The neighbor shown is an ASBR router. In an inter-AS VPRN implementation using model A, the autonomous-system boundary routers are directly connected and exchange customer VPN reachability across the AS boundary. Each ASBR maintains VPN routing and forwarding information for the VPNs it supports, and the eBGP peering between the boundary routers carries the required labeled IPv4 reachability for those VPNs. Consequently, a BGP neighbor displayed on an ASBR in this model represents the neighboring AS boundary device participating in the inter-AS VPN connection. The peer is not merely an ordinary provider-edge router hidden behind the boundary; it is the router at the adjacent AS edge that terminates the inter-AS relationship and performs the associated VPN route handling. This is the defining topology and control-plane role of inter-provider model A. Nokia's VPRN documentation describes inter-AS VPN support and the ASBR role in exchanging VPN reachability, while RFC 4364 defines the inter-provider VPN model in which AS boundary routers distribute VPN routes between autonomous systems. See Nokia SR OS Layer 3 VPN documentation and RFC 4364, BGP/MPLS IP Virtual Private Networks.

QUESTION NO: 57

Click on the exhibit below.

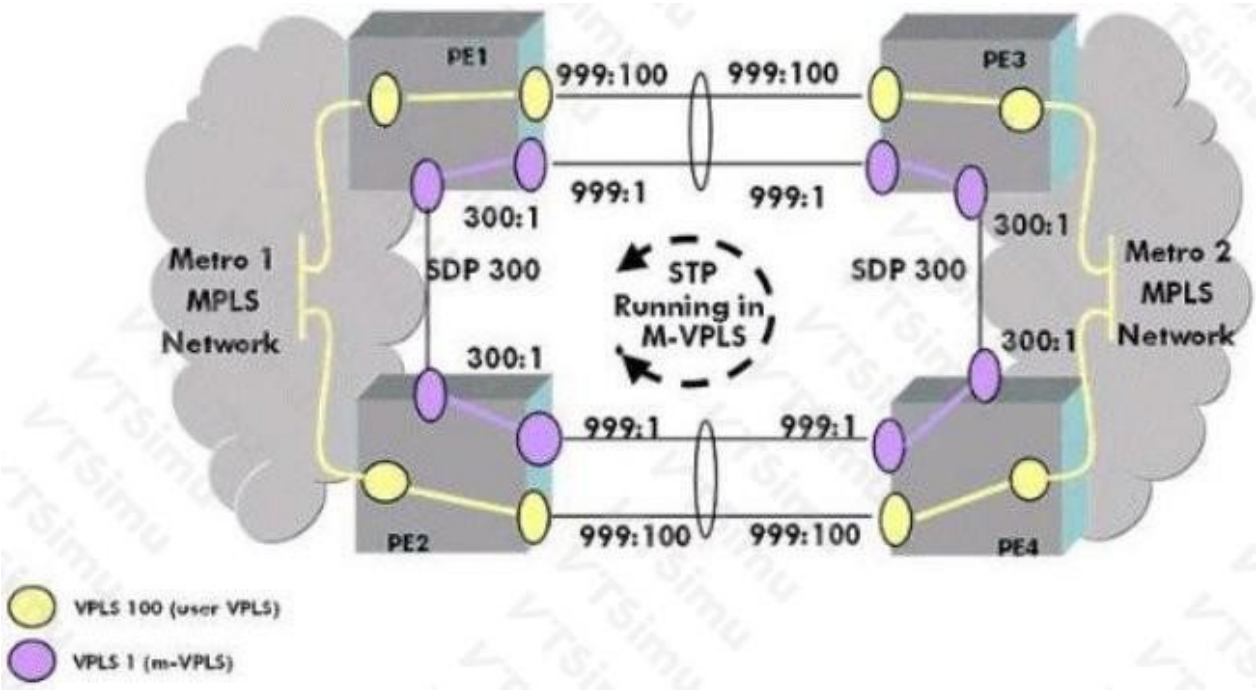

You have been requested to add load balancing across the redundant spoke-SDPs in the network shown. What further configuration is required? (Choose 2)

A. An additional management VPLS.

B. Another set of LSPs between PE1 and PE3 and between PE2 and PE4.

C. Another set of SDPs between PE1 and PE3 and between PE2 and PE4.

D. Load balancing has to be enabled within the management VPLS.

ANSWER: A C

Explanation:

Load balancing over redundant spoke-SDP paths requires the redundant topology to be represented by separate service constructs rather than simply relying on the existing protection relationship. An additional management VPLS is required so that the participating PE routers can exchange the service-management and forwarding-state information needed to coordinate the redundant spoke-SDP arrangement. This management function supports a loop-free, consistent active forwarding design while the redundant paths are used for traffic distribution.

Another set of SDPs between PE1 and PE3 and between PE2 and PE4 is also required. These SDP bindings provide the additional service transport adjacencies needed to form the second forwarding path across the redundant portions of the topology. The SDPs can use the MPLS transport infrastructure already available between the relevant PE routers; the essential requirement for this design is the additional SDP service bindings and the management VPLS coordination. Nokia's Layer 2 Services documentation describes VPLS service constructs and spoke-SDP operation, while the SR OS documentation portal provides the release-specific configuration references: Nokia SR OS documentation and Nokia 7750 SR Layer 2 Services Guide.

QUESTION NO: 58

Click the exhibit button below. Given that the slope-policy (below) has been enabled and applied on a network port, which of the following statements are TRUE? (Choose three)

```
A:sr1a>config>qos>slope-policy# info detail
----------------------------------------------
            description "Slope policy"
            high-slope
                start-avg 50
                max-avg 70
                max-prob 60
                no shutdown
            exit
            low-slope
                start-avg 10
                max-avg 20
                max-prob 80
                no shutdown
            exit
            time-average-factor 0
----------------------------------------------
```

A. All out-of-profile traffic will be dropped before any in-profile is dropped.

B. The time average factor of 0 will discard all packets arriving in the shared buffer pool.

C. When the shared buffer utilization reaches 51%, in-profile packets might be dropped.

D. The highest probability with which an out-of-profile packet can be dropped is 20%.

E. When an in-profile packet arrives and the shared buffer utilization is at 69%, the packet will have approximately a 60% likelihood of being discarded.

ANSWER: A C E

Explanation:

A slope policy applies Random Early Detection behavior according to shared-buffer utilization and packet profile. The configured out-of-profile slope begins acting at a lower buffer-utilization point than the in-profile slope and reaches full discard before the in-profile slope begins. Therefore, all out-of-profile traffic is discarded before the policy can begin discarding in-profile traffic. This preserves the preferred, in-profile traffic during congestion.

The in-profile slope starts at 50% shared-buffer utilization. Once utilization exceeds that threshold, such as at 51%, the applicable discard probability is no longer zero, so an arriving in-profile packet might be dropped. At 69% utilization, the packet is near the configured maximum-average threshold for the in-profile slope. Interpolating across the configured slope gives a discard probability of approximately 60%, as stated. The time-average factor determines how the buffer-use average is calculated; it does not alter the configured slope thresholds and probabilities used for this result. Nokia describes slope policies as congestion-management mechanisms that use buffer occupancy and profile-aware discard behavior; see the Nokia SR OS documentation and Nokia's SR OS product documentation overview.

QUESTION NO: 59

What needs to be configured in a VPLS in order to forward customer STP BPDUs transparently?

A. Spanning tree must be enabled within the VPLS.

B. Spanning tree BPDUs must be untagged.

C. Spanning tree BPDUs must be tagged.

D. Spanning tree must be disabled within the VPLS.

ANSWER: D

Explanation:

Spanning tree must be disabled within the VPLS. With the VPLS service's own spanning-tree function disabled, the provider edge does not participate in, terminate, or make forwarding decisions based on the customer's Spanning Tree Protocol control traffic. Customer Bridge Protocol Data Units are therefore handled as customer Ethernet frames and carried through the VPLS along with the customer's Layer 2 traffic. This lets customer switches at different VPLS attachment points exchange BPDUs and independently form their own STP topology, while the service remains transparent to that customer control plane.

This is an important distinction between the provider service's loop-prevention behavior and the customer's loop-prevention domain. The VPLS spanning-tree setting controls whether the VPLS itself runs STP-related processing. Disabling that service function preserves customer STP transparency, allowing the customer's STP instances to see the VPLS as a bridged Ethernet transport. Nokia's Layer 2 services documentation describes VPLS behavior and its Ethernet forwarding model; see the Nokia SR OS VPLS documentation and the Nokia SR OS Ethernet services documentation.

QUESTION NO: 60

Click the exhibit.

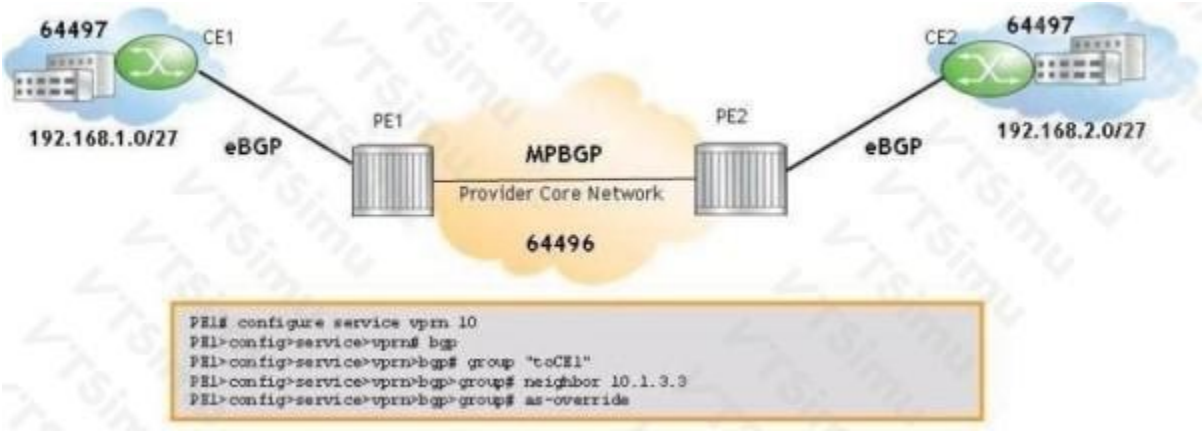

AS-override is configured on PE1 to eliminate the BGP loops. What is the expected AS-Path in the VPN-IPv4 route 64496:10:192.168.2.0/27 received by PE1?

A. 64496 64497

B. 64496

C. 64497

D. 64497 64497

ANSWER: C

Explanation:

64497 is the expected AS-Path for the VPN-IPv4 route when PE1 receives it through the provider MP-BGP VPN control plane. The route originated from the customer site in autonomous system 64497, so that customer AS is present in the path carried with the VPN-IPv4 advertisement. AS override is an egress PE-to-CE function: it is used when a PE advertises a VPN route to a CE and replaces occurrences of the CE's own AS in the advertised path with the provider PE's AS. This lets a CE accept a route learned from another site that uses the same customer AS, avoiding the normal BGP own-AS loop check. It does not alter the route's original customer AS path merely because PE1 receives the VPN-IPv4 route from the remote PE. Thus, before PE1 advertises the route onward to its attached CE and applies its configured override policy, the route received by PE1 retains AS 64497 as its AS-Path. Nokia's Layer 3 VPN documentation describes VPN route exchange and PE-to-CE BGP behavior in its Layer 3 VPN guide; related BGP path-processing concepts are covered in the BGP documentation.

QUESTION NO: 61

Click the exhibit.

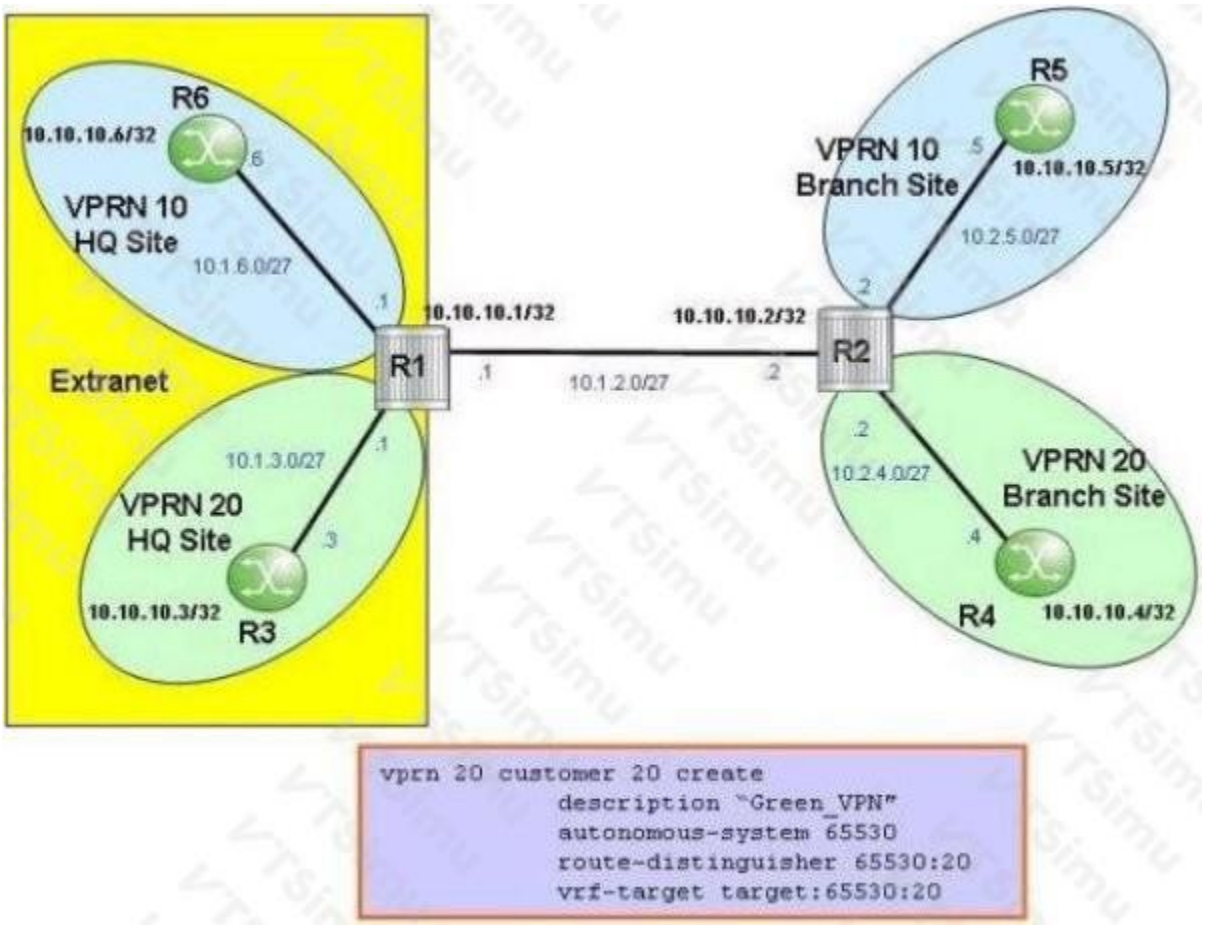

An extranet VPRN is configured on Alcatel-Lucent 7750 SRs. On which router(s) should the displayed configuration be applied?

A. Router R1

B. Router R2

C. Routers R1 and R2

D. Router R3

E. Router R4

ANSWER: B

Explanation:

Router R2 is correct because extranet VPRN configuration is applied at the provider-edge router that must perform the controlled exchange of routes between the participating customer VPRNs. The displayed policy and VPRN configuration therefore belongs on the shared service edge, where the relevant VPRN instances coexist and where import/export controls can be enforced. This router is responsible for advertising only the intended customer routes into the extranet and for importing the permitted extranet routes into the appropriate VPRN routing context. In an MPLS L3VPN/VPRN deployment, the provider-edge router associates customer routes with route targets and applies the required import and export policy before distributing VPN routes through MP-BGP. This makes the service edge, rather than a customer-facing site that has no need to participate in the inter-VPRN route exchange, the correct placement for the configuration. Nokia's VPRN documentation describes VPRN as a Layer 3 VPN service implemented on service routers and documents route-policy control for VPN route exchange. See the Nokia 7750 SR VPRN Guide and the Nokia 7750 SR BGP Guide for the VPN route-target and policy concepts used by this design.

QUESTION NO: 62

Click the exhibit.

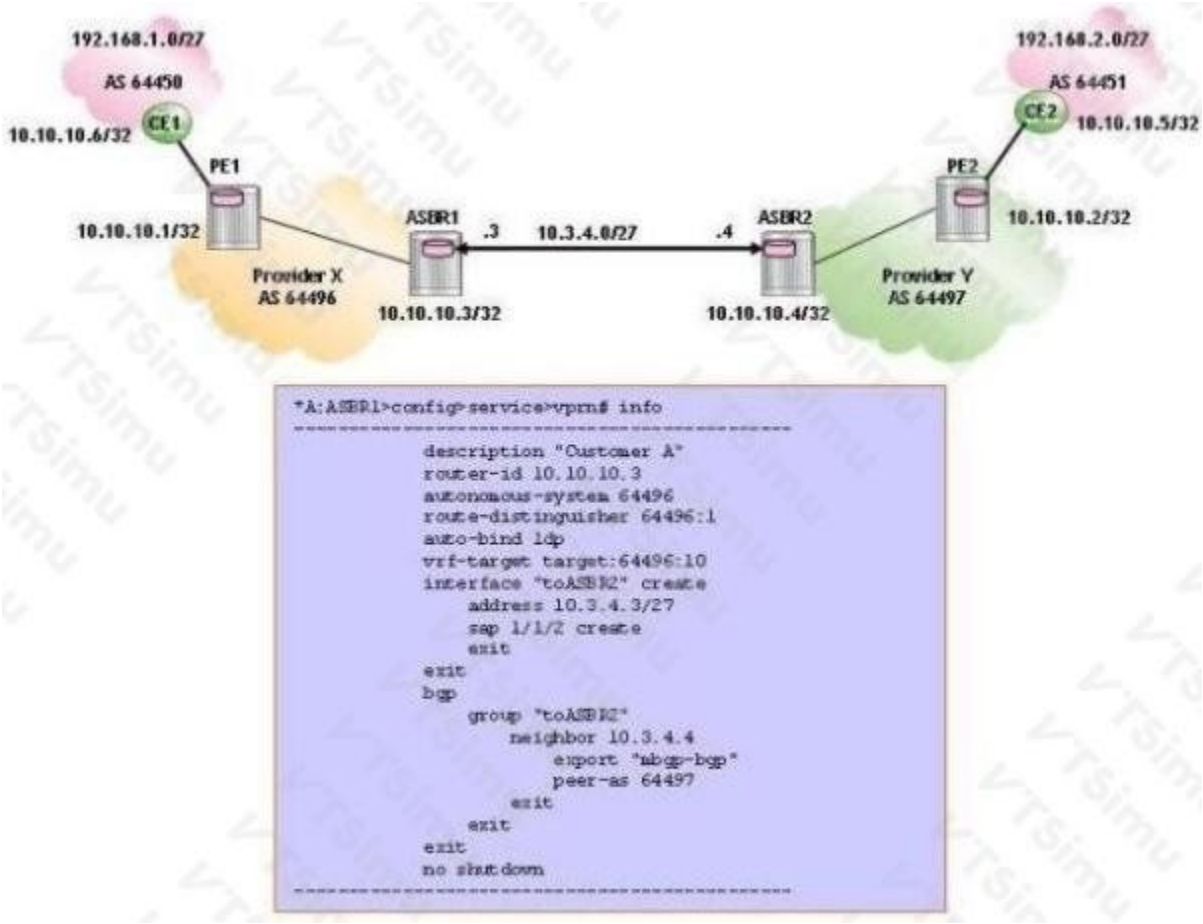

Which of the following about the inter-AS model A VPRN configuration on ASBR1 is TRUE?

A. The route-distinguisher value must be 64496:10 for successful VPRN operation.

B. An export policy is not required for successful VPRN operation.

C. The "family vpn-ipv4" command is required under the BGP configuration for successful VPRN operation.

D. The neighbor address should be 10.10.10.4 under the BGP configuration for successful VPRN operation.

E. The VPRN configuration on ASBR1 is correct.

ANSWER: E

Explanation:

The VPRN configuration on ASBR1 is correct. Inter-AS VPRN model A uses a direct customer-style routing connection between the autonomous-system border routers. Each border router places the inter-AS attachment in the relevant VPRN and exchanges ordinary IPv4 customer prefixes with its peer using eBGP in that VPRN context. This model does not require VPN-IPv4 route exchange across the AS boundary; the VPN route targets and route distinguishers remain significant within each provider's VPN service architecture, while the ASBR-to-ASBR session carries the customer IPv4 reachability needed to connect the two VPN portions.

The exhibited configuration follows that model-A design: it has the VPRN service context and a direct external BGP relationship toward the remote ASBR through the inter-AS VPRN interface. The route distinguisher is locally assigned service metadata rather than a value that must match a fixed value for end-to-end operation. Nokia's Layer 3 services documentation describes VPRN service operation and routing within a VPRN, including BGP-based customer routing, at Nokia SR OS VPRN documentation. Nokia's SR OS documentation portal, which provides the applicable routing and Layer 3 service guides by release, is available at Nokia SR OS documentation.

QUESTION NO: 63

Click the exhibit.

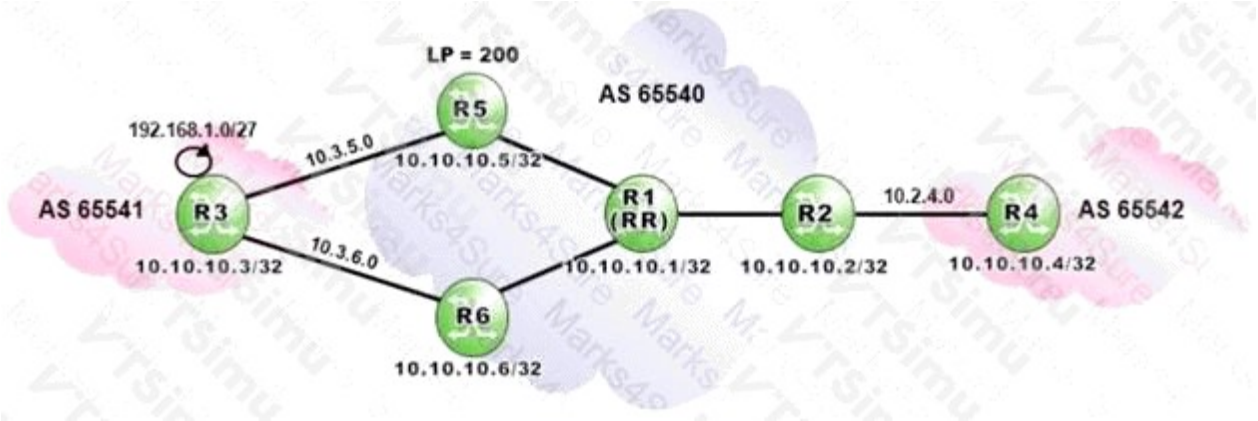

Router R1 is a route reflector with clients R2, R5 and R6. Prefixes advertised by router R5 have a local preference of 200. Router R3 advertises the prefix 192.168.1.0/27 to routers R5 and R6.

Assuming that router R6 is configured with " advertise-external " , what is the expected output of " show router bgp routes " when executed on router R5?

A.

Flag	Network	LocalPref	MED
	Nexthop	Path-Id	VPNLabel
	As-Path		
u*>i	192.168.1.0/27	200	None
	10.10.10.5	None	-
	65541		

B.

Flag	Network	LocalPref	MED
	Nexthop	Path-Id	VPNLabel
	As-Path		
u*>i	192.168.1.0/27	None	None
	10.3.5.3	None	-
	65541		
i	192.168.1.0/27	200	None
	10.10.10.5	None	-
	65541		
i	192.168.1.0/27	100	None
	10.10.10.6	None	-
	65541		

C.

Flag	Network	LocalPref	MED
	Nexthop	Path-Id	VPNLabel
	As-Path		
u*>i	192.168.1.0/27	None	None
	10.3.5.3	None	-
	65541		
i	192.168.1.0/27	200	None
	10.10.10.5	None	-
	65541		

D.

Flag	Network	LocalPref	MED
	Nexthop	Path-Id	VPNLabel
	As-Path		
u*>i	192.168.1.0/27	200	None
	10.10.10.5	None	-
	65541		
i	192.168.1.0/27	100	None
	10.10.0.6	None	-
	65541		

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: C

Explanation:

Option C is correct. The `advertise-external` setting permits BGP to advertise a best external route in addition to the overall best route when the overall best route is an internally learned path. This behavior is particularly relevant in a route-reflector design: R6 can make its externally learned route for 192.168.1.0/27 available through the reflection topology even where its selected route would otherwise prevent that external route from being advertised.

On R5, the displayed BGP entries and best-path indication must still follow normal BGP path selection. The locally configured preference of 200 on routes advertised by R5 is propagated as LOCAL_PREF within the AS and is evaluated before subsequent BGP tie-breakers. The route output represented by **Option C** correctly reflects both facts: the external path is present because of `advertise-external`, while the selected/usable route state is determined using the BGP decision process, including LOCAL_PREF and the reflected-path attributes.

This is consistent with Nokia SR OS BGP route-reflector and best-external advertisement behavior, documented in the Nokia SR OS BGP documentation, and with the route-reflection model in RFC 4456.

QUESTION NO: 64

Which of the following AS Paths will match the regular expression "65100."?

A. "65100"

B. "65100 65200"

C. "65200 65100"

D. "65100 65250 65200"

ANSWER: B

Explanation:

The AS path **"65100 65200""** is the intended match for the expression `65100.`. In SR OS AS-path regular-expression syntax, an autonomous-system number written in the expression matches that specific AS, while the period is a wildcard matching one AS-path element. Therefore, the expression identifies a path consisting of AS 65100 followed by one additional AS. The displayed path has exactly that structure: AS 65100 is followed by AS 65200. This type of matching is useful in routing policy when an operator needs to identify a particular neighboring or transit-AS relationship while allowing the second AS value to vary. AS-path matching is evaluated against the ordered AS sequence carried by BGP, preserving the order in which autonomous systems appear in the route's AS_PATH attribute. Nokia's BGP routing-policy documentation describes using AS-path expressions to select routes based on those AS sequences; the underlying AS_PATH attribute and its ordered semantics are defined by BGP. See the Nokia SR OS BGP documentation and RFC 4271.

QUESTION NO: 65

Why would a selective export policy be applied to a customer VPRN service? (Choose three)

A. To limit the number of routes advertised from a customer site

B. To limit the number of routes learned from other customer sites

C. To limit connectivity to selected customer networks

D. To perform packet filtering

ANSWER: A B C

Explanation:

A selective export policy controls the distribution of customer VPN routes rather than forwarding individual packets. "To limit the number of routes advertised from a customer site" is correct because an export policy can match routes in a VPRN routing table and determine which eligible customer prefixes are exported into the VPN route-distribution mechanism. This lets the provider avoid advertising unnecessary prefixes to the rest of the VPN.

“To limit the number of routes learned from other customer sites” is also a valid service-design objective of selective route distribution. By controlling which site routes are made available across the VPN, the service can ensure that a site receives only the routes it needs. “To limit connectivity to selected customer networks” follows directly from this behavior: if a route to a remote customer subnet is not distributed to a site, that site has no VPN control-plane route for reaching that subnet. This supports hub-and-spoke, extranet, and other partial-mesh VPRN designs. Nokia VPRN services use VPN route import/export policy controls for this type of route distribution. See Nokia’s VPRN configuration documentation and the VPN route-distribution model in RFC 4364.

QUESTION NO: 66

Which of the following rate-limiting approaches are used on the Alcatel-Lucent 7750 SR? (Choose two)

A. Shapeless policing at SAP-ingress only.

B. Shapeless policing at SAP-ingress and egress.

C. Biased round robin queuing at ingress only.

D. Soft shaping at all network interfaces.

E. Soft shaping at network ingress only.

ANSWER: A D

Explanation:

The 7750 SR QoS architecture uses **shapeless policing at SAP-ingress only** as an ingress enforcement mechanism. Packets are measured against the configured ingress service access point rate parameters, and traffic that exceeds the applicable policer profile is handled immediately according to the configured QoS behavior rather than being delayed in a shaping queue. This makes it appropriate for enforcing a service/customer traffic contract as traffic enters the router.

The platform also uses **soft shaping at all network interfaces**. Soft shaping provides rate adaptation through the network-interface QoS and scheduling hierarchy, allowing the system to manage transmission rates while retaining scheduler behavior and prioritization. This mechanism is associated with network-interface traffic management rather than being limited to network ingress. Together, these mechanisms distinguish immediate ingress policing at service access points from scheduler-based rate control on network-facing interfaces. Nokia’s SR OS QoS documentation describes the platform QoS model, including SAP QoS policies, network QoS policies, meters, queues, schedulers, and rate-control behavior. See the Nokia SR OS documentation portal and the Nokia Service Router Operating System overview.

QUESTION NO: 67

Click the exhibit button below. Given the scheduler-policy (below), which of the following can be done to guarantee that the best-effort traffic receives its configured CIR value?

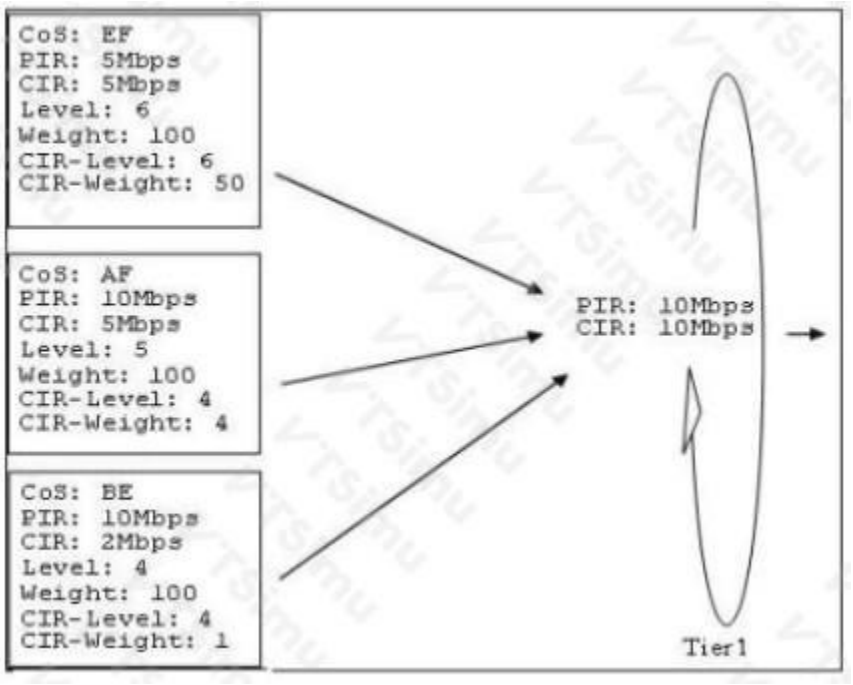

A. Change the CIR-level value for the best-effort queue to 8.

B. Do not change anything; it is already guaranteed to receive its CIR

C. Change the CIR-level value for the best-effort queue to 1.

D. Change the CIR-weight value of the AF queue to 0.

ANSWER: A D

Explanation:

Change the CIR-level value for the best-effort queue to 8 and change the CIR-weight value of the AF queue to 0 are valid ways to protect the best-effort queue's configured CIR in the scheduler policy shown. In SR OS scheduling, CIR levels and CIR weights control the order and distribution of bandwidth during CIR servicing. Moving best-effort to CIR level 8 places it in the appropriate CIR-priority treatment for this policy, preventing the competing AF configuration from taking the bandwidth allocation that best-effort requires to reach its committed rate.

Setting the AF queue's CIR weight to 0 removes AF's weighted participation at that CIR allocation point. This prevents AF from competing for the CIR bandwidth that must remain available to best-effort. Either change therefore eliminates the specific CIR-allocation contention shown in the policy and allows best-effort to receive its configured committed information rate, assuming the parent scheduler and port have sufficient available capacity to honor the configured commitments.

Nokia's QoS documentation describes scheduler-policy configuration, including CIR/PIR scheduling priorities and weights, in the SR OS Quality of Service Guide. The applicable command syntax and scheduler behavior are also documented in the SR OS scheduler-policy documentation.

QUESTION NO: 68

Which of the following statements regarding RSVP-TE LSPs are true? (Choose 3)

A. RSVP-TE can signal an explicitly routed path through the MPLS network.

B. RSVP-TE can reserve bandwidth for an LSP.

C. RSVP-TE supports fast reroute protection.

D. RSVP-TE requires all PE-CE routing protocols to use RSVP messages.

E. RSVP-TE distributes VPN routes between PE routers.

ANSWER: A B C

Explanation:

RSVP-TE can signal an explicitly routed path through the MPLS network. The ingress router can include an Explicit Route Object in an RSVP Path message to constrain the LSP to selected nodes or links rather than relying only on normal IGP shortest-path forwarding.

RSVP-TE can reserve bandwidth for an LSP. The RSVP signaling process carries traffic parameters and can cause each participating router to check available reservable bandwidth before admitting the LSP. This supports constraint-based traffic engineering.

RSVP-TE supports fast reroute protection. RSVP-TE can establish local repair paths so that a protected LSP can be rapidly redirected around a failed link or node. The point of local repair switches traffic to the backup path while a new end-to-end LSP is signaled if required.

RSVP-TE extensions for MPLS traffic engineering are specified in RFC 3209, and RSVP-TE fast reroute procedures are specified in RFC 4090.

QUESTION NO: 69

Click on the exhibit below.

CE-A sends frames with 2 tags. If the SAP on PE-A is " sap 1/1/1:1 .* " and the SAP on PE-B is " sap 1/1/1:100.200 " how many tags will there be on the frame when it egresses from the SAP on PE-B?

A. 0

B. 1

C. 2

D. 3

ANSWER: D

Explanation:

3 is correct. The ingress SAP configured as `sap 1/1/1:1.*` is a QinQ SAP: the first VLAN value identifies the outer service-delimiting tag, while the wildcarded inner VLAN accommodates the customer tag. When CE-A sends its double-tagged frame, the outer tag with value 1 is consumed as the SAP/service encapsulation tag. The existing inner customer VLAN tag remains associated with the customer frame as it is carried through the service.

At PE-B, the egress SAP is explicitly configured as `sap 1/1/1:100.200`. This QinQ egress encapsulation adds an outer VLAN tag of 100 and an inner VLAN tag of 200. These two egress tags are added in addition to the customer VLAN tag retained from the CE-A frame. Consequently, the frame leaving PE-B contains the preserved customer tag plus the two tags imposed by the PE-B SAP, for a total of three VLAN tags. This behavior illustrates how SR OS SAP encapsulation can carry customer VLAN information while adding provider-side VLAN stacking on egress. Nokia's SR OS service documentation describes SAP encapsulation and QinQ service behavior in the SR OS documentation portal; additional SR OS product context is available at Nokia Service Router Operating System.

QUESTION NO: 70

Click the exhibit.

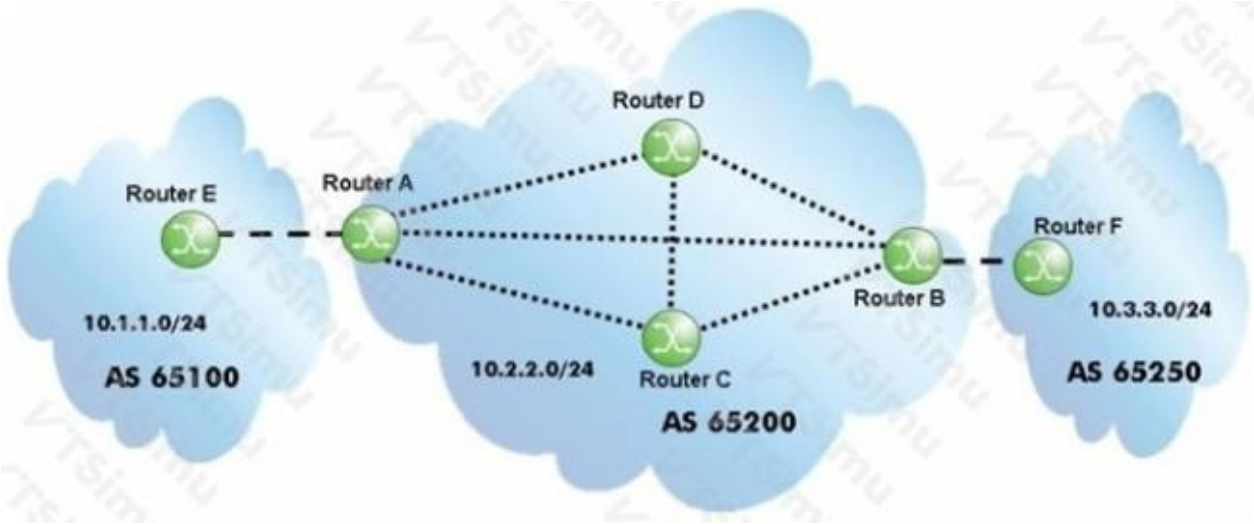

After router A receives the BGP update for the 10.3.3.0/24 prefix, which routers will the route be propagated to?

A. All routers with which it has a BGP session.

B. All routers with which it has a BGP session, except the router it received the update from, which is router B.

C. All routers with which it has a BGP session, except the router it received the update from, which is router C.

D. All routers with which it has a BGP session, except the router it received the update from, which is router D.

E. Only eBGP neighbors.

ANSWER: E

Explanation:

Only eBGP neighbors is correct because the update received by router A is learned through an iBGP relationship. BGP's iBGP split-horizon behavior prevents an iBGP-learned route from being advertised onward to another iBGP peer. This rule avoids routing loops within an autonomous system, since iBGP does not modify the AS_PATH when it propagates routes internally. Router A may, however, advertise its selected route to its external BGP peers, subject to normal BGP export policy and assuming the route is eligible for advertisement. Therefore, the applicable propagation set for this received prefix consists of router A's eBGP neighbors rather than its iBGP peers. In a full production design, routes learned through iBGP are normally distributed to all required internal routers using a full iBGP mesh, route reflectors, or confederations; route reflectors apply additional rules to safely relax the ordinary iBGP advertisement restriction. The base BGP behavior relevant here is defined in RFC 4271, which states that routes learned from an internal peer must not be advertised to another internal peer. See RFC 4271, Section 5.1.2 and Nokia's SR OS BGP documentation.

QUESTION NO: 71

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:

An SLA is an agreement between a customer and a provider that dictates the treatment of customer traffic across the provider's network. A service-level agreement establishes the measurable service commitments associated with a subscribed service, such as the traffic handling and performance objectives that the provider undertakes to deliver. In an IP/MPLS service context, these commitments commonly align with defined classes of service and the QoS behavior applied to the customer's traffic.

An SLA allows a customer to pre-mark traffic and ensure that traffic is treated as per the agreement within the provider's network. More precisely, the SLA can define the conditions under which customer markings are accepted, mapped, remarked, or otherwise used for classification, together with the corresponding forwarding treatment. This gives the customer and provider a shared, contractual basis for associating marked traffic with an agreed QoS class. The IETF's differentiated-services architecture describes how traffic classification and marking support differentiated forwarding behavior, while MEF service standards describe the use of performance objectives and service attributes in carrier services. See IETF RFC 2475 and MEF Service Standards.

QUESTION NO: 72

Which of the following statements regarding MPLS label processing in a VPRN are true? (Choose 3)

A. The ingress PE imposes an outer transport label and an inner VPN label.

B. A P router normally swaps the outer transport label.

C. The VPN label identifies the destination VPRN forwarding context at the egress PE.

D. A P router must maintain a separate IP routing table for every VPRN.

E. The CE must advertise MPLS labels directly to the remote PE.

ANSWER: A B C

Explanation:

At the ingress PE, **an outer transport label and an inner VPN label are imposed** for a packet sent to a remote VPRN site. The outer label identifies the LSP through the provider core, while the inner VPN label identifies the service forwarding context at the egress PE.

Within the provider core, **a P router normally swaps the outer transport label** and forwards the packet toward the next hop. The P router does not need to inspect the customer IP header or maintain per-VPRN routing information. At the egress PE, **the VPN label identifies the destination VPRN forwarding context**, allowing the PE to select the appropriate customer-facing interface and route the packet toward the CE. This separation between transport forwarding and VPN forwarding is a key scalability property of BGP/MPLS IP VPNs. See RFC 4364.

QUESTION NO: 73

MC-LAG provides link and node protection from an MDU to 3 or more PE devices.

A. TRUE

B. FALSE

ANSWER: B

Explanation:

FALSE is correct because Multi-Chassis Link Aggregation Group (MC-LAG) is a dual-homing resiliency mechanism: an access device such as an MDU can form a LAG toward a pair of provider-edge routers that operate together as an MC-LAG system. This arrangement protects the access connection against loss of an individual member link and against failure of either one of the two participating PE nodes. Both PE routers coordinate forwarding and LAG state through their interchassis control mechanisms, allowing the MDU to retain connectivity through the surviving PE when a link or one PE fails.

The defining deployment model is therefore two PE devices, not three or more. An MDU may have multiple uplinks in a broader network design, but those links are not a single MC-LAG association spanning three or more PEs. Nokia's SR OS documentation describes MC-LAG as a solution involving a pair of chassis that present a resilient LAG attachment to an access device. This dual-chassis architecture is what supplies both link and node redundancy while maintaining an Ethernet LAG operational model at the MDU. See Nokia's MC-LAG documentation and the SR OS Layer 2 Services Guide.

QUESTION NO: 74

The operation of a VPLS can be best described by which of the following statements?

A. A frame should be sent to the PE that connects to the target site whenever possible.

B. A frame should be flooded as little as possible.

C. Customer frames should be switched based on the destination MAC address.

D. All of the above.

E. Only (a) and (b) are true.

ANSWER: D

Explanation:

All of the above is correct because a VPLS emulates a multipoint Ethernet LAN across an MPLS provider network. Provider edge routers perform MAC learning for customer source MAC addresses and use the learned forwarding information to make unicast forwarding decisions. When the destination customer MAC address is known, the frame is sent over the pseudowire or spoke-SDP toward the provider edge router serving that destination site, rather than being replicated to every VPLS attachment point. This delivers the expected LAN behavior while using provider-network resources efficiently.

Flooding remains necessary for traffic whose destination cannot yet be determined from the forwarding database, as well as for broadcast and applicable multicast traffic. VPLS therefore limits flooding to the service's required forwarding scope and avoids it whenever a learned destination MAC entry permits directed forwarding. In short, the service switches customer Ethernet frames according to destination MAC addresses, forwards known unicast frames toward the appropriate remote site, and minimizes flooding through MAC learning. These behaviors are the fundamental forwarding model described for VPLS in RFC 4762 and in Nokia's Layer 2 service documentation. See RFC 4762: Virtual Private LAN Service and Nokia VPLS overview documentation.

QUESTION NO: 75

Click the exhibit button below. Given the output of the #show pools 1/2/2 network-egress command on a GigE port, what can the service provider deduce? (Choose two)

```
A:sr1a>config# show pools 1/2/2 network-egress

===============================================================
Pool Information
===============================================================
Port           : 1/2/2
Application    : Net-Egr            Pool Name          : default
Resv CBS       : Sum
---------------------------------------------------------------
-------------------------------------------------
Utilization           State    Start-Avg Max-Avg   Max-Prob
---------------------------------------------------------------
-------------------------------------------------
High-Slope            Down       70%       90%       80%
Low-Slope             Up         10%       50%       80%

Time Avg Factor     : 7
Pool Total     : 20480 KB
Pool Shared    : 12288 KB           Pool Resv          : 8192 KB

Pool Total In Use   : 704 KB
Pool Shared In Use  : 192 KB              Pool Resv In Use   :
512 KB
WA Shared In Use    : 1 KB

Hi-Slope Drop Prob  : 0                  Lo-Slope Drop Prob : 0
---------------------------------------------------------------
FC-Maps          MBS       Depth   A.CIR       A.PIR
                 CBS               O.CIR       O.PIR
---------------------------------------------------------------
be l2            192       190     0           1000000
                 0                 0           Max
af               10240     456     250000      1000000
                 1792              250000      Max
l1               5120      0       250000      1000000
                 512               250000      Max
h2               10240     0       1000000     1000000
                 1792              Max         Max
ef               10240     0       1000000     1000000
                 1792              Max         Max
h1               512       0       100000      10000
                 512               100000      10000
nc               5120      0       100000      1000000
                 512               100000      Max
```

A. QoS is not configured on the router.

B. Packets in forwarding classes "be" and "l2" are being queued.

C. Traffic belonging to forwarding class "ef" is not experiencing queuing delays.

D. Traffic belonging to forwarding class "l1" will never make use of the shared buffer pool.

E. Out-of-profile traffic in the shared buffer pool is being dropped.

ANSWER: B C

Explanation:

The correct deductions are that packets in the `be` and `l2` forwarding classes are being queued, and that `ef` traffic is not experiencing queuing delay at the time represented by the display. The network-egress pool output reports buffer and queue use for the forwarding classes associated with the port. Nonzero queued-buffer or queued-packet values for `be` and `l2`

demonstrate that traffic in those classes is resident in egress queues, awaiting transmission. This is direct evidence of queueing on the port rather than merely of QoS policy configuration.

The `ef` forwarding class has no queued occupancy in the shown pool statistics. Consequently, its packets are being serviced without waiting in an egress queue during the observed interval. This is consistent with expedited-forwarding treatment, where delay-sensitive traffic is normally assigned a high-priority service treatment. Queue occupancy is dynamic, so this conclusion describes the displayed measurement, not a permanent guarantee about all future `ef` traffic. Nokia's QoS documentation describes egress buffer pools and the relationship between forwarding-class queues, buffer allocation, and congestion behavior. See the Nokia SR OS QoS overview and the Nokia SR OS QoS pools documentation.

QUESTION NO: 76

Click the exhibit.

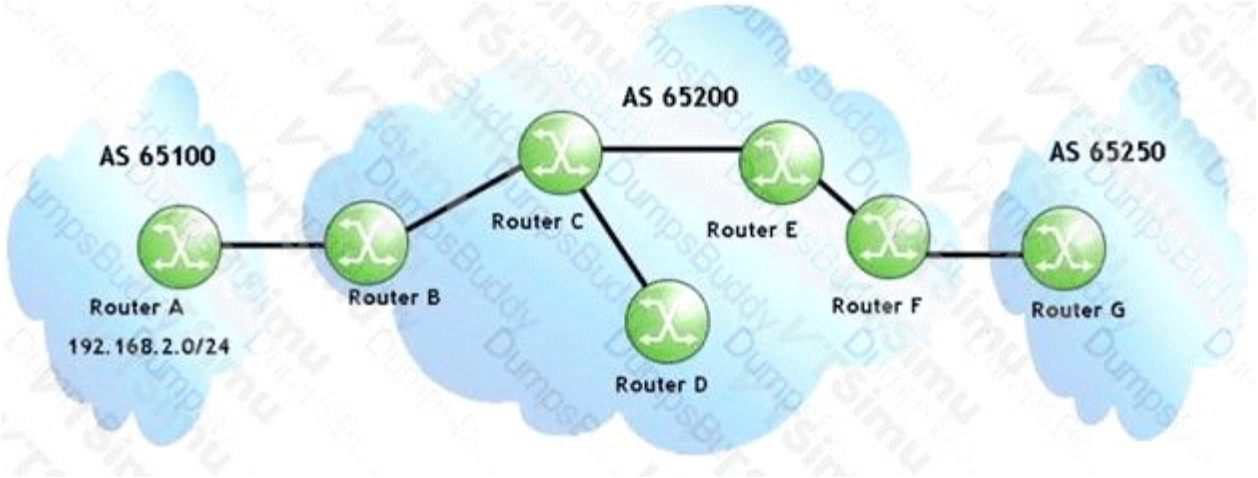

Routers C and E are route reflectors. When the update for prefix 192.168.2.0/24 is received at router F, what will be the Originator_ID attribute?

A. The router ID of router A.

B. The router ID of router B.

C. The router ID of router C.

D. It will not be set.

ANSWER: B

Explanation:

The router ID of router B is correct. In a route-reflector topology, the Originator_ID is a non-transitive BGP attribute used to preserve the identity of the BGP speaker that originally advertised a route into the reflected-client environment. When router C reflects the route originated by router B, it adds an Originator_ID equal to router B's BGP router ID if the attribute is not already present. The attribute remains associated with that route as it is subsequently propagated through additional route reflectors, including router E, and is therefore still present when router F receives the update. This provides route-reflection loop prevention: should router B later receive its own route back through a reflected path, it can identify the route as one it originated and reject it. The Originator_ID is specifically intended to represent the original route originator, rather than the route reflector currently advertising the update or the immediately preceding BGP neighbor. This behavior is defined by the BGP route-reflection mechanism in RFC 4456, which specifies that a route reflector creates the Originator_ID using the router ID of the route's originator. Nokia SR OS BGP and route-reflection documentation is available from the Nokia SR OS documentation portal.

QUESTION NO: 77

VPN-IPv4 addressing is made unique by defining an address structure consisting of:

A. A 64-bit Route Target followed by a 32-bit IPv4 address.

B. A 64-bit Route Distinguisher followed by a 32-bit IPv4 address.

C. A 32-bit Route Distinguisher followed by a 32-bit IPv4 address.

D. A 32-bit Route Target followed by a 32-bit IPv4 address.

ANSWER: B

Explanation:

A 64-bit Route Distinguisher followed by a 32-bit IPv4 address is correct. In BGP/MPLS IP VPNs, customer IPv4 address spaces can overlap because each customer VPN is logically isolated. To ensure that otherwise identical customer prefixes can be carried and selected uniquely in Multiprotocol BGP, the provider constructs a VPN-IPv4 address by prepending an 8-octet (64-bit) Route Distinguisher to the 4-octet (32-bit) IPv4 prefix. This produces a 12-octet address family value and allows the same IPv4 prefix, such as 10.0.0.0/8, to be advertised separately for different VPNs. The Route Distinguisher identifies the distinct VPN route space; it is not itself a policy mechanism for importing or exporting routes. Route distribution policy is instead controlled through BGP extended communities known as Route Targets. The VPN-IPv4 format and its use in Multiprotocol BGP are defined by RFC 4364, while the Route Distinguisher encoding is specified in RFC 4364 and RFC 4360-related BGP extended-community context. See RFC 4364, BGP/MPLS IP Virtual Private Networks and Nokia SR OS VPRN overview.

QUESTION NO: 78

Click on the exhibit below.

Given the following output what is the most likely reason the SDP binding is operationally down? (Choose 2)

A. The sdp id is different on the far-end.

B. LDP is disabled on the far-end.

C. The vc-id is different on the far-end.

D. The mesh-sdp is missing on the far-end.

E. The mesh-sdp is administratively shutdown on PEL

F. There is an MTU mismatch.

ANSWER: C D

Explanation:

An SDP binding requires a matching service attachment at both ends of the pseudowire. **The vc-id is different on the far-end.** is a valid cause because the VC ID is the pseudowire identifier exchanged between the two provider-edge routers. Both sides must use the same VC ID for the remote endpoint to associate the received pseudowire signaling with the intended service binding. If the identifiers differ, no matching far-end binding is found and the binding remains operationally down.

The mesh-sdp is missing on the far-end. is also a valid cause. A mesh SDP binding is a point-to-point service connection between VPLS peers, and an equivalent mesh-SDP service binding must exist on the remote router. Without that far-end service binding, there is no remote pseudowire endpoint with which to establish the service connection. The local router therefore cannot bring the binding to an operationally up state, even if the underlying transport SDP is otherwise available. Nokia's VPLS documentation describes mesh SDPs as the mechanism used for inter-PE VPLS connectivity and documents the requirement for consistent service and SDP-binding configuration: Nokia SR OS VPLS documentation. See also the Nokia SR OS Layer 2 services overview.

QUESTION NO: 79

What are the advantages of running a physical ring topology with a mesh of mesh-SDPs? (Choose 2)

A. STP can be used to break the loop.

B. STP is not required to break the loop because of split-horizon rules.

C. If the traffic is primarily multicast, a full mesh of SDPs is a simple and efficient design.

D. If the traffic is primarily Unicast, a full mesh of SDPs is a simple and efficient design.

ANSWER: B D

Explanation:

“STP is not required to break the loop because of split-horizon rules.” is correct because mesh service distribution points implement split-horizon behavior within a VPLS service. Traffic received through one mesh service distribution point is not forwarded through another mesh service distribution point in that same service. This forwarding rule prevents data-plane loops even when the underlying physical transport has a ring arrangement, so a spanning-tree protocol is not needed as the loop-prevention mechanism for the mesh service distribution point topology.

“If the traffic is primarily Unicast, a full mesh of SDPs is a simple and efficient design.” is also correct. A full mesh provides direct logical connectivity between provider-edge routers. For predominantly unicast traffic, the ingress router can send traffic directly toward the egress router over the applicable service distribution point, avoiding unnecessary transit through intermediate provider-edge routers. This makes forwarding straightforward and reduces the risk of inefficient, multi-hop service forwarding. Nokia’s Layer 2 service documentation describes the split-horizon and forwarding behavior used by mesh service distribution points; see the Nokia SR OS documentation portal and Nokia’s SR OS product documentation resources.

QUESTION NO: 80

Which protocol is used to exchange customer VPRN routes between PE devices?

A. OSPF

B. ISIS

C. MP-BGP

D. BGP

E. Targeted LDP

ANSWER: C

Explanation:

MP-BGP is used to exchange customer VPRN routes between provider-edge devices. In Nokia SR OS VPRN deployments, each PE maintains VRF-style routing information for its attached customer sites. Routes learned from a customer-facing interface or customer routing protocol are converted into VPN-IP routes and advertised to other PEs through Multiprotocol BGP. MP-BGP carries the VPN route family, including the route distinguisher that makes overlapping customer address spaces unique and the route-target extended communities used to control import and export policy between VPRN instances. The receiving PE uses these attributes to identify the appropriate VPRN routing table and installs the route with the transport tunnel information needed to reach the advertising PE. This separation allows the provider backbone to distribute routes for many independent customer VPNs while retaining isolation between them. Nokia’s VPRN documentation describes BGP VPN route distribution and route-target-based import/export mechanisms as core VPRN functions. See the Nokia SR OS VPRN documentation and the IETF’s RFC 4364: BGP/MPLS IP Virtual Private Networks.

QUESTION NO: 81

What is the function of the Cluster_List attribute?

A. It is used for loop detection within the same cluster.

B. It is used for loop detection between clusters.

C. It is used for loop detection in a confederation.

D. It is used to identify the route reflectors within an AS.

ANSWER: B

Explanation:

It is used for loop detection between clusters. In BGP route reflection, a cluster consists of one or more route reflectors that share a common cluster ID. When a route reflector advertises a reflected route, it adds its cluster ID to the non-transitive `CLUSTER_LIST` path attribute. As the route is reflected through additional route-reflector clusters, each reflector prepends its own cluster ID, creating a record of the clusters the route has traversed.

Before accepting a reflected update, a route reflector examines the `CLUSTER_LIST`. If its local cluster ID is already present, the update has returned to a cluster it previously crossed and is discarded. This prevents a reflected BGP route from circulating indefinitely among route-reflector clusters while retaining the normal AS path unchanged inside an autonomous system. The mechanism is particularly important in designs with redundant route reflectors and multiple interconnected clusters, where several reflection paths can otherwise exist. RFC 4456 defines this behavior and specifies that a route reflector must ignore an update containing its local cluster ID in the `CLUSTER_LIST`: RFC 4456: BGP Route Reflection. Nokia's BGP route-reflector documentation also describes cluster IDs and cluster-list-based loop prevention: Nokia SR OS BGP documentation.

QUESTION NO: 82

Which of the following statements are TRUE regarding the building of a QoS-enabled network? (Choose two)

A. It is important that all packets are classified in the same way at all SAP-ingress points.

B. Once a packet is classified to a forwarding class, it cannot change across the network.

C. Packet enters a node will be treated as per the local hop's configured policy.

D. Native IP packets entering a network port will be marked at the first network egress by default

E. Native IP packets entering a network port will not be marked at the first network egress by default.

F. Native IP packets entering a network port will be marked at the first network egress, regardless of the re-marking configuration.

ANSWER: C E

Explanation:

"Packet enters a node will be treated as per the local hop's configured policy." is correct because SR OS QoS is implemented on a per-node, per-interface basis. At each hop, the device applies its locally configured classification, policing or marking behavior, forwarding-class handling, queue mapping, scheduling, and congestion-management policy. This hop-by-hop model allows an operator to enforce the QoS design at every ingress and egress point while using forwarding classes and packet markings to carry service intent through the network. "Native IP packets entering a network port will not be marked at the first network egress by default." is also correct. Native IP traffic received at a network ingress can be classified into an internal forwarding class, but that classification does not by itself cause the original IP header DSCP value to be rewritten automatically at the first network egress. IP remarking is an explicit QoS-policy action and is used when the operator intends to impose or restore a particular DSCP marking domain. This default preserves the received native IP marking unless a configured marking policy directs otherwise. Nokia's SR OS QoS documentation describes the per-hop QoS functions and policy-controlled packet marking behavior; the DiffServ architecture also defines DSCP processing as a per-hop behavior.

References: Nokia SR OS Documentation; RFC 2474 — Differentiated Services Field.

QUESTION NO: 83

Which of the following statements regarding LDP discovery and session establishment are true? (Choose 3)

A. LDP Hello messages use UDP port 646.

B. An LDP session uses TCP port 646.

C. Targeted LDP can establish an LDP session between non-directly connected routers.

D. Link LDP Hello messages use TCP to provide reliable discovery.

E. LDP label mappings are exchanged only through UDP Hello messages.

ANSWER: A B C

Explanation:

LDP Hello messages use UDP port 646. LDP peers use Hello messages for discovery and to communicate transport-address information. These link Hello messages are sent using UDP port 646 to the all-routers multicast address on the applicable interface.

An LDP session uses TCP port 646. After discovery, the LDP peers establish a reliable TCP connection for LDP session messages, including Initialization, KeepAlive, and Label Mapping messages. TCP provides ordered and reliable delivery for label-distribution information.

Targeted LDP can establish an LDP session between non-directly connected routers. Targeted Hello messages are unicast toward a specified peer and are commonly used for applications such as pseudowire signaling between PE routers. Unlike basic link discovery, the peers do not need to share a directly connected subnet.

LDP discovery and session procedures are specified in RFC 5036. Nokia SR OS MPLS and LDP configuration information is available through the Nokia SR OS documentation portal.

QUESTION NO: 84

Click the exhibit button below.

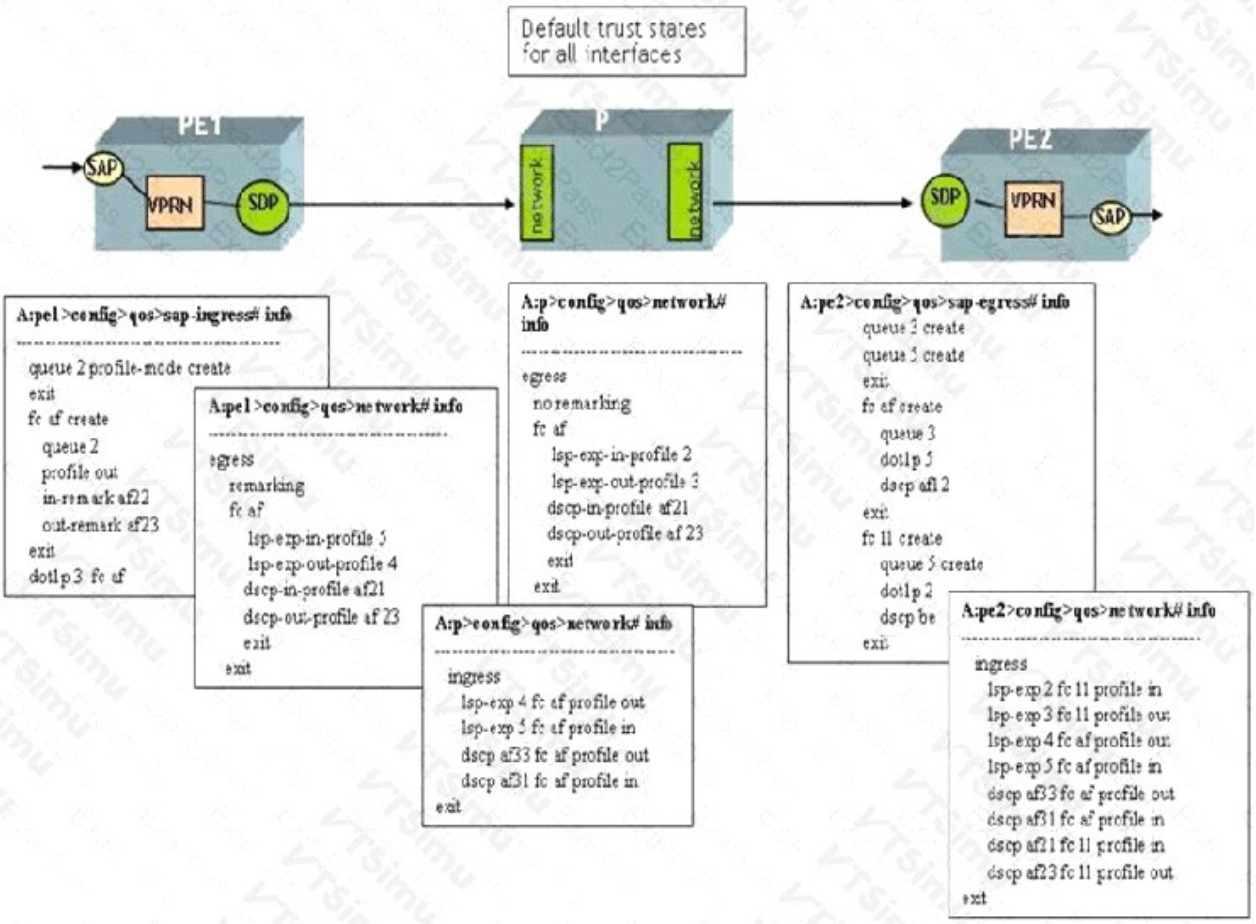

All interfaces are using their default trust states and MPLS is used as the transport tunnel. The SAP-ingress, SAP-egress, and network QoS policies have been configured as shown below. Assume that the default network-queue policy is used on each router.

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. The DSCP value will be set to af23.

D. The DSCP value will be set to af12.

E. The dot1p value will be set to 5.

F. The dot1p value will be set to 2.

G. The dot1p value will be set to 3.

ANSWER: D E

Explanation:

The DSCP value will be set to af12 and The dot1p value will be set to 5 are correct. QoS processing is performed independently at service ingress, in the MPLS network, and at service egress. At PE 1, the SAP-ingress policy classifies the received customer frame using the applicable trusted marking and maps it to the forwarding class and profile defined by the policy. MPLS transport then carries the packet across the provider network using the network QoS treatment appropriate to that forwarding class; it does not make the customer-facing Ethernet and IP remarking at the far-end SAP.

At PE 2, the SAP-egress policy supplies the customer-facing remarking values associated with the packet's resolved forwarding class and profile. The configured IP remarking entry produces the AF12 DSCP value, while the configured IEEE 802.1p remarking entry produces priority value 5. Thus, the packet sent from PE 2 toward the customer has DSCP AF12 in its IP header and dot1p priority 5 in its Ethernet VLAN tag. This illustrates the separation between internal MPLS QoS handling and egress customer marking described in Nokia's QoS documentation. See the Nokia SR OS QoS overview and the Nokia SR OS QoS policy documentation.

QUESTION NO: 85

Click the exhibit button below. Given the output of the #show pools 1/2/2 network-egress command on a GigE port, what can the service provider deduce? (Choose two)

```
A:sr1a>config# show pools 1/2/2 network-egress

===============================================================
Pool Information
===============================================================
Port           : 1/2/2
Application    : Net-Egr            Pool Name          : default
Resv CBS       : Sum
---------------------------------------------------------------
-----------------------------------------------
Utilization           State    Start-Avg Max-Avg   Max-Prob
---------------------------------------------------------------
-----------------------------------------------
High-Slope            Down       70%       90%       80%
Low-Slope             Up         10%       50%       80%

Time Avg Factor     : 7
Pool Total     : 20480 KB
Pool Shared    : 12288 KB           Pool Resv          : 8192 KB

Pool Total In Use   : 704 KB
Pool Shared In Use  : 192 KB              Pool Resv In Use   :
512 KB
WA Shared In Use    : 1 KB

Hi-Slope Drop Prob  : 0                  Lo-Slope Drop Prob : 0
---------------------------------------------------------------
FC-Maps          MBS        Depth   A.CIR        A.PIR
                 CBS                O.CIR        O.PIR
---------------------------------------------------------------
be l2            192        190     0            1000000
                 0                  0            Max
af               10240      456     250000       1000000
                 1792               250000       Max
l1               5120       0       250000       1000000
                 512                250000       Max
h2               10240      0       1000000      1000000
                 1792               Max          Max
ef               10240      0       1000000      1000000
                 1792               Max          Max
h1               512        0       100000       10000
                 512                100000       10000
nc               5120       0       100000       1000000
                 512                100000       Max
```

A. QoS is not configured on the router.

B. Packets in forwarding classes "be" and "12" are being queued.

C. Traffic belonging to forwarding class "ef" is not experiencing queuing delays.

D. Traffic belonging to forwarding class "11 "will never make use of the shared buffer pool.

E. Out-of-profile traffic in the shared buffer pool is being dropped.

ANSWER: B C

Explanation:

The displayed network-egress pool statistics show that the queues associated with forwarding classes *be* and *l2* contain queued packets. A nonzero packet count in a queue is direct evidence that packets have entered that egress queue and are awaiting transmission, which is the relevant indication of queuing for these forwarding classes. Nokia SR OS uses egress queues to hold traffic when transmission scheduling and available egress bandwidth require packets to wait before they can be sent.

The forwarding class *ef* is not experiencing queuing delay in the shown output because its queue has no packets waiting in it. Therefore, packets classified into *ef* are being transmitted without an observed backlog at the time the counters were collected. This conclusion is based on the queue occupancy shown by the command, rather than on an assumption that the class can never queue. Queue and buffer-pool statistics are point-in-time operational measurements and can change as offered load changes.

Nokia documents egress queueing, buffer pools, and forwarding-class handling in the SR OS QoS guide: Nokia SR OS Quality of Service Guide. The SR OS documentation library is available at Nokia SR OS Documentation.

QUESTION NO: 86

How does an IOM3 improve the queue usage over an IOM2 in an Alcatel-Lucent 7750 SR? (Choose two)

A. The IOM3 adds more memory to both forwarding complexes which increases the total number of queues available.

B. The IOM3 only contains a single forwarding complex which reduces the number of queues used.

C. The IOM3 allows multiple forwarding classes to be assigned to a single queue.

D. The IOM3 allows queues to be dynamically allocated between ingress and egress.

E. The IOM3 has a larger CPU capable of tracking active hardware queue utilization and adjusting memory allocation accordingly.

ANSWER: B D

Explanation:

The IOM3 only contains a single forwarding complex which reduces the number of queues used. This architecture avoids the queue-resource inefficiency that results when queue capacity must be considered separately for multiple forwarding complexes. With one forwarding complex, queue resources are managed from a unified hardware context, allowing deployed services and ports to make more effective use of the available queue capacity.

The IOM3 allows queues to be dynamically allocated between ingress and egress. This is the key queue-efficiency improvement: instead of relying on a fixed division of queue resources by traffic direction, the IOM3 can assign available queue capacity according to actual ingress and egress requirements. Consequently, capacity that is not required in one direction can be available for demand in the other direction, improving utilization and reducing the likelihood that one resource partition is exhausted while capacity remains unused elsewhere. This behavior is especially useful where subscriber, service, and port traffic patterns are asymmetric or change over time. Nokia's SR OS quality-of-service documentation describes the platform queueing model and hardware-specific queue-resource behavior; the hardware documentation provides the corresponding IOM architecture context. See the Nokia SR OS documentation portal and the Nokia 7750 SR documentation archive.

QUESTION NO: 87

Click on the exhibit below.

Which ports will become root ports? (Choose 3)

A. Bridge B 1/1/3

B. Bridge B 1/1/2

C. Bridge C 1/1/2

D. Bridge C 2/1/1

E. Bridge D 1/1/1

F. Bridge D 1/1/3

G. Bridge E 1/1/1

ANSWER: A D F

Explanation:

In an RSTP/MSTP topology, every non-root bridge selects exactly one root port: the port that offers the lowest-cost active path toward the elected root bridge. Based on the bridge IDs, port costs, and inter-switch links shown in the referenced exhibit, **Bridge B 1/1/3**, **Bridge C 2/1/1**, and **Bridge D 1/1/3** are the interfaces receiving the superior BPDUs for their respective non-root bridges. Each therefore supplies its bridge's best path to the root and transitions to the root-port role after spanning-tree convergence.

Root-port selection first compares the cumulative root-path cost advertised in received BPDUs. Where paths have the same cost, the standard BPDU comparison process uses the upstream bridge ID, upstream port ID, and then the local receiving port ID as tie-breakers. The root bridge itself has no root port, because it is the reference point for the spanning-tree instance. Nokia SR OS implements these standard spanning-tree port-role concepts for loop-free Layer 2 forwarding. See the Nokia SR OS Layer 2 Services documentation and the IEEE 802.1D standard overview.