PassQueen
a.k.a dumpsqueen.com

Nokia NRS II Composite Exam

Nokia 4A0-C01

Version Demo

Total Demo Questions: 104

Total Premium Questions: 1040

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, OSPF Routing Protocol	214
Topic 2, IS-IS Routing Protocol	119
Topic 3, BGP Routing Protocol	1
Topic 4, Multicast Routing	3
Topic 5, MPLS	317
Topic 6, MPLS VPNs	236
Topic 7, Quality of Service	7
Topic 8, High Availability	3
Topic 9, Mix Questions	140
Total	1040

QUESTION NO: 1

Choose the Distance Vector Protocols from the list below:

A. OSPF

B. RIPv1

C. RIPv2

D. IS-IS

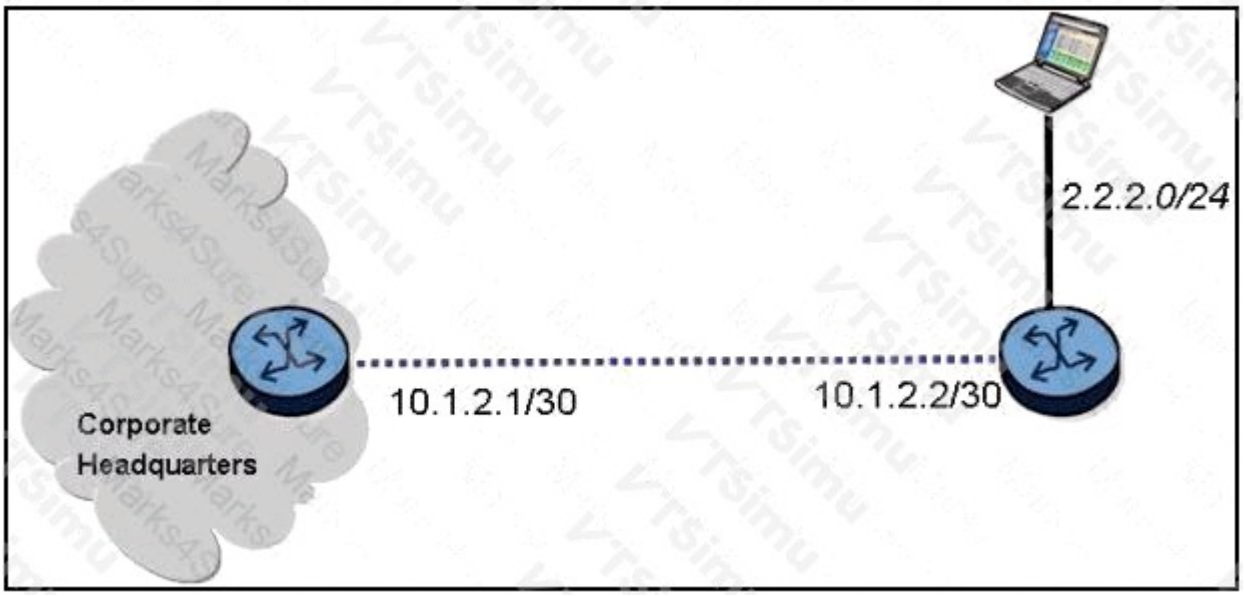
ANSWER: B C

Explanation:

RIPv1 and RIPv2 are distance-vector interior gateway protocols. Each router uses Routing Information Protocol metrics based on hop count and shares routing information with directly connected neighbors. The receiving router uses the advertised distance together with the next-hop information to select or update routes, which is the defining operational model of a distance-vector protocol. RIP's metric has a maximum usable value of 15 hops; a metric of 16 represents an unreachable destination. RIPv1 is specified in RFC 1058 and uses the original RIP route format. RIPv2 retains RIP's distance-vector algorithm while adding capabilities such as subnet-mask carriage, next-hop information, route tags, and multicast route updates, as specified in RFC 2453. These enhancements make RIPv2 suitable for classless addressing environments, but they do not change its fundamental distance-vector nature. Therefore, both RIPv1 and RIPv2 belong in the selected set. The protocol behavior, metric definition, and route-update processing for RIPv1 are documented in RFC 1058; the version 2 extensions and continued use of the distance-vector algorithm are documented in RFC 2453.

QUESTION NO: 2

Click the exhibit button.



Given the diagram in the exhibit, what would be the correct static route configuration on the router on the right, if it is an Alcatel-Lucent 7750 SR?

A. config > static-route 0.0.0.0/0 next-hop 10.1.2.2

B. config > router > static-route 0.0.0.0 255.255.255.255 next-hop 10.1.2.1

C. config > router > static-route 0.0.0.0/0 next-hop 10.1.2.2

D. config > router > static-route 0.0.0.0/0 next-hop 10.1.2.1

ANSWER: D

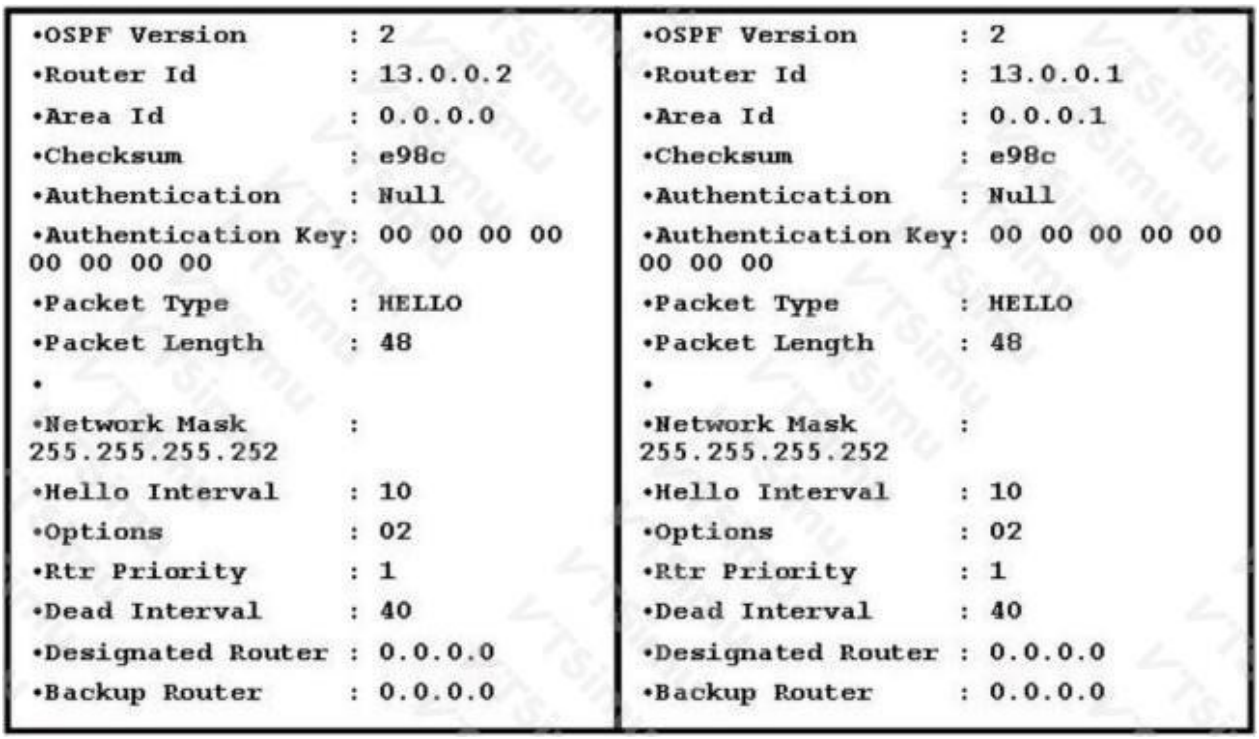
Explanation:

`config > router > static-route 0.0.0.0/0 next-hop 10.1.2.1` is the correct configuration. On the router at the right side of the point-to-point link, the default route must use the address of the adjacent router as its next hop. The connected peer is addressed as `10.1.2.1`, so packets destined for networks not already represented by a more-specific route are forwarded to that neighbor. The destination prefix `0.0.0.0/0` represents the IPv4 default route and is valid CIDR notation for a route that matches all IPv4 destinations when no longer prefix match exists.

In SR OS, IPv4 static routes are configured in the base routing context beneath `configure router`, using the `static-route` command followed by the destination prefix and a `next-hop` address. The next-hop address must be reachable through a directly connected interface or through an existing route; here, the point-to-point subnet provides that reachability. The static route becomes the router's gateway of last resort and enables the right-side router to send off-link traffic toward the router on the left. Nokia's SR OS routing documentation describes static-route configuration under the router context and the forwarding use of next-hop addresses: Nokia 7750 SR Router Configuration Guide and Nokia SR OS Static Routes documentation.

QUESTION NO: 3

Click the exhibit button.



•OSPF Version	: 2	•OSPF Version	: 2
•Router Id	: 13.0.0.2	•Router Id	: 13.0.0.1
•Area Id	: 0.0.0.0	•Area Id	: 0.0.0.1
•Checksum	: e98c	•Checksum	: e98c
•Authentication	: Null	•Authentication	: Null
•Authentication Key: 00 00 00 00 00 00 00 00		•Authentication Key: 00 00 00 00 00 00 00 00	
•Packet Type	: HELLO	•Packet Type	: HELLO
•Packet Length	: 48	•Packet Length	: 48
•		•	
•Network Mask 255.255.255.252	:	•Network Mask 255.255.255.252	:
•Hello Interval	: 10	•Hello Interval	: 10
•Options	: 02	•Options	: 02
•Rtr Priority	: 1	•Rtr Priority	: 1
•Dead Interval	: 40	•Dead Interval	: 40
•Designated Router	: 0.0.0.0	•Designated Router	: 0.0.0.0
•Backup Router	: 0.0.0.0	•Backup Router	: 0.0.0.0

Two neighboring OSPF routers have exchanged the packets shown. What state is the adjacency in?

A. Full; both routers are exchanging Hellos to maintain the adjacency.

B. The adjacency is down.

C. The adjacency is in exstart state.

D. The adjacency is in the exchange state.

ANSWER: C

Explanation:

The adjacency is in exstart state. After bidirectional neighbor communication has been established, OSPF enters ExStart to negotiate which router will act as the master for database synchronization and to establish the initial Database Description

(DBD) sequence number. The DBD packets exchanged during this negotiation use the initial, more, and master/slave control bits. These fields identify the negotiation phase before the routers begin the normal transfer of link-state database-summary information.

Once master/slave roles and the initial sequence number are agreed, the adjacency proceeds to Exchange. In that later phase, the routers continue sending DBD packets containing summaries of link-state advertisements, using the negotiated sequencing behavior. The packet exchange shown is characteristic of the preliminary DBD master/slave negotiation, so it identifies ExStart rather than the later database-summary exchange process. This state is an essential part of OSPF's reliable link-state database synchronization procedure. RFC 2328 defines the ExStart state and specifies the Database Description packet negotiation in Sections 10.6 and 10.8: OSPF neighbor states and Database Description packet process.

QUESTION NO: 4

Which of the following cannot be used as a matching criterion in route policy statements?

A. Source IP address

B. TCP source port

C. Prefix list

D. Protocol type

ANSWER: B

Explanation:

TCP source port is the correct answer because an SR OS route policy evaluates routing-information attributes, rather than fields found in individual data packets. Route policies are applied when routes are received, advertised, imported, exported, or redistributed. Their match conditions therefore operate on information carried by, or associated with, a route: for example, route prefixes, protocol origin, next-hop information, autonomous-system path data, community values, route tags, and related routing attributes.

A TCP source port exists only in the TCP header of a specific IP packet. It is a transport-layer value used by endpoints to identify an application flow or session, and it is not an attribute of an IP route in the routing table or a route advertisement. Consequently, a route-policy evaluation has no TCP source-port value available to inspect. Policies that need to classify or filter traffic using TCP or UDP ports are packet-filter, ACL, firewall, or service-policy functions; they are not route-policy functions.

Nokia's SR OS routing-policy documentation describes route policies as mechanisms for matching route characteristics and applying routing actions. See the Nokia SR OS Route Policies documentation and the route-policy overview.

QUESTION NO: 5

Which of the LSA types stay within an OSPF area and are not exported outside of the area? (Choose all that apply)

A. Router LSA

B. Network LSA

C. Summary Network LSA

D. Summary Router LSA

ANSWER: A B

Explanation:

Router LSA and Network LSA have area flooding scope, so they are retained within the OSPF area in which they originate. Every OSPF router generates one Router LSA for each area to which it is attached. This LSA describes the router's links,

interface costs, and neighboring relationships as seen from within that specific area. It provides the topology information each router needs to run the SPF calculation and build an intra-area shortest-path tree.

Network LSA is also limited to its originating area. On a broadcast or non-broadcast multi-access segment that has at least one fully adjacent neighbor, the designated router originates this LSA. It identifies the attached routers on that shared transit network, allowing routers in the same area to model the multi-access topology correctly during SPF processing.

These LSAs are the fundamental intra-area topology records: their information is used locally by routers participating in the area rather than being carried as the original LSAs across an area boundary. RFC 2328 defines Router LSAs in section 12.4.1 and Network LSAs in section 12.4.2, and specifies their area-level flooding scope. See RFC 2328, Router-LSAs and RFC 2328, Network-LSAs.

QUESTION NO: 6

Initially, all ports on an Alcatel-Lucent 7750 SR have 10Gbps bandwidth. RSVP configuration limits reservable bandwidth to 30 percent on all interfaces. An LSP is signaled reserving 1 Gbps bandwidth. How much unreserved bandwidth is left on the interface?

A. 10 Gbps

B. 7 Gbps

C. 3 Gbps

D. 2 Gbps

ANSWER: D

Explanation:

2 Gbps is correct. The interface has a physical bandwidth of 10 Gbps, but RSVP-TE does not make all of that capacity available for RSVP reservations in this scenario. The configured reservable-bandwidth limit is 30 percent, so the RSVP reservable pool is calculated as 30 percent of 10 Gbps: 0.30 × 10 Gbps = 3 Gbps. This value represents the maximum bandwidth that RSVP can advertise and allocate for RSVP-signaled LSPs on the interface; it is distinct from the port's total line rate.

When the LSP is signaled, it requests and reserves 1 Gbps from that 3 Gbps RSVP reservable pool. The remaining unreserved RSVP bandwidth is therefore 3 Gbps minus 1 Gbps, which equals 2 Gbps. This remaining value is the bandwidth available for a subsequent RSVP-TE LSP reservation, assuming no other RSVP reservations or constraints affect the interface. Nokia SR OS RSVP-TE documentation describes RSVP interface bandwidth parameters and their role in limiting bandwidth available to RSVP LSPs; see the Nokia SR OS documentation portal and the SR OS MPLS and RSVP-TE documentation.

QUESTION NO: 7

Which of the following statements are TRUE regarding sdp-ping? (Choose 3)

A. Provides a mechanism to determine the hops an SDP traverses.

B. Provides in-band uni-directional connectivity tests.

C. Provides in-band round-trip connectivity tests.

D. Tests ability of reaching the far-end IP address of an SDP ID within the SDP encapsulation.

E. Provides information regarding the exact MTU supported between service ingress and service termination.

ANSWER: B C D

Explanation:

sdp-ping is an SR OS OAM diagnostic used to verify forwarding connectivity for a configured Service Distribution Point. It sends an in-band OAM probe through the SDP transport encapsulation rather than relying solely on ordinary IP reachability outside the service transport path. Consequently, it validates whether the remote endpoint associated with the SDP can be reached using that SDP's encapsulated forwarding path. This is especially useful when troubleshooting whether a service can use its intended MPLS tunnel or GRE-based SDP transport.

The function supports a unidirectional connectivity test, allowing the operator to test delivery toward the far-end SDP endpoint. It also supports round-trip connectivity testing when the far end returns the OAM response, allowing confirmation that the complete forward-and-return OAM path is operational. Together, these capabilities distinguish an SDP-level forwarding test from a basic IP ping: the result reflects the data-plane transport path selected for the SDP and its remote termination.

Nokia's SR OS documentation describes SDP operational and OAM functions in the Nokia SR OS documentation portal. The underlying MPLS echo-request/echo-reply OAM model used for connectivity verification is standardized in RFC 8029.

QUESTION NO: 8

Which of the following is a function of the control plane in an MPLS router?

A. Label swapping

B. Pushing a label on an unlabeled packet

C. Label allocation

D. Routing lookups

ANSWER: C

Explanation:

Label allocation is a control-plane function because MPLS requires routers to create and distribute label bindings that associate a forwarding equivalence class with a locally meaningful MPLS label. The control plane uses routing information to determine reachable forwarding equivalence classes and then employs a label-distribution mechanism, such as LDP, RSVP-TE, BGP, or SR-related control-plane procedures, to advertise or learn the applicable label bindings. These bindings are used to build the Label Forwarding Information Base (LFIB), which is subsequently installed for packet handling. The MPLS architecture distinguishes this decision and signaling activity from the forwarding operation itself: the control plane determines the label bindings and forwarding state, while the data plane applies that precomputed state to packets. Therefore, label allocation is the appropriate control-plane responsibility. RFC 3031 describes MPLS label bindings and their distribution as core MPLS architectural functions: RFC 3031, Multiprotocol Label Switching Architecture. For the LDP protocol behavior used to distribute label mappings, see RFC 5036, LDP Specification.

QUESTION NO: 9

Which of the following statements regarding IS-IS Level 1 routing are true? Choose two answers.

A. Level 1 routing is limited to an IS-IS area.

B. Level 1 neighbors must share a common area address.

C. A Level 1 router maintains the complete Level 2 topology.

D. Level 1 adjacencies can form between routers in unrelated areas.

E. Level 1 routers advertise Type 5 LSAs for external routes.

ANSWER: A B

Explanation:

Level 1 IS-IS routing is scoped to a single IS-IS area. Level 1 routers maintain a Level 1 link-state database containing the topology and reachability information for their own area. A Level 1 adjacency requires the routers to share at least one common area address.

A Level 1 router normally uses the closest attached Level 1/Level 2 router as an exit point for destinations outside its area. It does not maintain the complete Level 2 topology unless it also operates at Level 2. IS-IS Level 1 and Level 2 operation is described in RFC 1195.

QUESTION NO: 10

A transport tunnel can either use an MPLS or GRE encapsulation.

A. TRUE

B. FALSE

ANSWER: A

Explanation:

TRUE is correct. In Nokia SR OS service architecture, a transport tunnel is the packet-forwarding path used to carry service traffic between endpoints across an IP/MPLS network. The tunnel's encapsulation identifies how the service packet is carried through the provider core. MPLS encapsulation is used when the transport infrastructure forwards labeled packets, allowing the service to use the MPLS transport and its associated label-switched paths. GRE encapsulation is also supported for IP-based tunneling scenarios, where the original packet is wrapped in a Generic Routing Encapsulation header for transport across an IP network. Thus, MPLS and GRE represent valid alternative encapsulation methods for a transport tunnel; the selected method depends on the configured service and transport design. Nokia's SR OS documentation organizes the relevant service and tunnel behavior under its service and routing configuration material; see the Nokia SR OS documentation portal. GRE itself is standardized as an IP encapsulation mechanism in RFC 2784, while MPLS provides the label-based transport model used in provider networks.

QUESTION NO: 11

Click on the exhibit.



Router R2 uses IS-IS to advertise the network 192.168.3.0/24 to R1. How can router R1 discard the route?

A. Use an import policy and a prefix list on router R2.

B. Use an export policy and a prefix list on router R1.

C. Use an import policy and a prefix list on router R1.

D. Router R1 cannot discard the route because IS-IS does not support import policies.

ANSWER: D

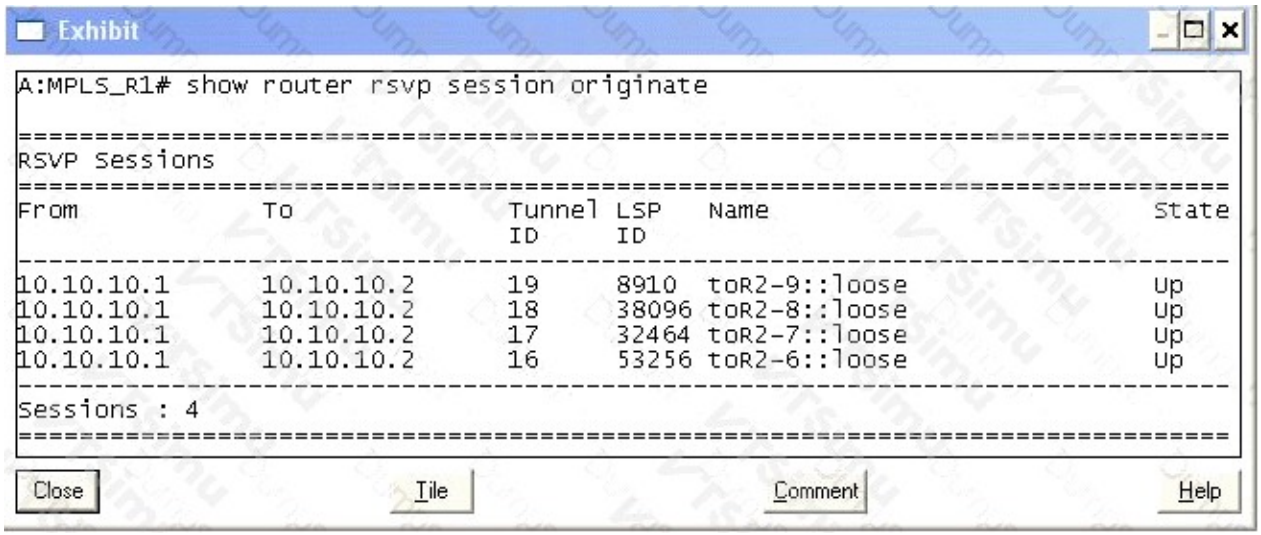
Explanation:

Router R1 cannot discard the route because IS-IS does not support import policies. IS-IS is a link-state IGP: R1 receives and retains the link-state information advertised by R2 in its IS-IS link-state database, then runs SPF to calculate the resulting routes. On SR OS, routing-policy control for IS-IS applies when routes are exported or redistributed into IS-IS; it is not an

inbound per-prefix import filter that R1 can use to suppress a specific route learned through the IS-IS protocol itself. Consequently, R1 cannot use a local IS-IS import policy and prefix list to reject 192.168.3.0/24 while continuing normal IS-IS operation. If the prefix must not be advertised through IS-IS, filtering must be applied at the originating/advertising router when the prefix is being injected or exported into IS-IS, or the network design and IS-IS advertisement must be changed accordingly. This behavior follows the link-state operation defined for integrated IS-IS and the SR OS IS-IS configuration model. See the Integrated IS-IS RFC 1195 and the Nokia SR OS documentation portal.

QUESTION NO: 12

Click on the exhibit button below.



Which fields in this output uniquely identify the RSVP session belonging to a single LSP? (Choose three)

A. Tunnel ID

B. Name

C. LSP ID

D. To

ANSWER: A C D

Explanation:

To, **Tunnel ID**, and **LSP ID** are the fields used to identify the RSVP-TE signaling context for an individual LSP in this output. The **To** value identifies the RSVP tunnel destination (the egress endpoint). The **Tunnel ID** identifies the RSVP-TE tunnel at that endpoint, allowing the control plane to distinguish one tunnel from another. The **LSP ID** then identifies the particular signaled LSP instance associated with that tunnel. This distinction is important because RSVP-TE maintains state for a tunnel and for the LSP sender state within it; the LSP identifier enables a router to distinguish a specific LSP instance, including during signaling refreshes or make-before-break operations. The LSP name is an operationally useful, locally configured label, but the protocol-level RSVP identification relies on signaling object fields rather than a descriptive name. RFC 3209 defines RSVP-TE SESSION and sender-template information, including the tunnel and LSP identifiers used by RSVP-TE signaling. See RFC 3209, RSVP-TE and Nokia's MPLS RSVP documentation.

QUESTION NO: 13

Choose the following components that are necessary for the configuration of a mirror service. (Choose 2)

A. A mirror destination to indicate which ports to monitor.

B. A mirror source to indicate which ports to monitor.

C. A mirror destination to indicate on which sap or sdp traffic should be replicated.

D. A mirror source to indicate which SAP or SDP traffic should be replicated.

ANSWER: A D

Explanation:

A mirror service requires a mirror source and a mirror destination. **A mirror source to indicate which SAP or SDP traffic should be replicated.** is required because the source definition selects the service access point or service distribution path traffic, along with the applicable direction, from which SR OS makes copies of packets. This identifies the traffic subject to mirroring. **A mirror destination to indicate which ports to monitor.** is also required because the destination defines the physical egress port that receives the replicated packet stream for observation by an analyzer, probe, or packet-capture device. Together, these elements establish the complete replication path: traffic is selected at the service attachment or tunnel endpoint and sent to the designated monitoring port. Nokia SR OS supports mirroring of service traffic for operational visibility, troubleshooting, and lawful or diagnostic packet analysis; the mirror destination is the monitored-output interface rather than the service traffic selector. The mirror-service configuration therefore must associate the selected source traffic with its destination port. Nokia's SR OS service documentation describes mirror services and their source and destination configuration model in the Mirror Services documentation. The broader SR OS documentation library is available from the Nokia SR OS documentation portal.

QUESTION NO: 14

Port 1/1/1 has been configured as an access port with encapsulation dot1q on an Alcatel-Lucent 7750 SR. Which of the following statements are TRUE? (Choose 2)

A. The SAP 1/1/1 will forward all traffic transparently.

B. The SAP 1/1/1:0.* will forward all untagged traffic and traffic tagged with VLAN 0.

C. The SAP 1/1/1:* will forward untagged traffic but not tagged traffic.

D. TheSAP1/1/1:400 will forward customer traffic tagged with VLAN 400

E. The SAP 1/1/1:0 will forward untagged traffic and traffic tagged with VLAN 0

ANSWER: D E

Explanation:

With `encap-type dot1q`, a Service Access Point uses a VLAN identifier in its SAP encapsulation value to classify incoming customer traffic. The SAP `1/1/1:400` therefore identifies the customer VLAN whose IEEE 802.1Q tag has VLAN ID 400. Frames carrying that VLAN tag are mapped to the service associated with this SAP, making "TheSAP1/1/1:400 will forward customer traffic tagged with VLAN 400" correct.

SR OS also reserves the dot1q SAP identifier `0` for traffic that has no customer VLAN classification: untagged frames and priority-tagged frames whose VLAN ID is zero. A priority-tagged Ethernet frame retains an 802.1Q header but uses VLAN ID zero, so it is handled with the untagged traffic class rather than as a regular customer VLAN. Consequently, `1/1/1:0` is the appropriate SAP identifier for both untagged traffic and VLAN-0-tagged traffic. This is why "The SAP 1/1/1:0 will forward untagged traffic and traffic tagged with VLAN 0" is also correct. Nokia's SR OS Layer 2 Services documentation describes SAP encapsulation identifiers and their role in service traffic classification: SAP encapsulation identifiers. IEEE 802.1Q VLAN tag semantics, including priority-tagged frames, are summarized by the IEEE 802.1Q standard.

QUESTION NO: 15

How many SAPs can be configured on a port configured with NULL encapsulation?

A. 1

B. 4096

C. None. This encapsulation is not supported.

D. The number is limited only by the capacity of the router

ANSWER: A

Explanation:

1 is correct. On Nokia SR OS, a port using *null* encapsulation has no Layer 2 encapsulation identifier available to distinguish traffic for multiple service access points. The entire physical port is therefore represented by a single SAP, conventionally configured using a null SAP identifier such as `port-id:*`. All frames received on that port are associated with that one SAP and its attached service context. This model is appropriate when the port is dedicated to one service and there is no VLAN-tag, QinQ tag, or other encapsulation value on which SR OS could perform service demultiplexing. A SAP is the binding between a customer-facing interface and a service; when encapsulation is null, the port itself is that sole binding point. Consequently, exactly one SAP can be configured for a port configured with NULL encapsulation. Nokia's Layer 2 service documentation describes SAP encapsulation as the mechanism used to identify traffic entering a service and documents null encapsulation as a non-multiplexed port use case. See the Nokia SR OS Layer 2 Services Guide and the Nokia SR OS Interface Configuration Guide.

QUESTION NO: 16

What is the default preference value for a static route in the Alcatel-Lucent 7750 SR router?

A. 0

B. 5

C. 10

D. 15

ANSWER: B

Explanation:

The default preference for a static route on an Alcatel-Lucent 7750 SR, now part of the Nokia 7750 SR product family, is **5**. In SR OS, route preference is the administrative value used to select between routes to the same destination that have been learned from different route sources. A lower preference value is more preferred. Consequently, the low default assigned to a static route allows an operator-configured route to take precedence over many dynamically learned alternatives when both provide reachability for the same prefix. The static-route configuration supports an explicit preference setting when a design requires different route-selection behavior, but omitting that setting retains the platform default of 5. This behavior is part of the base-router static-route implementation and is relevant when interpreting the active route in the routing table or deliberately designing primary and backup forwarding paths. Nokia's SR OS documentation describes static-route configuration and the route-preference attribute in its router configuration material; see the Nokia SR OS static routes documentation and the Nokia SR OS routing documentation.

QUESTION NO: 17

The Alcatel-Lucent 7750 SR supports which two of the OSPF area types below?

A. Not So Stubby Areas

B. Level 1 Areas

C. Stub Areas

D. Partially Stubby Areas

ANSWER: A C

Explanation:

Not So Stubby Areas and **Stub Areas** are supported OSPF area types on the 7750 SR. A stub area limits the inter-area routing information distributed into the area and uses a default route to reach destinations outside the area. This reduces the size of the link-state database and routing table maintained by routers within that area, which is useful where a simple path toward the rest of the OSPF domain is appropriate.

A Not-So-Stubby Area (NSSA) retains the operational benefits of a stub-style area while permitting an autonomous-system boundary router inside that area to inject external reachability. NSSA external information is originated as Type 7 LSAs and is translated to Type 5 external LSAs when it is advertised into normal OSPF areas. This behavior is defined by the NSSA extension to OSPF and is implemented by the 7750 SR OSPF feature set. Nokia's 7750 SR routing documentation covers OSPF configuration and supported area behavior in its Service Router documentation library. The protocol definition for NSSA operation is available in RFC 3101.

QUESTION NO: 18

Which of the following LSA types stay within an OSPF area, and are not flooded outside of the area? Choose two answers.

A. Router LSA

B. Network LSA

C. Summary router LSA

D. Type 5 (AS external) LSAs

ANSWER: A B

Explanation:

Router LSA and **Network LSA** are the OSPF link-state advertisements whose flooding scope is limited to a single area. A Router LSA, also called a Type 1 LSA, is originated by every router in an area and describes that router's links, interfaces, costs, and adjacencies within the area's topology database. This information allows all routers in the same area to independently calculate the intra-area shortest-path tree.

A Network LSA, or Type 2 LSA, is originated by the designated router on a broadcast or NBMA multi-access network. It represents the shared network segment and identifies the routers attached to it. Like Router LSAs, Network LSAs are area-scoped: an ABR does not flood them into another OSPF area. Instead, an ABR uses the detailed intra-area topology learned from these LSAs to calculate and advertise inter-area reachability where required. This area boundary behavior supports OSPF's hierarchical design, limiting topology-database size and SPF-processing scope. RFC 2328 defines Type 1 and Type 2 LSAs and specifies their area flooding scope in its LSA descriptions and flooding rules: RFC 2328: OSPF Version 2. The IETF's OSPF working-group resources are available at IETF OSPF Working Group.

QUESTION NO: 19

Click on the exhibit.

Given the contents of the Router LSA, which of the following best describes the local topology of the router?

A. The router has a system interface, two point-to-point interfaces and one broadcast interface.

B. The router has a system interface, two passive interfaces, two point-to-point interfaces and one broadcast interface.

C. The router has a system interface, two point-to-point interfaces and one virtual link.

D. The router has a system interface, two passive interfaces, two point-to-point interfaces and one virtual link.

ANSWER: A

Explanation:

The router has a system interface, two point-to-point interfaces and one broadcast interface. A Router-LSA describes the links that originate on the advertising router, allowing its local OSPF topology to be inferred from the link records in the LSA. The system interface is represented as a stub-style destination associated with the router itself, while the two point-to-point link descriptions identify adjacencies to neighboring routers over point-to-point media. The remaining link is a transit-network connection, which corresponds to a broadcast multi-access network represented in OSPF by a network pseudonode and its Network-LSA. Taken together, those records describe one system interface, two point-to-point interfaces, and one broadcast interface. This interpretation follows the Router-LSA link-type definitions in OSPF version 2: point-to-point links describe router-to-router connectivity, transit links identify attachment to a multi-access network, and stub links advertise reachable prefixes. Nokia SR OS OSPF uses these standard LSA semantics when advertising topology information into an area. See RFC 2328, Router-LSAs and Nokia SR OS OSPF documentation.

QUESTION NO: 20

Which of the following statements are TRUE about configuring a distributed Epipe service? (Choose 2)

A. The vc-id on the spoke-sdp must be explicitly configured.

B. The vc-id on the mesh-sdp must be explicitly configured.

C. The vc-id on the spoke-sdp can be configured to match the service id.

D. The vc-id on the mesh-sdp can be configured to match the service id.

E. The vc-id on the spoke-sdp does not have to be explicitly configured and will use the service id.

F. The vc-id on the mesh-sdp does not have to be explicitly configured and will use the service id.

ANSWER: A C

Explanation:

In a distributed Epipe, the service is extended between provider edge routers through a spoke SDP binding. The pseudowire endpoint must be identified consistently at both ends of the SDP, so the spoke-SDP VC ID is explicitly configured as part of the binding. This VC ID is the virtual-circuit label that associates the local pseudowire endpoint with its corresponding remote endpoint; it is therefore a required element of the distributed Epipe configuration.

The VC ID selected for the spoke SDP may be the same numeric value as the service ID. This is a common and valid operational convention because it makes the pseudowire identifier easy to correlate with the service during provisioning, monitoring, and troubleshooting. The VC ID remains an explicitly configured spoke-SDP parameter even when its value matches the service ID; matching the values does not make the VC ID implicit.

Nokia's SR OS service documentation describes Epipe as a point-to-point service and documents the use of spoke SDP bindings to extend services across the packet network. See the Nokia SR OS Epipe service overview and the Nokia SR OS service SDP bindings documentation.

QUESTION NO: 21

In an OSPF Hello packet, which of the following fields must match for all neighbor routers on the segment? Choose three answers.

A. Area ID W B.

B. Hello and Dead intervals PC.

C. Stub flag

D. Checksum values

E. The list of neighbors

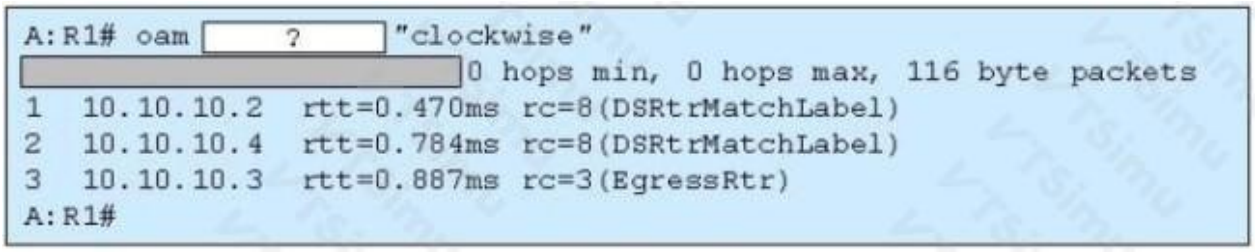
ANSWER: A B C

Explanation:

OSPF routers use Hello packets to discover neighbors and to verify that they share the fundamental parameters required to form an adjacency on the same link. The Area ID must be identical because an OSPF interface belongs to one area, and neighbors exchange routing information within that common area. Hello and Dead intervals must also agree: the Hello interval defines the expected frequency of Hellos, while the Router Dead Interval defines how long a router waits before declaring a neighbor unavailable. Matching timers provides a consistent neighbor-liveness relationship. The Stub flag is part of the Hello packet's Options field and reflects area capabilities and type. Routers on a common segment must agree on the relevant stub-area capability so that they operate with consistent area behavior. RFC 2328 specifies the Hello packet format, including the Area ID, HelloInterval, RouterDeadInterval, Options field, and neighbor list, and specifies the compatibility checks performed when a Hello is received. These matching values allow routers to recognize one another as eligible OSPF neighbors before progressing toward adjacency and database synchronization. See RFC 2328, section 9.5 and RFC 2328, Hello packet format.

QUESTION NO: 22

Click the exhibit.



Which OAM tool produced this output?

A. vprn-ping

B. sdp-ping

C. traceroute

D. lsp-trace

ANSWER: D

Explanation:

lsp-trace is the Nokia SR OS MPLS OAM utility used to trace the forwarding path of a label-switched path. It sends MPLS echo requests for the target LSP and uses TTL-expiry behavior at successive label-switching routers to identify the transit hops and the egress reached by the LSP. Consequently, its output is structured around MPLS/LSP path discovery rather than ordinary IP hop discovery. This makes it appropriate when the displayed results identify an MPLS forwarding path, including responses from LSRs and details associated with MPLS labels, router interfaces, and the traced FEC. The command is especially useful for validating the actual data-plane path of RSVP-TE, LDP, BGP-labeled, or other MPLS-signaled transport tunnels, and for locating a break or unexpected forwarding point along that path. Nokia documents LSP trace as part of its MPLS OAM functionality and describes its use of MPLS echo-request/echo-reply processing for LSP path verification and troubleshooting. See the Nokia SR OS OAM Guide and the MPLS Guide for the related MPLS OAM concepts.

QUESTION NO: 23

From the list below identify two matching criteria supported for IP filters on the Alcatel-Lucent 7750 SR: (Choose two)

A. DSCP marking

B. Class of Service Marking

C. DSAP marking

D. ICMP type

ANSWER: A D

Explanation:

IP filters on the 7750 SR can classify packets using Layer 3 and Layer 4 header fields. **DSCP marking** is a supported IPv4/IPv6 filter match criterion because the Differentiated Services Code Point is carried in the IP header and can be used to identify traffic that has received a particular QoS marking. A filter entry can therefore match a configured DSCP value or DSCP set and apply the associated filter action.

ICMP type is also supported because ICMP control and diagnostic traffic contains a type field that identifies the ICMP message category, such as echo request or destination unreachable. Matching that field allows the router to classify and control specific ICMP traffic types without treating all IP traffic identically. These criteria are part of the protocol-aware matching capabilities available in SR OS IP filter policies, which are used in contexts such as interfaces, services, and control-plane protection.

Nokia's SR OS Router Configuration Guide documents IP filter match fields and filter-entry configuration: SR OS filter policies. The broader SR OS documentation library is available at Nokia SR OS documentation.

QUESTION NO: 24

Which of the following statements is TRUE concerning the advertisement of VRF routes on a 7750 SR?

A. By default, a 7750 SR does not advertise all routes in the VRF to a CE router.

B. The default 7750 SR behavior advertises all routes in a VRF to other PE routers.

C. By default, a 7750 SR does not advertise all routes in the VRF to other PE routers.

D. By default, a 7750 SR advertises all routes in the VRF to a CE router.

ANSWER: A B

Explanation:

By default, a 7750 SR does not advertise all routes in the VRF to a CE router. Route advertisement toward a CE is governed by the routing context and the applicable routing-protocol behavior and policy. In particular, the PE does not simply send the complete contents of its VPRN/VRF routing table to every attached CE. This avoids indiscriminate distribution of customer routes and permits the operator to control CE-facing routing information through protocol configuration and import/export policy.

The default 7750 SR behavior advertises all routes in a VRF to other PE routers. In an MPLS Layer 3 VPN, the PE distributes VPN-IPv4/VPN-IPv6 routes for the VPRN using MP-BGP, carrying route-target information that lets receiving PEs determine which VRF should import each route. This PE-to-PE control-plane exchange is what makes routes learned at one VPN site available to other sites that belong to the same VPN. Nokia's VPRN documentation describes this MP-BGP VPN route distribution model and the role of VRF policies in controlling route exchange. See the Nokia SR OS VPRN documentation and the Nokia SR OS BGP documentation.

QUESTION NO: 25

A new router is added to a broadcast network. What is used by IS-IS as the tie breaker for selecting the DIS if the priorities are the same?

A. The router with the highest system ID.

B. The router with the highest loopback address.

C. The Hello packet with the highest sequence number.

D. The existing DIS remains the DIS.

E. The router with the highest interface MAC address.

ANSWER: E

Explanation:

The router with the highest interface MAC address is correct. On an IS-IS broadcast LAN, such as Ethernet, routers elect a Designated Intermediate System (DIS) independently for each IS-IS level. The primary election criterion is the configured IS-IS interface priority advertised in IS-IS Hello (IIH) packets. When participating routers advertise the same priority, IS-IS uses the Subnetwork Point of Attachment (SNPA) as the deterministic tie breaker. For an Ethernet broadcast interface, the SNPA is the interface's MAC address, and the numerically highest MAC address wins the DIS election.

This selection method ensures that every router on the LAN reaches the same result from the information carried in received IIHs. The elected DIS performs key LAN functions, including generating pseudonode link-state advertisements to represent the multiaccess segment in the IS-IS link-state database. A newly connected router can therefore become DIS when its advertised priority is equal to the current candidates' priority but its Ethernet MAC address is numerically higher. The IS-IS election behavior and the use of the highest SNPA for a priority tie are specified in the Integrated IS-IS protocol definition: RFC 1195. The underlying IS-IS LAN election procedures are also documented in RFC 1142.

QUESTION NO: 26

Which of the following LDP messages are used in the successful establishment of LDP sessions? Choose three answers

A. Label request message

B. Notification message

C. Address message

D. Keepalive message

E. Initialization message

F. Hello message

ANSWER: D E F

Explanation:

"Hello message," "Initialization message," and "Keepalive message" are the messages involved in successfully establishing an LDP session. LDP peer discovery begins when routers exchange Hello messages over UDP port 646. These Hellos advertise the router's LDP identifier and transport address, allowing eligible peers on a link or targeted adjacency to discover one another. Once discovery identifies a peer, the routers establish a TCP connection, normally using port 646, to carry the reliable LDP session protocol. Each peer then sends an Initialization message to negotiate session parameters, including protocol version, keepalive settings, label-distribution mode, and other capability-related values. After a valid Initialization message is received, each peer sends a Keepalive message. Receipt of that Keepalive completes initialization of the session; subsequent Keepalive messages maintain the established session when no other LDP messages are being exchanged. This sequence is defined by the LDP specification: discovery via Hello messages, TCP session establishment, Initialization exchange, and Keepalive exchange. See the LDP session-establishment procedures in RFC 5036, especially its discovery and session-initialization sections, and Nokia's LDP documentation.

QUESTION NO: 27

Customer 'A' has sites on 3 different Alcatel-Lucent PE routers. The routers are connected to each other through an IP/MPLS network in a full mesh fashion. Customer 'A' requires a VPLS service. Which of the following statements are TRUE? (Choose2)

A. Each router requires 1 VPLS service.

B. Each router requires 2 VPLS services.

C. Each router requires 1 SDP with 2 LSPs, one to each of the other 2 routers.

D. Each router requires 2 SDPs.

E. Each router requires 3 SDPs

ANSWER: A D

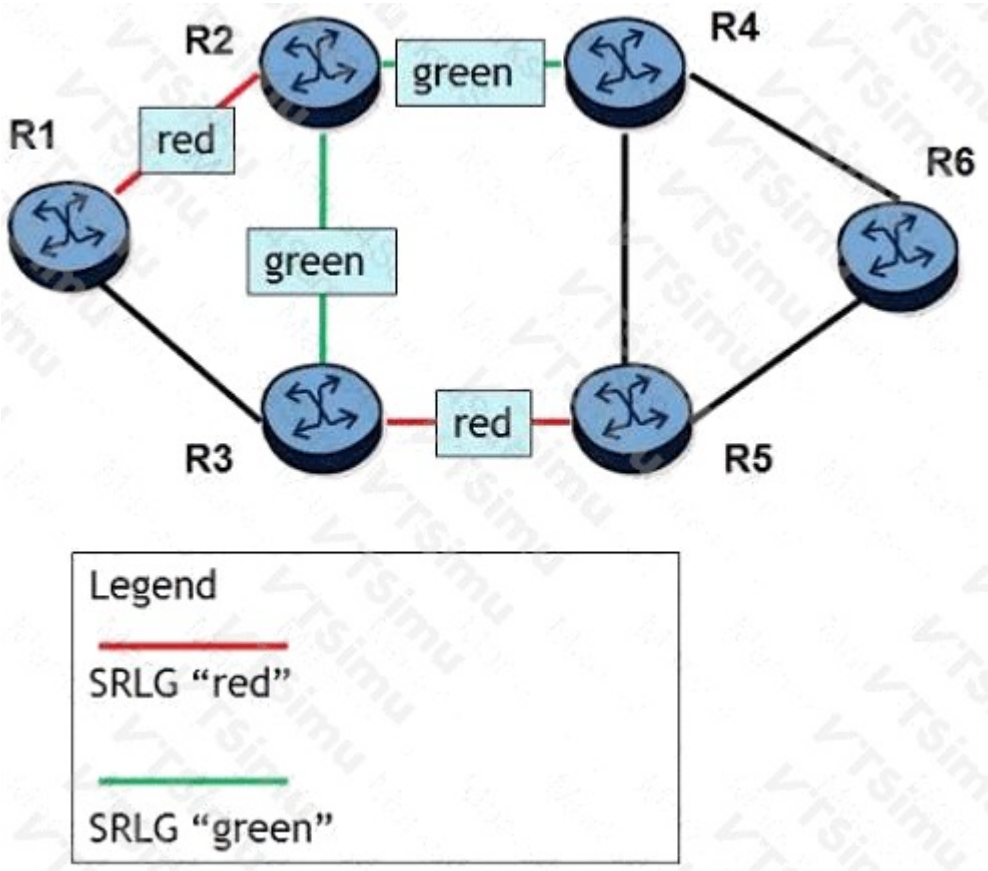
Explanation:

“Each router requires 1 VPLS service.” is correct because a VPLS is a single multipoint Layer 2 service instance that represents the customer’s common bridged LAN across all participating provider-edge routers. The service is instantiated on every participating router using the same service identifier; it is not a separate VPLS service per remote site. Customer-facing SAPs and the VPLS mesh connectivity are associated with that one local service instance.

“Each router requires 2 SDPs.” is also correct. A full-mesh VPLS requires each provider-edge router to establish service connectivity to every other participating provider-edge router. With three routers total, each router has two remote peers. Nokia SR OS uses an SDP to define the service-transport relationship toward a remote provider-edge router, and the local VPLS service uses mesh SDP bindings over those SDPs. Thus, each router needs one SDP for each of its two remote peers, yielding two SDPs per router. This topology provides the required full-mesh pseudowire connectivity for the VPLS forwarding domain. Nokia’s Layer 2 service documentation describes VPLS service components and mesh SDP bindings; the SR OS service overview also covers SDP-based service transport. See Nokia SR OS VPLS documentation and Nokia SR OS Layer 2 service overview.

QUESTION NO: 28

Click on the exhibit.



An LSP ' s primary path takes R1-R3-R5-R6. Its fully loose secondary path is configured for SRLG. All links are of equal cost. Which path will the secondary path take?

A. The secondary path will take R1-R3-R2-R4-R6.

B. The secondary path will take R1-R2-R4-R6.

C. The secondary path will take R1-R2-R3-R5-R6.

D. The secondary path will take R1-R3-R2-R4-R5-R6.

ANSWER: A

Explanation:

The secondary path will take R1-R3-R2-R4-R6. SRLG protection is intended to avoid a shared physical-risk failure, rather than merely avoid every router or every link used by the primary LSP. During computation of the fully loose standby path, the ingress uses the SRLG information associated with the primary path and selects an end-to-end route that does not traverse resources belonging to the protected primary path's relevant shared-risk groups. In the displayed topology, R1-R3-R2-R4-R6 is the route that satisfies the SRLG-disjointness requirement for the primary path R1-R3-R5-R6.

"Fully loose" is important because it permits normal constrained shortest-path computation between the ingress and egress instead of imposing intermediate strict hops. Once the SRLG constraint has removed the shared-risk resources from eligibility, all remaining links have equal metric, so the selected route is the eligible end-to-end path shown in the exhibit. This behavior supplies resilience against a single underlying failure that could affect multiple logical links, which is the purpose of SRLG-aware MPLS traffic-engineering protection. SRLGs are defined for MPLS TE signaling in RFC 4202; Nokia SR OS MPLS and traffic-engineering documentation is available through the Nokia SR OS documentation portal.

QUESTION NO: 29

What is the significance of receiving a packet with the bottom MPLS label S bit set to 1?

A. It results in forwarding based on the next packet header.

B. It indicates that the label is the bottom of the MPLS label stack.

C. It indicates that the label is the top of the MPLS label stack.

D. Any LSR that receives an S bit set to 1 must discard the packet.

ANSWER: B

Explanation:

It indicates that the label is the bottom of the MPLS label stack. In the MPLS label-stack entry format, the S field is the Bottom of Stack bit. A value of 1 marks the current label-stack entry as the final (lowest) label in the MPLS stack; consequently, no additional MPLS label lies beneath it. This bit lets a label switching router identify the boundary between MPLS encapsulation and the payload carried after the label stack. The payload following that final label may be an IPv4 packet, IPv6 packet, pseudowire control word, or another protocol indicated by the forwarding context and applicable MPLS specifications. The S bit is carried in every 32-bit MPLS label-stack entry alongside the 20-bit label value, traffic-class field, and TTL field. It is a structural indicator of label-stack depth, not an instruction to discard the packet. RFC 3032 defines the label-stack encoding and specifies that the Bottom of Stack field is set to 1 for the last entry in the stack. Nokia's SR OS MPLS documentation also describes MPLS forwarding and label-stack processing in the context of its router implementation. See RFC 3032, MPLS Label Stack Encoding and Nokia SR OS MPLS overview.

QUESTION NO: 30

Click on the exhibit button below.

Given the configuration below, which of the following commands would enable fast reroute on the LSP? Choose two answers

A. This configuration already is configured for FRR.

B. fast-reroute facility

C. fast-reroute one-to-one

D. secondary

E. standby

ANSWER: B C

Explanation:

`fast-reroute facility` and `fast-reroute one-to-one` are the two SR OS MPLS LSP configuration commands that enable RSVP-TE Fast Reroute protection. Fast Reroute establishes local protection so that, when a protected link or next-hop node fails, the point of local repair can switch MPLS traffic to a pre-signaled detour or bypass path without waiting for end-to-end LSP re-signaling. The facility method uses a shared bypass tunnel that can protect multiple eligible LSPs, whereas the one-to-one method creates a detour LSP associated with the protected primary LSP. Both therefore represent valid FRR operating modes that may be selected under the applicable LSP configuration context. Their behavior follows the RSVP-TE local-repair models defined in RFC 4090, including facility backup and one-to-one detour protection. Nokia's SR OS documentation library contains the MPLS and RSVP-TE command and operational guidance for the relevant software release. See RFC 4090: Fast Reroute Extensions to RSVP-TE and the Nokia SR OS documentation portal.

QUESTION NO: 31

Which of the following statements regarding LSPs are true? (Choose two)

A. IGP-based LSPs are established solely from information contained in routing tables.

B. The constraint-based routed LSP is calculated at the ingress of the network, based on the router database and user defined constraints.

C. The constraint-based routed LSP is calculated at each hop of the LSP, based on the route table and user defined constraints.

D. IGP-based LSPs must use LDP signaling.

ANSWER: A B

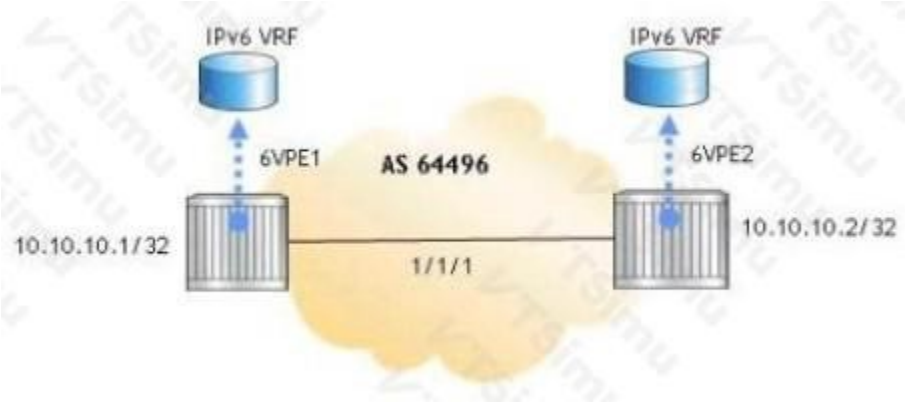
Explanation:

IGP-based LSPs are established solely from information contained in routing tables. This describes a basic MPLS forwarding path in which the label-switched path follows the next-hop information selected by the underlying interior gateway protocol. The route to the forwarding-equivalence-class destination determines the sequence of next hops, and label distribution associates MPLS labels with that routed path. This model does not require traffic-engineering constraints or a separately computed constrained path.

The constraint-based routed LSP is calculated at the ingress of the network, based on the router database and user defined constraints. In RSVP-TE, the head-end router performs constrained shortest-path-first processing using traffic-engineering topology information held in its traffic-engineering database. The calculation can account for configured requirements such as bandwidth, administrative groups, metrics, and explicit path requirements. After selecting an eligible end-to-end path, the ingress signals the LSP along that path; intermediate routers process the signaling and reserve applicable resources rather than independently recalculating the complete constrained route. Nokia's MPLS documentation describes both IGP/LDP-based transport and RSVP-TE constraint-based path computation. See the Nokia SR OS MPLS Guide and RFC 3209: RSVP-TE.

QUESTION NO: 32

Click on the exhibit button below.



Which of the following is the correct configuration for an MP-BGP session on 6VPE1?

A. PE1>config>router>bgp#info
---
group "multi-bgp"
family vpn-ipv6
neighbor 10.10.10.2/32
local-address 10.10.10.1
peer-as 64496
exit
exit

B. PE1>config>router>bgp#info
---
group "multi-bgp"
family ipv6 vpn-ipv4
neighbor 10.10.10.2
local-address 10.10.10.1
peer-as 64496
exit
exit

C. PE1>config>router>bgp#info
---
group "multi-bgp"
family vpn-ipv6
neighbor 10.10.10.2
local-address 10.10.10.1
peer-as 64496
exit
exit

D. PE1>config>router>bgp#info
---
group "multi-bgp"
family ipv6 vpn-ipv6
neighbor 10.10.10.2/32
local-address 10.10.10.1
peer-as 64496
exit
exit

ANSWER: C

Explanation:

The configuration containing `family vpn-ipv6`, neighbor `10.10.10.2`, local address `10.10.10.1`, and peer AS `64496` is correct for this MP-BGP peering. The VPN-IPv6 address family enables the BGP session to carry IPv6 VPN routes, including the route distinguisher and MPLS label information required for IPv6 Layer 3 VPN service transport. This is the appropriate MP-BGP family when 6VPE devices advertise VPN-IPv6 reachability across an IPv4 MPLS provider core.

On SR OS, the BGP neighbor is configured as an IP address rather than as an address/prefix-length value; therefore, the peer is specified as `10.10.10.2`. The `local-address 10.10.10.1` statement selects the local IPv4 endpoint for the BGP TCP session, while `peer-as 64496` identifies the remote autonomous system. This combines an IPv4 transport session with the VPN-IPv6 multiprotocol capability needed by the PE relationship. Nokia's BGP configuration documentation describes address-family activation under BGP groups and neighbors: Nokia SR OS BGP documentation. The VPN-IPv6 NLRI encoding and use of labels for IPv6 VPNs are standardized in RFC 4659.

QUESTION NO: 33

Click on the exhibit.

```
*A:MPLS_R5# show router mpls lsp "toR8-5" path detail

===============================================================================
MPLS LSP toR8-5 Path  (Detail)
===============================================================================
Legend :
    @ - Detour Available                 # - Detour In Use
    b - Bandwidth Protected              n - Node Protected
    s - Soft Preemption
===============================================================================
-------------------------------------------------------------------------------
LSP toR8-5 Path loose
-------------------------------------------------------------------------------
LSP Name     : toR8-5                          Path LSP ID : 27694
From         : 10.10.10.5                      To          : 10.10.10.8
Adm State    : Up                              Oper State  : Down
Path Name    : loose                           Path Type   : Primary
Path Admin   : Up                              Path Oper   : Down
OutInterface: n/a                              Out Label   : n/a
Path Up Time: 0d 00:00:00                      Path Dn Time: 0d 00:01:40
Retry Limit : 0                                Retry Timer : 30 sec
RetryAttempt: 4                                NextRetryIn : 21 sec
SetupPriori*: 7                                Hold Priori*: 0
Preference   : n/a
Bandwidth    : No Reservation                  Oper Bw     : 0 Mbps
Hop Limit    : 255                             Class Type  : 0
Backup CT    : None
MainCT Retry: n/a                              MainCT Retry: 0
    Rem      :                                     Limit   :
Oper CT      : None
Record Route: Record                           Record Label: Record
Oper MTU     : 0                               Neg MTU     : 0
Adaptive     : Enabled                         Oper Metric : 65535
Include Grps:                                  Exclude Grps:
None                                         None
Path Trans   : 2                               CSPF Queries: 1
Failure Code: looseHopsInFRRLsp                Failure Node: 10.10.10.5
ExplicitHops:
    No Hops Specified
Actual Hops :
    No Hops Specified
ResigEligib*: False
LastResignal: n/a                              CSPF Metric : 0
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

LSP " to R8-5 " requests fast reroute protection. What could be done to remove the Failure Code shown?

A. The LSP must be enabled with CSPF on the head-end router.

B. Fast reroute must be enabled on the downstream routers.

C. A mixture of strict and loose hops must be defined in the primary path.

D. CSPF and fast reroute must be enabled on all downstream routers.

ANSWER: A

Explanation:

The LSP must be enabled with CSPF on the head-end router. Fast reroute protection requires a path computation that can honor MPLS traffic-engineering constraints and select an eligible protected route. In SR OS, enabling CSPF for the ingress LSP allows the head-end to compute the RSVP-TE path using TE database information rather than relying only on an administratively specified route. This is especially important when the LSP requests local repair, because the signaled path and its protection requirements must be compatible with the topology and available TE resources. Once CSPF is enabled on the head-end LSP, the router can calculate and signal a route suitable for the requested fast-reroute behavior, removing the displayed failure condition when an eligible protected path exists. CSPF also ensures that the route calculation can apply constraints such as bandwidth, metrics, affinities, and explicit-hop requirements consistently with RSVP-TE signaling. Nokia's SR OS documentation describes CSPF as the mechanism used to calculate constrained MPLS TE LSP paths, while the RSVP fast-reroute specification defines the local-protection behavior requested through RSVP-TE signaling. See the Nokia SR OS documentation portal and RFC 4090, Fast Reroute Extensions to RSVP-TE.

QUESTION NO: 34

On an Alcatel-Lucent 7750 SR, which vc-type will forward service delimiting tags at the ingress PE?

A. vc-type ether

B. vc-type vlan

C. vc-type dot1q

D. vc-type null

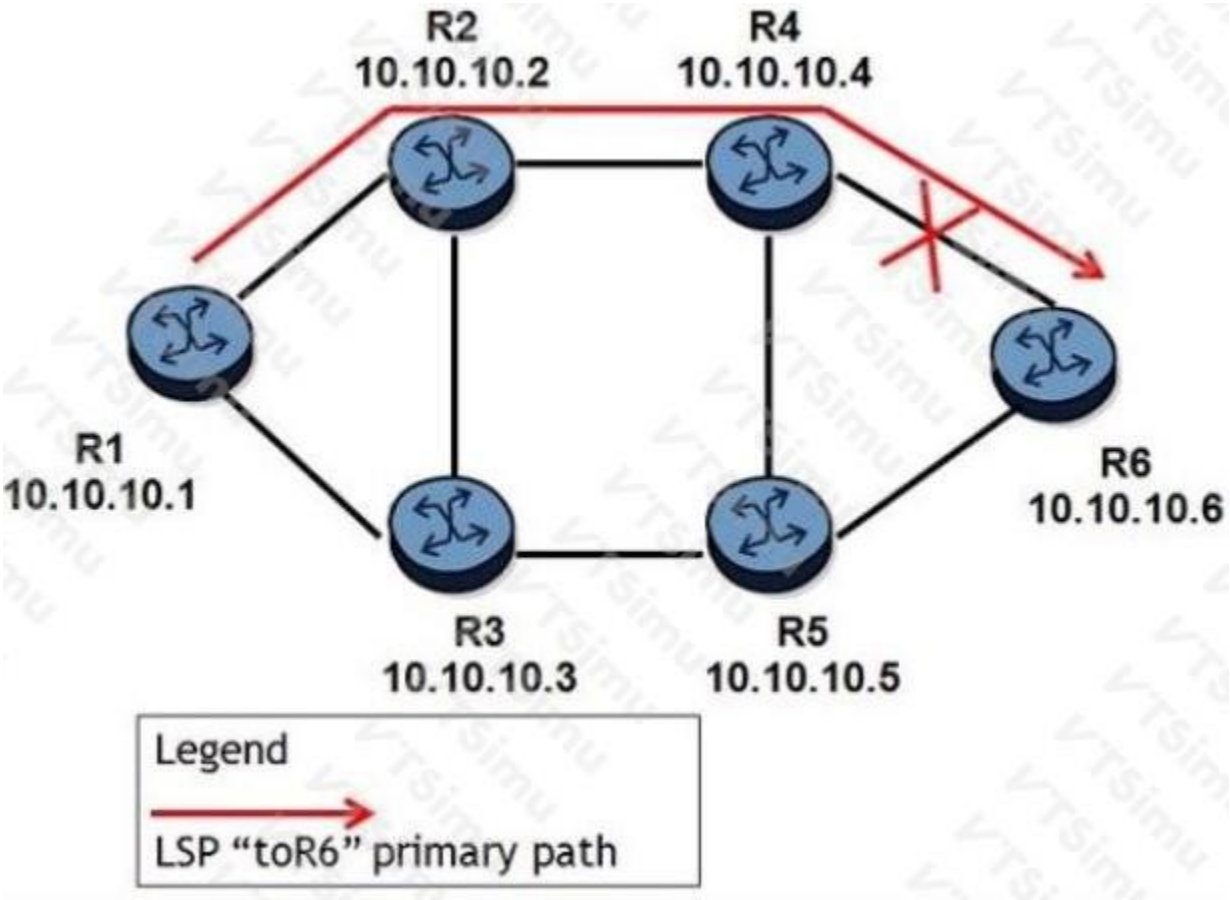
ANSWER: B

Explanation:

vc-type vlan is correct because it establishes an Ethernet VLAN-tagged pseudowire. In this mode, the ingress PE retains and transports the customer's service-delimiting IEEE 802.1Q VLAN tag across the pseudowire rather than treating the Ethernet frame as an untagged, raw Ethernet payload. This behavior is needed when the VLAN identifier is part of the customer frame information that must remain available to the far-end PE or customer-facing attachment circuit. The pseudowire type is also advertised through LDP signaling, allowing both PEs to agree on VLAN-tagged Ethernet encapsulation for the service. The underlying Ethernet pseudowire specification defines a tagged Ethernet mode in which the customer VLAN tag is preserved, matching the operational purpose of the 7750 SR **vc-type vlan** setting. Nokia SR OS service configuration uses the VC type to define the payload treatment and signaling associated with the service, so selecting the VLAN VC type ensures the ingress PE forwards the delimiting tag as required. See RFC 4448: Encapsulation Methods for Transport of Ethernet over MPLS Networks and Nokia Documentation Center.

QUESTION NO: 35

Click on the exhibit.



An LSP path is established on router R1 toward router R6. After router R4 detects that the link between routers R4 and R6 goes down, which of the following will affect the failure propagation time?

A. The time router R4 takes to send a PATH Error message toward the head-end router.

B. The time router R4 takes to detect RSVP Hello message timeout from router R6.

C. The time router R4 takes to send a RESV Error message to the head-end router.

D. The time router R1 takes to signal an alternate path through router R3.

ANSWER: A

Explanation:

The time router R4 takes to send a PATH Error message toward the head-end router affects failure propagation time because R4 is the point at which the established RSVP-TE LSP becomes broken after the R4-to-R6 adjacency fails. Once the failure has been detected, R4 generates an RSVP PathErr message identifying the path failure and sends it upstream along the reverse path-state route toward the ingress, R1. The time needed to originate, process, and forward that notification over each upstream hop contributes directly to the interval before the head-end learns that its currently signaled LSP is no longer usable. Upon receiving the PathErr, the head-end can initiate the configured recovery action, such as attempting to signal a replacement LSP through the available alternate route. RSVP-TE defines PathErr as the mechanism used to report path-related errors upstream to the sender, including failures encountered by a transit node while maintaining an explicitly routed LSP. This makes PathErr propagation the relevant control-plane notification phase after R4 has already detected the failed link. See RFC 3209, RSVP-TE and Nokia SR OS RSVP-TE documentation.

QUESTION NO: 36

Which of the following statements is NOT a valid reason to use MP-BGP as a routing protocol to transport VPRN routes?

A. The number of VPRN routes can become very large. MP-BGP can support a very large number of routes.

B. MP-BGP can carry routing information for a number of different address families.

C. MP-BGP is designed to exchange information between routers that are not directly connected. This feature keeps VPRN routing information out of P-routers.

D. MP-BGP can carry additional information attached to a route as an optional BGP attribute.

E. All of the above statements are valid reasons for using MP-BGP as a VPRN routing protocol.

ANSWER: E

Explanation:

All of the above statements are valid reasons for using MP-BGP as a VPRN routing protocol. MP-BGP is well suited to VPRN route distribution because its VPN address-family extensions scale to large numbers of customer routes while maintaining separation between VPN routing domains. It also supports multiple network-layer address families, allowing a common BGP framework to distribute VPN-IPv4, VPN-IPv6, and other relevant families where deployed.

Its ability to establish sessions between non-directly connected provider-edge routers is especially important in an MPLS VPN design: VPN control-plane information can be exchanged directly between PE routers without requiring provider core routers to participate in customer VPN routing. BGP path attributes also provide the extensible policy information needed for VPN operation. In particular, route-target extended communities identify VPN route import and export policy, enabling a receiving PE router to determine the appropriate VPRN routing-table placement for a received route. These capabilities collectively provide the scale, reachability model, multiprotocol support, and policy signaling required for VPRN service deployment. The MP-BGP extensions are specified in RFC 4760, while BGP/MPLS IP VPN route distribution and route-target communities are defined in RFC 4364.

QUESTION NO: 37

Which of the following devices receives labeled packets from the MPLS domain, removes the MPLS header, and forwards unlabeled packets outside the MPLS domain?

A. P router

B. LSR

C. eLER

D. iLER

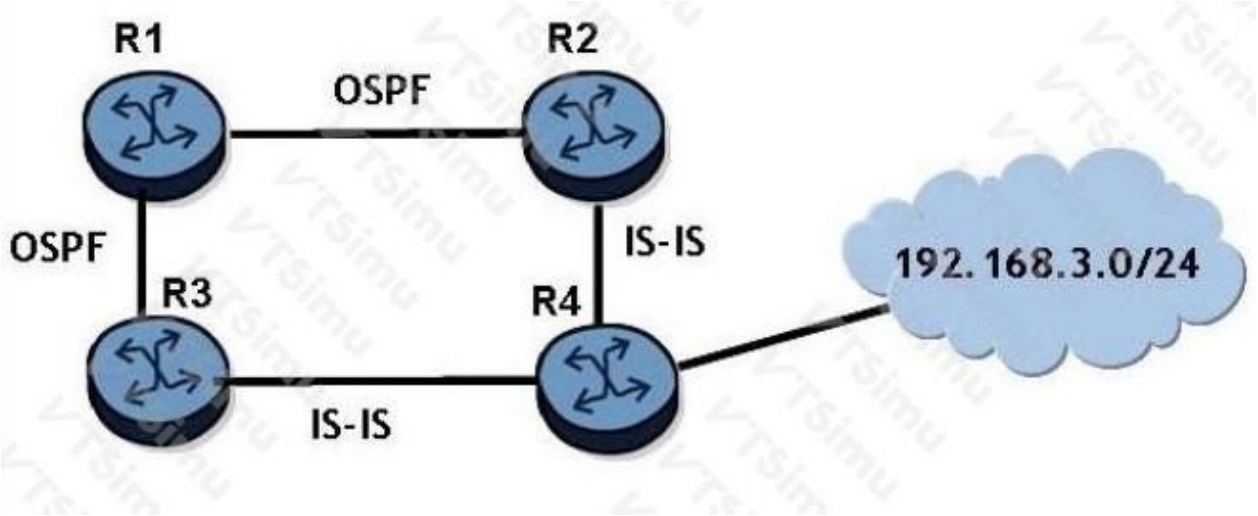
ANSWER: C

Explanation:

eLER is correct because an egress Label Edge Router sits at the point where traffic leaves an MPLS domain. Packets arriving at this router carry one or more MPLS label-stack entries that were imposed at MPLS ingress and used for forwarding through the provider backbone. The eLER uses the top label to identify the appropriate forwarding treatment, removes the MPLS label stack (the MPLS “pop” operation), and then forwards the resulting native IP or other payload packet toward the external, non-MPLS destination network. This operation allows the MPLS core to use label switching internally while presenting ordinary, unlabeled traffic at the network edge. In MPLS terminology, a Label Edge Router performs label imposition when traffic enters the MPLS domain and label disposition when traffic exits it; the egress role specifically performs disposition. The MPLS Architecture RFC describes this distinction between ingress and egress LER functions and the forwarding of packets after label-stack processing. See RFC 3031: MPLS Architecture and the MPLS forwarding terminology in RFC 3032: MPLS Label Stack Encoding.

QUESTION NO: 38

Click the exhibit button.



If router R2 redistributes the IS-IS route to 192.168.3.0/24 into OSPF, router R3 will receive two routes to 192.168.3.0/24. What will be the preference of these two routes? Assume that all IS-IS routers are L1/L2 capable and are in the same area. Choose two answers.

A. The OSPF internal preference

B. The IS-IS Level 1 internal preference

C. The IS-IS Level 2 internal preference

D. The OSPF external preference

E. The IS-IS Level 2 external preference

F. The IS-IS Level 1 external preference

ANSWER: B D

Explanation:

Router R3 learns 192.168.3.0/24 through its native IS-IS domain and also through OSPF after R2 redistributes that IS-IS route. Because all IS-IS routers are Level 1/Level 2 capable and belong to the same IS-IS area, the native IS-IS advertisement for the destination is an intra-area Level 1 route. Its route category is therefore **The IS-IS Level 1 internal preference**. The second route is created when R2 imports the IS-IS route into OSPF. Redistribution produces an Autonomous System External (ASE) OSPF route rather than an OSPF intra-area or inter-area route. Thus, on R3, that route has **The OSPF external preference**. Route preference identifies the administrative ranking assigned by SR OS to a route source/type before route metrics are considered among routes of the same preference. The router can consequently distinguish the directly learned IS-IS internal route from the redistributed OSPF external route when determining the active route. Nokia's SR OS routing documentation describes route preference and its role in route selection: SR OS Route Preference. Nokia's OSPF configuration documentation also covers external routes generated by redistribution: SR OS OSPF.

QUESTION NO: 39

Which of the following are functions of the MPLS control plane in a router? (Choose two)

A. Label allocation and management.

B. Label swapping.

C. Pushing a label on an unlabeled packet.

D. Label signaling.

E. Routing lookups.

ANSWER: A D

Explanation:

Label allocation and management is a control-plane function because the router must create, retain, advertise, withdraw, and associate MPLS labels with forwarding equivalence classes and forwarding information. This work establishes the label bindings that the forwarding plane will later use to handle packets at high speed. The MPLS architecture separates this label-binding information and forwarding state from the actual per-packet forwarding operation. RFC 3031 describes the Label Distribution Protocol role as distributing label bindings and the Label Information Base as maintaining label-related information used to construct forwarding behavior.

Label signaling is also a control-plane function. In MPLS networks, signaling protocols such as LDP or RSVP-TE communicate label bindings between label-switching routers. This permits routers to establish label-switched paths and install the corresponding label forwarding state. In Nokia SR OS deployments, MPLS control protocols are configured under the MPLS context and provide the signaling needed for MPLS services and transport paths. The resulting label information is programmed into forwarding tables, enabling the router's forwarding hardware to process traffic according to the established MPLS path. See RFC 3031: Multiprotocol Label Switching Architecture and Nokia's SR OS MPLS overview documentation.

QUESTION NO: 40

Click the exhibit button.

Router 2 advertises the network 192.168.3.0/24 to router R1 via IS-IS.

How can router R1 ensure that it discards the route?

A. Use a prefix list on router R1 to discard all ingress IS-IS routes.

B. Use an export policy and prefix list on router R1.

C. Use an import policy and a prefix list on router R1.

D. This is not possible because IS-IS cannot use import policies.

ANSWER: C

Explanation:

Use an import policy and a prefix list on router R1. In SR OS, an IS-IS import policy is evaluated for routes learned by IS-IS before they are installed in the routing table. R1 can define a prefix list that matches `192.168.3.0/24`, reference that prefix list in a policy statement, and specify a reject action for the matching route. Applying that policy as the IS-IS import policy prevents the advertised prefix from being accepted into R1's route table while allowing the router to continue processing other IS-IS information normally. The policy should include an explicit action for all nonmatching prefixes, usually acceptance, so that only the intended destination is filtered. This is an inbound routing-policy function: the decision is made when R1 receives and evaluates IS-IS-learned routes, rather than when R1 advertises routes to another router. Nokia SR OS routing policies use prefix lists as match criteria and support accept or reject actions to control route installation and advertisement. See Nokia's Route Policies documentation and the IS-IS documentation.

QUESTION NO: 41

What must be configured on the Alcatel-Lucent 7750 SR for RIP to advertise locally attached links?

A. A policy statement within the router identifying what is to be advertised

B. Access-lists denoting what is to be advertised

C. Default metrics for each link

D. A group and neighbor statement for each peer

ANSWER: A

Explanation:

A policy statement within the router identifying what is to be advertised is correct. On SR OS, RIP does not automatically advertise every locally connected route merely because RIP is enabled on an interface. Routes to be injected into RIP, including locally attached (direct) prefixes, are controlled through an export routing policy applied under the RIP routing context. The policy contains match criteria—for example, matching routes with the direct protocol—and an action that accepts those routes for export into RIP. This explicit policy-based export model lets the router advertise only intended connected networks and apply attributes such as a RIP metric when required. A policy statement therefore provides both the route-selection mechanism and the controlled redistribution point needed for local links to become RIP advertisements. Nokia's SR OS documentation describes RIP route export as policy controlled and documents routing-policy match and action processing in the routing protocol configuration material. See the Nokia SR OS documentation portal and the Nokia 7750 SR OS Infocenter for the RIP and routing-policy configuration references.

QUESTION NO: 42

Which of the following statements describe the major features of OSPF? Choose two answers.

A. Fast reroute capability

B. Control traffic prioritization

C. Route redistribution

D. Traffic engineering extensions

E. Cut through forwarding

ANSWER: C D

Explanation:

Route redistribution is a major OSPF capability because OSPF can import routes learned from other routing sources, such as static routes or other routing protocols, into an OSPF domain. These imported routes are advertised as AS-external routes, allowing OSPF routers to calculate reachability to destinations outside the OSPF autonomous system while retaining the protocol's link-state operation within the domain. This is a fundamental function for connecting OSPF networks to external routing domains.

Traffic engineering extensions are also a major OSPF feature. OSPF traffic engineering uses extensions that advertise additional link attributes, including available bandwidth and administrative characteristics, through opaque LSAs. MPLS traffic-engineering applications can use this topology and resource information to compute constraint-based paths rather than relying solely on the normal shortest-path calculation. The base OSPF specification defines the external-route mechanism, while the OSPF traffic-engineering extension defines the TE link-state advertisements and attributes used by TE-capable networks. See RFC 2328: OSPF Version 2 and RFC 3630: Traffic Engineering Extensions to OSPF Version 2.

QUESTION NO: 43

Click on the exhibit.

```
*A:R2# show router ospf database

===============================================================================
OSPF Link State Database (Type : All)
===============================================================================
Type    Area Id         Link State Id   Adv Rtr Id      Age  Sequence   Cksum
-------------------------------------------------------------------------------
Router  0.0.0.0         10.10.10.2      10.10.10.2      21   0x8000001a 0xae7d
Router  0.0.0.0         10.10.10.3      10.10.10.3      42   0x80000017 0x7282
Router  0.0.0.0         10.10.10.4      10.10.10.4      39   0x80000019 0x25ab
Router  0.0.0.0         10.10.10.5      10.10.10.5      13   0x80000017 0xc1ed
Router  0.0.0.0         10.10.10.6      10.10.10.6      18   0x80000022 0xa812
Network 0.0.0.0         10.2.4.2        10.10.10.2      52   0x80000001 0x5095
Network 0.0.0.0         10.2.5.2        10.10.10.2      21   0x80000001 0x5390
Network 0.0.0.0         1.1.1.1         10.10.10.3      1920 0x80000002 0x3991
Network 0.0.0.0         10.3.4.3        10.10.10.3      43   0x80000001 0x3ea3
Network 0.0.0.0         10.5.6.6        10.10.10.6      18   0x80000001 0xcc7
-------------------------------------------------------------------------------
No. of LSAs: 10
===============================================================================
```

Given the output, which of the following is the most accurate statement about the network?

A. There are five broadcast links and there are no point-to-point links.

B. There are five broadcast links and there are five point-to-point links.

C. There are five broadcast links. It's not possible to say how many point-to-point links there are.

D. It's not possible to say how many broadcast or point-to-point links there are.

ANSWER: C

Explanation:

There are five broadcast links. It's not possible to say how many point-to-point links there are. OSPF represents each transit broadcast or non-broadcast multi-access network with a Type 2 Network LSA, originated by that network's designated router. Thus, the five Network LSAs shown in the output identify five multi-access transit segments, commonly described in this context as broadcast links. A point-to-point OSPF interface does not originate or require a Type 2 Network LSA; its connectivity is instead described through the routers' Type 1 Router LSAs. Consequently, the displayed Network-LSA count establishes the number of broadcast transit networks, but it cannot be used to derive a count of point-to-point links. Determining the latter would require examining the Router LSAs or relevant OSPF interface and neighbor information, including each interface's configured network type. This behavior follows the OSPF LSA model defined in RFC 2328, which specifies that Network LSAs describe transit networks and are generated by the designated router. Nokia's OSPF implementation and operational model are documented in the Nokia SR OS OSPF documentation.

QUESTION NO: 44

Which of the following statements regarding OSPF route types are true? Choose two answers.

A. Intra-area routes are preferred over inter-area routes.

B. Inter-area routes are preferred over external routes.

C. External routes are preferred over intra-area routes.

D. Type 3 LSAs advertise external routes originated by an ASBR.

E. All OSPF route types have equal preference.

ANSWER: A B

Explanation:

OSPF route selection prefers intra-area routes over inter-area routes, and it prefers inter-area routes over external routes. An intra-area route is learned from the topology database of the local area. An inter-area route is learned through an ABR using Type 3 summary information. External routes are introduced by an ASBR and are represented by Type 5 LSAs in normal areas or Type 7 LSAs within an NSSA.

These route-type preferences are part of the OSPF path-selection process defined in RFC 2328, section 16.4.

QUESTION NO: 45

In an OSPF Hello packet what fields must match all neighbor routers on the segment? (Choose 3)

A. Area ID

B. Hello and Dead Intervals

C. Stub flag

D. DR and BDR addresses

E. The list of neighbors

ANSWER: A B C

Explanation:

OSPF routers can establish and maintain neighbor adjacencies only when key Hello parameters are compatible. **Area ID** must be identical because an OSPF interface belongs to one area, and Hellos received for a different area are not accepted as belonging to the same OSPF domain segment. **Hello and Dead Intervals** must also agree: the Hello interval controls the expected rate of received Hellos, while the Dead interval defines how long a router can remain silent before the neighbor is declared down. Matching timers ensures both routers use the same liveness expectations. The **Stub flag**, carried as the external-routing-capability setting in the OSPF Options field, must match as well. This setting identifies whether the attached area supports external LSAs; routers on a shared segment must therefore agree on the area's stub capability before becoming neighbors. These checks protect OSPF area consistency and prevent routers with incompatible operational assumptions from forming an adjacency. RFC 2328 specifies the Hello packet format and the required neighbor-compatibility checks, including Area ID, timers, and the stub-area capability bit: RFC 2328: OSPF Version 2.

QUESTION NO: 46

How have Traffic Engineering extensions to IS-IS been implemented?

A. By defining new LSAs for the IS-IS protocol.

B. By utilizing the OSPF opaque LSAs to maintain interoperability with OSPF routing domains

C. By defining new TE extension TLVs to the IS-IS protocol.

D. By enabling CSPF algorithm extensions that allow link state TE to be implemented.

ANSWER: C

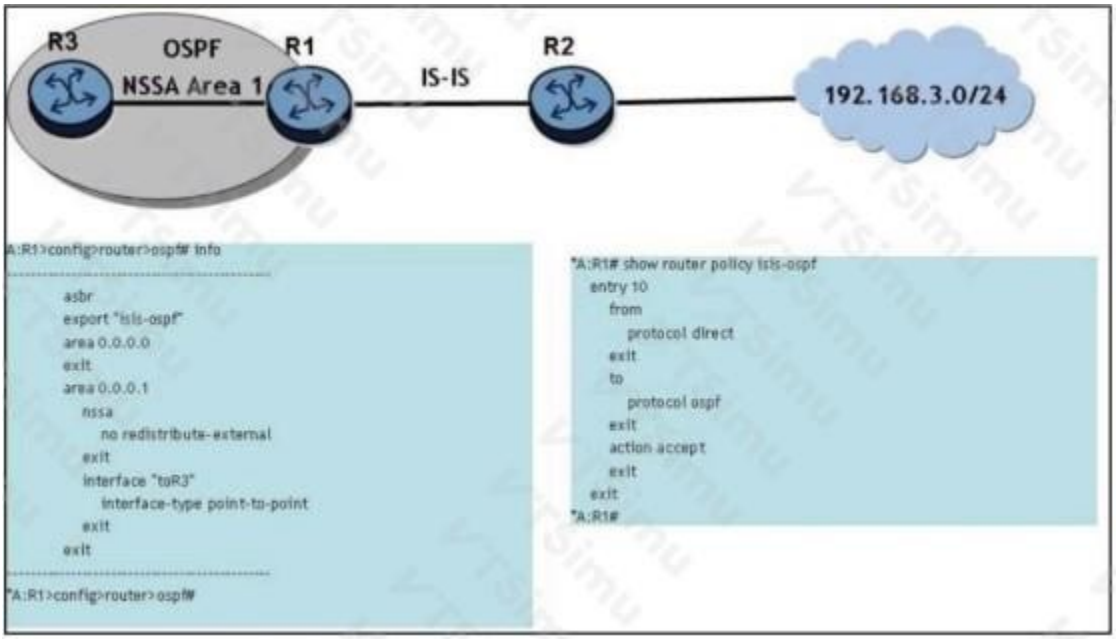
Explanation:

By defining new TE extension TLVs to the IS-IS protocol is correct. IS-IS is designed to carry extensible link-state information using type-length-value (TLV) elements within its Link State Protocol Data Units. Traffic Engineering extends this existing mechanism by defining additional TLVs and sub-TLVs that advertise TE-specific link attributes. These attributes form the Traffic Engineering Database used for constrained path computation, including information such as administrative groups, available bandwidth at each priority, and TE metrics. Because the TE information is distributed with the IS-IS link-state database, every participating router can build a consistent topology and resource view for MPLS-TE path calculation. RFC 5305 specifically defines the IS-IS extensions for traffic engineering and describes the Extended IS Reachability TLV

together with TE-related sub-TLVs. This approach preserves the fundamental IS-IS flooding and shortest-path architecture while supplying the additional constraints needed by CSPF and RSVP-TE signaling. See RFC 5305, IS-IS Extensions for Traffic Engineering and the IETF RFC 5305 text.

QUESTION NO: 47

Click the exhibit button



The 192 168 3.0/24 network is learned on router R1 via IS-IS Given the OSPF configuration shown, and assuming that the OSPF adjacency between routers R1 and R3 is up, why is the 192 168.3.0/24 route not in router R3's route table? (Choose two)

A. The route policy should be applied as in import policy

B. The no-redistribute-external command is used on router R1

C. The route policy is incorrect. It should say "from protocol ISIS" rather than "from protocol direct".

D. It is not possible to be an ASBR and an NSSA. The ASBR configuration should be removed.

E. The interface between routers R1 and R3 needs to be in OSPF area 0 rather than OSPF area 1 because a backbone area must always exist.

ANSWER: B C

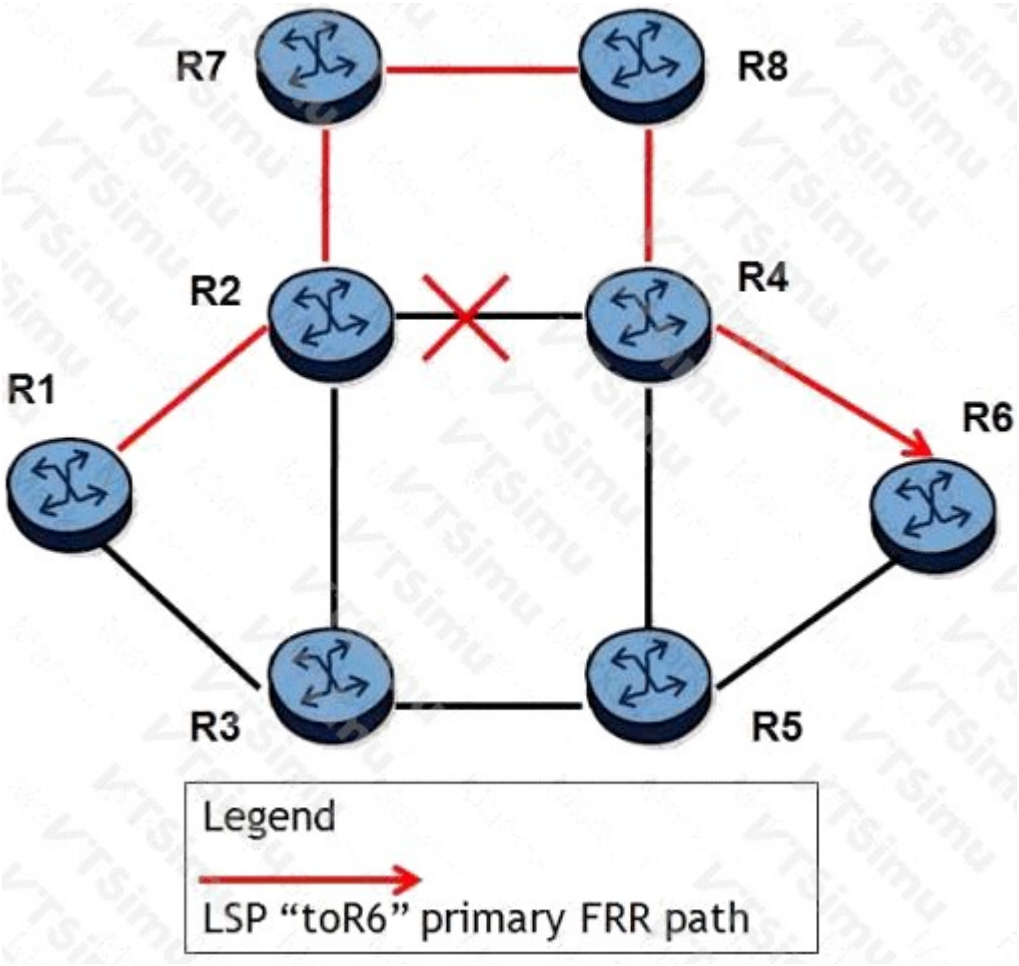
Explanation:

The route must first be selected by the OSPF export policy on R1 before it can be redistributed from IS-IS into OSPF. The policy shown matches routes learned from the direct protocol, whereas 192.168.3.0/24 was learned through IS-IS. Therefore, changing the match condition to `from protocol ISIS` allows that IS-IS route to be accepted by the export policy and considered for OSPF external redistribution. Nokia SR OS routing policies use protocol matching to control which source routes are eligible for export into another routing protocol.

In addition, the `no-redistribute-external` setting configured for the NSSA prevents external routes from being redistributed into that NSSA. Even once the IS-IS route is correctly matched by the export policy, this NSSA control must permit external redistribution for R3 to receive and install the redistributed prefix. With both conditions corrected, R1 operates as the ASBR, originates the appropriate NSSA external route, and advertises reachability for 192.168.3.0/24 toward R3 through the established OSPF adjacency. Nokia's SR OS documentation covers OSPF route export and NSSA external-route behavior in the OSPF configuration guide and routing-policy processing in the routing policy guide.

QUESTION NO: 48

Click on the exhibit.



A fully loose LSP " toR6 " is enabled with FRR protection. All links have the same cost. After the link between R2 and R4 goes down, FRR protection repairs the LSP and traffic traverses on R1-R2-R7-R8-R4-R6. Which of the following about this CSPF-enabled LSP on an Alcatel-Lucent 7750 SR is TRUE?

A. The LSP will stay on the current path indefinitely unless another failure occurs.

B. The LSP will stay on the current path until the resignal timer expires.

C. The LSP will switchover to R1-R2-R3-R5-R6 after the resignal timer expires.

D. The LSP will switchover to R1-R3-R5-R6 after the retry timer expires.

ANSWER: D

Explanation:

The LSP will switchover to R1-R3-R5-R6 after the retry timer expires. Fast reroute provides immediate local protection at the point of failure, so traffic can continue across the pre-established repair path while the control plane processes the topology change. Because the LSP is fully loose, its ingress router is not constrained to preserve the former hop sequence when it subsequently establishes a replacement path. CSPF can therefore calculate a new end-to-end route using the updated traffic-engineering database, which excludes the failed R2-to-R4 adjacency. With equal link costs, R1-R3-R5-R6 is the shortest available CSPF route to R6 shown in the topology. The retry timer governs when the ingress retries LSP establishment after the affected path condition, allowing the replacement RSVP-TE LSP to be signaled using that newly computed route. Once the replacement LSP is successfully established, traffic moves from the temporary FRR repair path to R1-R3-R5-R6. Nokia SR OS MPLS documentation describes CSPF-based LSP path calculation and LSP signaling behavior in its SR OS documentation portal. RSVP-TE local-repair behavior is also standardized in RFC 4090.

QUESTION NO: 49

Which of the following statements best describes the flooding of traffic on a PE when traffic is received on a mesh SDP?

A. Traffic is flooded to all SAPs in the service.

B. Traffic is flooded to all SAPs and spoke SDPs in the service.

C. Traffic is flooded to all SAPs and mesh SDPs in the service.

D. Traffic is flooded to all SAPs, spoke SDPs. and mesh SDPs in the service.

E. The traffic is not flooded.

ANSWER: B

Explanation:

Traffic is flooded to all SAPs and spoke SDPs in the service. In a Nokia VPLS service, a mesh SDP represents a network-facing connection to another PE in the VPLS mesh. When the PE receives unknown-unicast, broadcast, or applicable multicast traffic through that mesh SDP, it replicates the frame toward locally attached customer-facing SAPs and toward spoke SDPs. A spoke SDP typically connects the VPLS to a remote edge, access, or hub-and-spoke portion of the service, so it is an eligible flooding destination when the frame arrives from the mesh side.

This behavior implements split horizon for the full-mesh PE interconnection. The receiving mesh SDP is part of the VPLS mesh domain, and traffic received from that domain must not be replicated onto other mesh SDPs. Preventing mesh-to-mesh flooding avoids loops and duplicate copies in the MPLS core while still delivering flooded traffic to customer-facing and spoke-connected portions of the same service. This is the standard VPLS split-horizon model described in Nokia SR OS Layer 2 service documentation and in the VPLS specification. See the Nokia SR OS documentation portal and RFC 4762, Section 10.

QUESTION NO: 50

Which of the following options are used to configure an MPLS signaling type to be used on an SDP? (Choose 2)

A. configure service sdp sdp-id gre create

B. configure service sdp sdp-id create

C. ldp

D. lsp lsp-name

ANSWER: C D

Explanation:

For an MPLS SDP, SR OS supports selecting either LDP signaling or a specifically named RSVP-TE LSP as the transport-signaling method. The `ldp` command configures the SDP to use Label Distribution Protocol, allowing the SDP to use an LDP-signaled transport tunnel to its far-end address. This is appropriate where the MPLS transport network distributes labels through LDP.

The `lsp lsp-name` command associates the SDP with a configured MPLS LSP by name. This selects an explicitly provisioned RSVP-TE LSP for the SDP, which is commonly used when traffic-engineered paths, constraints, or reserved bandwidth are required. These settings are made within the MPLS context of the SDP after the SDP itself has been created. Nokia's SR OS service configuration documentation describes SDPs as the transport binding used by services and documents MPLS SDP configuration, while the CLI reference provides the command hierarchy and syntax for SDP MPLS signaling parameters. See the Nokia SR OS Services Overview and the Nokia SR OS CLI Reference.

QUESTION NO: 51

Click on the exhibit.



Which of the following best describes the format of data traffic sent to the server as it exits from the PC?

A. The destination IP address will be the system IP address of router R1 and the destination MAC address will be the MAC address of the router R1 port connected to the PC.

B. The destination IP address will be the IP address of the router R1 port connected to the PC and the destination MAC address will be the MAC address of the router R1 port connected to the PC.

C. The destination IP address will be the IP address of the server and the destination MAC address will be the MAC address of the server.

D. The destination IP address will be the IP address of the server and the destination MAC address will be the MAC address of the router R1 port connected to the PC.

ANSWER: D

Explanation:

The destination IP address will be the IP address of the server and the destination MAC address will be the MAC address of the router R1 port connected to the PC. IP addressing is end-to-end: when the PC creates an IP packet for a server located on a different IP subnet, the packet retains the server's IP address as its destination throughout the routed path. The PC uses its routing decision, normally its configured default gateway, to determine that R1 is the next hop for traffic to the server's subnet.

Before transmitting the packet on its local Ethernet segment, the PC encapsulates it in an Ethernet frame. Ethernet delivery is local to that LAN, so the frame must be addressed to the MAC address of the next-hop router interface connected to the PC. The PC learns that MAC address through ARP if it is not already present in its ARP cache. R1 removes the incoming Layer 2 frame, routes the IP packet using the server destination address, and applies a new Layer 2 encapsulation for the next link. This distinction between persistent Layer 3 destination addressing and per-hop Layer 2 addressing is a fundamental routed-network behavior. See RFC 1812: Requirements for IP Routers and RFC 826: Address Resolution Protocol.

QUESTION NO: 52

What can the sdp-ping command be used for? (Choose 2)

A. To test the ability of reaching the far-end IP address of an SDP within the SDP encapsulation

B. To determine the path MTU to the far-end IP address over an SDP.

C. To determine the service MTU of the service using an SDP.

D. To determine the presence of hosts using the service

ANSWER: A B

Explanation:

The `sdp-ping` command is an MPLS SDP OAM tool that tests forwarding across a configured service distribution path. It sends test packets toward the SDP far-end while applying the SDP/MPLS encapsulation, allowing an operator to validate that the remote SDP endpoint is reachable through the transport tunnel rather than merely verifying ordinary IP reachability in the base routing table. This makes it useful when troubleshooting an SDP that is administratively up but may have an underlying MPLS or transport-path forwarding issue.

The command can also perform path-MTU discovery toward the far-end IP address across the SDP. By sending progressively sized probes and observing the result, it identifies the largest packet that can traverse the SDP path without fragmentation. This is valuable for ensuring that the transport network can carry the encapsulated service traffic and for diagnosing MTU mismatches between SDP endpoints or along the MPLS path. Nokia SR OS documentation describes SDP operational and diagnostic functions in its service-management material and command references; see the Nokia SR OS Services Overview and the SR OS MD-CLI Command Reference.

QUESTION NO: 53

If the Alcatel-Lucent 7750 SR is rebooted after an "admin save" has been performed, which of the following about mirror services is TRUE?

A. Only the mirror source must be re-configured.

B. Only the mirror destination must be re-configured.

C. Both the mirror destination and mirror source must be re-configured.

D. Neither the mirror destination nor the mirror source must be re-configured.

ANSWER: A

Explanation:

Only the mirror source must be re-configured. On the 7750 SR, an `admin save` preserves the configured mirror destination, so its definition is available again when the router loads its saved configuration during startup. The mirror destination supplies the persistent egress or capture point used by the mirroring function. Mirror-source configuration, however, is intentionally not retained by the saved configuration in this behavior. It must therefore be recreated after the reboot before traffic can again be copied to the destination. This distinction is important operationally: after planned maintenance or an unexpected restart, engineers should verify that required mirror sources have been restored before relying on packet captures, lawful-intercept workflows, troubleshooting traces, or traffic analysis. The saved destination can be reused when recreating the source, which avoids having to rebuild the full mirror-service endpoint configuration. Nokia's Mirror Services documentation describes mirror-service configuration and operational behavior for SR OS platforms; see the Nokia SR OS Mirror Services Guide. The wider SR OS documentation library is available from Nokia Documentation Center.

QUESTION NO: 54

Why does IPv6 not use ARP?

A. The ARP functionality has been replaced by the Neighbor Discovery protocol.

B. The ARP functionality is not needed due to the increased IPv6 address space.

C. ICMPv6 reduced its role to connectivity checking and auto-configuration only.

D. ICMFV6 replaced the ARP functionality with Multicast Listener Discovery (MLD).

ANSWER: A

Explanation:

The ARP functionality has been replaced by the Neighbor Discovery protocol. IPv6 Neighbor Discovery (ND) performs the Layer 3-to-Layer 2 address-resolution function that ARP provides in IPv4, but it does so as part of ICMPv6. When an IPv6 node needs to send traffic to a neighbor on the local link, it sends an ICMPv6 Neighbor Solicitation message, normally directed to the target's solicited-node multicast address. The target responds with a Neighbor Advertisement containing its link-layer address. The sender can then create or update its neighbor-cache entry and encapsulate traffic in the appropriate Layer 2 frame.

ND is broader than simple address resolution. It also supports neighbor reachability detection, router and prefix discovery, redirect handling, duplicate-address detection, and certain address-configuration processes. Integrating these related local-link functions into ICMPv6 gives IPv6 a unified neighbor-management mechanism rather than retaining the separate ARP protocol used by IPv4. The IETF defines the ND protocol and its Neighbor Solicitation/Advertisement behavior in RFC 4861, Neighbor Discovery for IP version 6. The ICMPv6 message framework used by ND is specified in RFC 4443, Internet Control Message Protocol for IPv6.

QUESTION NO: 55

Which of the following commands is used to enable an E-pipe service?

A. config>service= epipe enable

B. config>service~ vpws enable

C. config>service= no shutdown

D. config>port= no shutdown

E. config>service= epipe no shutdown

ANSWER: E

Explanation:

The command `config>service= epipe <service-id> no shutdown` is used to administratively enable an Epipe service. In Nokia SR OS, a service has an administrative state, and Epipe services are created in the shutdown state unless they are explicitly brought up. After entering the context of the required Epipe service, the `no shutdown` command changes that service's administrative state to enabled, allowing it to become operational when its required service configuration and attachment circuits are also available. The service ID identifies the particular Epipe instance to activate, so the command applies only to that specified point-to-point Ethernet service. This reflects the standard SR OS configuration workflow: create and configure the service, then issue `no shutdown` in the service context to place it into service. Nokia's Epipe documentation describes Epipe as a point-to-point Ethernet service and documents its service-level administrative controls. See the Nokia SR OS Epipe services documentation and the Nokia SR OS CLI reference.

QUESTION NO: 56

Click on the exhibit button below.

Consider the output shown; what table is displayed?

A. RIB

B. FIB

C. LIB

D. LFIB

ANSWER: D

Explanation:

The displayed table is the **LFIB** (Label Forwarding Information Base). In an MPLS-capable Nokia SR OS router, the LFIB contains the forwarding-plane label entries used to process labeled packets. Each entry associates an incoming MPLS label with the forwarding action required by the router, such as swapping to an outgoing label, popping a label, or pushing labels before transmission toward a next hop or service. This is why output showing label values together with MPLS forwarding operations, outgoing labels, next hops, or interfaces identifies the LFIB rather than a conventional IP route table. The LFIB is derived from MPLS control-plane information and is the table consulted for label-based packet forwarding. Nokia's MPLS documentation describes the SR OS MPLS forwarding and label-management functions: Nokia SR OS Documentation. The MPLS architecture and the role of label forwarding are also defined in RFC 3031.

QUESTION NO: 57

Which of the following encapsulation types on an Ethernet port support the provisioning of multiple SAPs? (Choose 2)

A. Dot1Q

B. Q-in-Q

C. ATM

D. NULL

E. Nil

ANSWER: A B

Explanation:

Dot1Q and Q-in-Q support multiple SAPs on an Ethernet port because they provide VLAN tag values that the router can use to distinguish separate service attachment points. With Dot1Q encapsulation, SAPs can be provisioned against distinct customer VLAN IDs on the same physical Ethernet port. Each VLAN-tagged traffic stream is therefore mapped to its own SAP and may be associated with the applicable service configuration. Q-in-Q extends this model by using stacked VLAN tags, allowing SAP identification using customer and service VLAN tag information. This enables a provider to multiplex a greater number of logically separate customer connections over a shared Ethernet interface, while preserving the required service separation. Nokia SR OS service provisioning identifies Ethernet SAPs using encapsulation values and supports VLAN-based SAP multiplexing where the port encapsulation permits tagged Ethernet frames. These capabilities are fundamental to delivery of Ethernet services such as Epipe and VPLS over shared access-facing or network-facing Ethernet links. Further details on SAPs and Ethernet encapsulation are available in the Nokia SR OS SAP documentation and the Nokia Ethernet services overview.

QUESTION NO: 58

Which of the following about dot1Q SAPs is FALSE?

A. The dot1Q default SAP (port: *) can be used to pass VLAN tags transparently through a network.

B. The dot1Q null SAP (port: 0) accepts all tagged frames and untagged frames.

C. The dot1Q default SAP (port: *) accepts all untagged frames and any frames with tag values that are not used as a service-delimiting value on another SAP.

D. A dot1q null SAP (port: 0) and a dot1q default SAP (port: *) cannot both be defined on the same port.

ANSWER: B

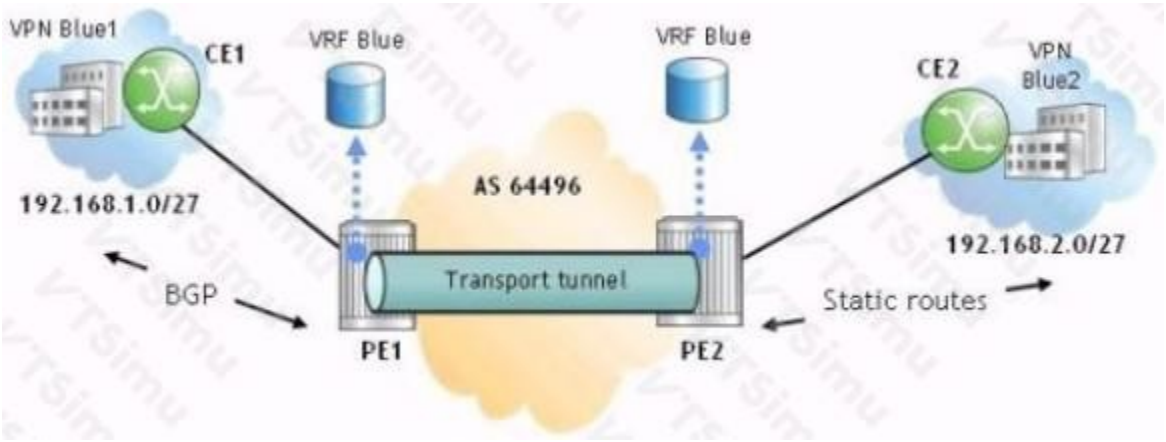
Explanation:

The statement "The dot1Q null SAP (port: 0) accepts all tagged frames and untagged frames." is false because a dot1Q null SAP is specifically the service access point for *untagged* Ethernet traffic on a port. In SR OS SAP syntax, the value `0` represents the absence of an IEEE 802.1Q VLAN tag; consequently, frames received without a VLAN header match this SAP. Tagged frames carry a VLAN ID and are classified using a matching VLAN-specific SAP or, where configured, a

default SAP that handles otherwise-unmatched VLAN-tagged traffic. This distinction is fundamental to SR OS service demultiplexing: the encapsulation identifier used in the SAP determines the incoming frame type that can be associated with the service. A null SAP therefore provides a clear attachment point for native or untagged customer traffic, rather than acting as a catch-all attachment for both tagged and untagged traffic. Nokia's Services Overview documentation describes SAP encapsulation and the handling of dot1Q SAP identifiers; the SR OS documentation library is available at Nokia SR OS Documentation, including the Service Access Points documentation.

QUESTION NO: 59

Click the exhibit.



Routes are exchanged between CE1 and PE1 using BGP. CE2 and PE2 use static routes to forward traffic to and from the VPRN. Which of the following is FALSE?

A. An export policy is required on PE1 to advertise routes to CE1.

B. An export policy is required on CE1 to advertise local routes to PEL

C. An export policy is required to advertise routes between the PE routers.

D. No export policy is required on PE2 or CE2 due to the static routes.

ANSWER: C

Explanation:

"An export policy is required to advertise routes between the PE routers." is the false statement. In an SR OS VPRN, PE-to-PE distribution is performed by MP-BGP using VPN route families. After a route is installed in the applicable VPRN routing context, the PE creates and distributes the VPN route with the service route distinguisher and route-target information. Remote PEs use the route-target association to identify routes belonging to the VPRN and import them into the appropriate service. This PE-to-PE VPN route exchange is a fundamental VPRN control-plane function and does not inherently require a BGP export policy merely to make the exchange occur. An administrator can configure VRF import/export policies to selectively control which service routes are advertised or accepted, but such filtering is optional rather than a base requirement for PE-to-PE route advertisement. The use of static routing at a CE-PE attachment does not change the PE-to-PE signaling mechanism; the PE still distributes eligible VPRN routes through MP-BGP. Nokia's SR OS documentation describes VPRN service routing and VPN route distribution at the Nokia SR OS documentation portal; the underlying VPN route model is also defined in RFC 4364.

QUESTION NO: 60

Which of the following statements best describes an OSPF type 3 LSA?

A. An OSPF type 3 LSA is an intra-area LSA, generated by the ABR only when route summarization is configured.

B. An OSPF type 3 LSA is an inter-area LSA, generated by the ABR only when route summarization is configured.

C. An OSPF type 3 LSA is an inter-area LSA, generated by an ASBR to describe external routes.

D. An OSPF type 3 LSA is an inter-area LSA, generated by the ABR to describe a route to neighbors outside the area.

ANSWER: D

Explanation:

An OSPF type 3 LSA is an inter-area LSA, generated by the ABR to describe a route to neighbors outside the area. Type 3 LSAs are commonly called summary-LSAs. An Area Border Router originates them to advertise reachability information learned from one OSPF area into another area. For example, when an ABR has routes for networks within one attached area, it originates type 3 LSAs into its other attached areas so routers there can calculate inter-area paths toward those destinations. The LSA identifies the destination network and mask and supplies the metric used for inter-area route calculation; it does not carry the full topology of the source area. This behavior preserves OSPF's hierarchical design, in which detailed link-state information is maintained within an area while other areas receive summarized reachability information. Route summarization can affect how prefixes are represented in type 3 advertisements, but type 3 LSAs are not limited to deployments where manual summarization is configured. RFC 2328 specifies that area-border routers originate summary-LSAs to describe destinations outside the receiving area and defines their use during inter-area route calculation. See RFC 2328, Summary-LSAs and RFC 2328, Summary-LSA origination.

QUESTION NO: 61

In a VPLS service, the mesh-sdp vc-id is identical to the service id by default.

A. TRUE

B. FALSE

ANSWER: A

Explanation:

TRUE is correct. In Nokia SR OS, a VPLS service uses mesh-SDP bindings to establish the full mesh of pseudowires toward remote provider-edge routers. A mesh-SDP binding is identified with an SDP ID and a virtual-circuit ID (VC ID). When the VC ID is not explicitly configured in the mesh-SDP statement, SR OS derives its default value from the local VPLS service ID. Consequently, a VPLS configured with service ID 100 and a mesh-SDP for which no VC ID is supplied uses VC ID 100 by default. This default behavior simplifies configuration because the service identifier can be used consistently as the pseudowire identifier across the VPLS mesh. The corresponding remote endpoint must use the matching VC ID for the pseudowire to be associated correctly. Operators may explicitly configure a different VC ID when service-numbering or interconnection requirements make that necessary; however, the stated default remains that the mesh-SDP VC ID equals the service ID. Nokia's service documentation describes VPLS mesh-SDP bindings and their VC-ID behavior in the SR OS VPLS documentation; related service configuration concepts are also covered in the SR OS Services Overview.

QUESTION NO: 62

Which of the following statements about VPRN is NOT true? (Choose 2)

A. VPRN is an IP service that connects multiple sites in a single routed domain over the provider-managed IP/MPLS network.

B. Each VPRN consists of a set of customer sites that must connect to the same PE router

C. In a VPRN service, each customer router becomes a routing peer to other customer routers

D. VPRN provides customer routers with transparent IP connectivity without knowledge of the MPLS core

ANSWER: B C

Explanation:

The statements “Each VPRN consists of a set of customer sites that must connect to the same PE router” and “In a VPRN service, each customer router becomes a routing peer to other customer routers” are not true. A VPRN is a Layer 3 VPN in which customer sites can attach to different provider edge routers throughout the provider IP/MPLS network. The service is specifically designed to extend a customer’s routed VPN across multiple provider edge devices while maintaining a separate VPN routing and forwarding context for that customer. There is no requirement for every site in the VPRN to be attached to one provider edge router.

Customer routers also do not form routing adjacencies directly with remote customer routers through the MPLS core. A customer edge router establishes its routing relationship with its locally attached provider edge router, using static routing or a PE-CE routing protocol such as BGP, OSPF, RIP, or IS-IS where supported and configured. The provider edge routers distribute VPN reachability across the backbone using VPN route mechanisms, allowing remote customer prefixes to be reachable without exposing or requiring customer routers to participate in the core’s MPLS control plane. This model is defined by the BGP/MPLS IP VPN architecture in RFC 4364 and is consistent with Nokia’s VPRN service documentation.

QUESTION NO: 63

Which of the following describes what happens when facility backup fast reroute is enabled for a LSP? Choose three answers.

A. Each PLR will signal a bypass tunnel that avoids the next-hop node or link.

B. When a failure occurs, the PLR will push an additional label onto the label stack and start forwarding traffic onto its bypass tunnel.

C. The ingress router will include a RRO object in the path message to indicate the routers that should be PLRs and MPs along the protected path.

D. The ingress router will set the local protection desired flag in the session attribute object of the path message to inform all routers along the protected LSP path to establish bypass tunnels.

E. When the bypass tunnel is in use, the MP will include the RRO object in its PathErr message, with the local protection in use flag set to notify the ingress router that it is using the bypass tunnel.

ANSWER: A B D

Explanation:

Facility backup fast reroute protects an RSVP-TE LSP locally at points along its path, allowing traffic to be redirected in tens of milliseconds without waiting for the ingress router to calculate and signal a replacement end-to-end LSP. Each potential point of local repair establishes an RSVP-TE bypass tunnel that avoids the protected next-hop resource: either the adjacent link or the next-hop router. The bypass is merged back into the original LSP path at a merge point downstream of the protected resource. When the protected resource fails, the point of local repair redirects affected traffic by adding the bypass forwarding label to the existing MPLS label stack. This preserves the original LSP label so that, after the bypass reaches the merge point, forwarding can continue over the original protected LSP. The ingress signals that local protection is desired by setting the Local Protection Desired flag in the Session Attribute object carried in the Path message. This requests local protection along the LSP path and enables routers capable of providing protection to associate the LSP with suitable bypass tunnels. These behaviors are defined for RSVP-TE fast reroute in RFC 4090, with RSVP-TE Path-message and Session Attribute processing specified in RFC 3209.

QUESTION NO: 64

If an Alcatel-Lucent 7750 LSR receives a labeled packet with a outer label TTL of 250 and an inner label TTL of 251, which of the following statements is true?

A. The packet will be label switched; the outer label TTL will be 249 and the inner label TTL will be 250.

B. The packet will be label switched; the outer label TTL will be 249 and the inner label TTL will be 251.

C. The packet will be label switched; both the outer label TTL and the inner label TTL will be 250.

D. The packet will be label switched; both the outer label TTL and the inner label TTL will be 249, because they must be the same.

ANSWER: B

Explanation:

The packet will be label switched; the outer label TTL will be 249 and the inner label TTL will be 251. MPLS TTL processing is applied to the label at the top of the received label stack—the label that is used by the LSR to make its forwarding decision. Before forwarding, the 7750 LSR decrements that top-label TTL by one. Thus, an incoming outer TTL of 250 becomes 249. The inner label is not exposed or processed during this label-swap operation, so its TTL remains independently encoded as 251. MPLS labels each carry their own 8-bit TTL field; the TTL values in adjacent stack entries do not need to match. This behavior provides loop protection for the forwarding path represented by each active label-stack level, while allowing a transport label and a service or VPN label to retain distinct TTL values. The MPLS label-stack encoding and top-label TTL forwarding behavior are specified in RFC 3032, MPLS Label Stack Encoding. Nokia's SR OS MPLS documentation is available through the Nokia SR OS Documentation Center.

QUESTION NO: 65

Which of the following statements regarding IS-IS point-to-point operation are true? Choose three answers.

A. Point-to-Point IS-to-IS Hello PDUs are exchanged on the circuit.

B. A DIS is elected on the circuit for each IS-IS level.

C. Point-to-point operation does not require a DIS election.

D. The three-way handshake can be used to verify bidirectional connectivity.

E. Only Level 1 LAN Hello PDUs can be used on a point-to-point circuit.

ANSWER: A C D

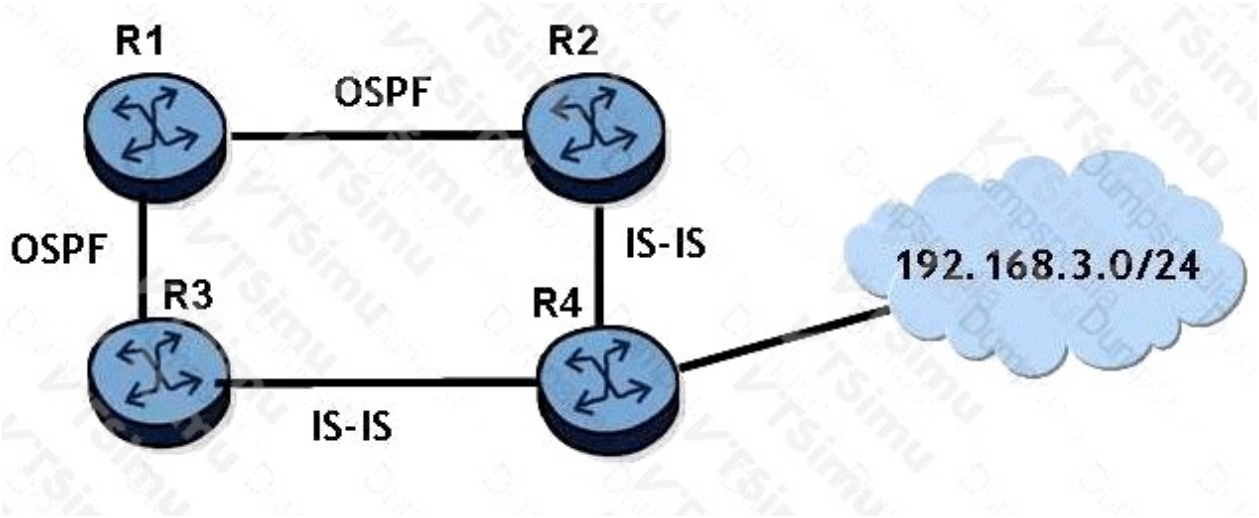
Explanation:

IS-IS uses Point-to-Point IS-to-IS Hello PDUs on a point-to-point circuit. These differ from the Level 1 and Level 2 LAN Hello PDUs used on broadcast networks. Because exactly two routers are attached to a point-to-point circuit, no Designated Intermediate System election is required.

IS-IS point-to-point adjacency formation can use the three-way handshake extension to confirm bidirectional connectivity and detect adjacency changes. The point-to-point model and three-way handshake procedures are specified in RFC 5303. General IS-IS PDU operation is defined in RFC 1195.

QUESTION NO: 66

Click the exhibit button.



If router R2 redistributes the IS-IS route to 192.168.3.0/24 into OSPF, router R3 will receive two routes to 192.168.3.0/24. What will be the preference of these two routes? Assume that all IS-IS routers are L1/L2 capable and are in the same area. Choose two answers.

A. The OSPF internal preference

B. The IS-IS Level 1 internal preference

C. The IS-IS Level 2 internal preference

D. The OSPF external preference

E. The IS-IS Level 2 external preference

F. The IS-IS Level 1 external preference

ANSWER: B D

Explanation:

Router R3 learns 192.168.3.0/24 through its native IS-IS domain and also through OSPF after R2 redistributes that IS-IS route. Because the IS-IS routers are all Level 1/Level 2 capable and belong to the same IS-IS area, the native IS-IS advertisement is an IS-IS Level 1 internal route at R3. In SR OS, an internal IS-IS route is installed with the protocol preference associated with IS-IS Level 1 internal routes; the default preference is 15.

The route injected by R2 into OSPF is not an OSPF intra-area or inter-area route. Redistribution creates an OSPF autonomous-system external advertisement, so R3 classifies the second route as an OSPF external route. SR OS assigns OSPF external routes their external-route preference, which defaults to 150. Route preference is considered before protocol metric when routes from different sources compete for the same prefix, and the lower preference is favored. Consequently, R3 has an IS-IS Level 1 internal route and an OSPF external route for this destination; under default settings, the IS-IS route is preferred for forwarding. Nokia's SR OS routing documentation describes route preference and the protocol-specific default values, while its OSPF documentation covers external routes generated through redistribution. See Nokia SR OS route-policy and route-preference documentation and Nokia SR OS OSPF documentation.

QUESTION NO: 67

Which of the following are major features of OSPF? (Choose 2.)

A. Fast reroute capability.

B. Control traffic prioritization.

C. Route redistribution.

D. Traffic engineering extensions.

E. Cut through forwarding.

ANSWER: C D

Explanation:

Route redistribution. OSPF supports the import of routes learned from other routing domains or protocols into an OSPF autonomous system. These imported destinations are advertised as AS-external routes using AS-external link-state advertisements, allowing OSPF routers to calculate reachability to networks that are external to the OSPF domain. This is a key practical capability for connecting OSPF with static routing, other IGPs, or BGP-based routing environments. RFC 2328 specifies the handling and calculation of AS-external routes: RFC 2328, OSPF Version 2.

Traffic engineering extensions. OSPF has standardized extensions that advertise traffic-engineering attributes, including available bandwidth and link administrative attributes. These attributes are carried in opaque LSAs so that a traffic-engineering application, commonly MPLS traffic engineering, can construct paths that meet specified resource constraints rather than relying only on the normal shortest-path metric. The standardized OSPF traffic-engineering extension is defined in RFC 3630, Traffic Engineering Extensions to OSPF Version 2. Together, route redistribution and TE extensions represent important OSPF control-plane functions used in service-provider and enterprise routing deployments.

QUESTION NO: 68

Which of the following statements regarding RSVP refresh reduction are true1? (Choose two)

A. Message-IDs replace individual refresh messages.

B. Refresh reduction applies only to reservation messages.

C. RSVP generates a summary-refresh message for each LSP.

D. An LSP state change causes an incremented Message-ID.

ANSWER: A D

Explanation:

Message-IDs replace individual refresh messages. and **An LSP state change causes an incremented Message-ID.** are correct. RSVP refresh reduction, specified in RFC 2961 and used with RSVP-TE, reduces the recurring control-plane overhead of traditional RSVP soft-state refreshes. Rather than repeatedly transmitting full PATH and RESV refresh messages for unchanged state, RSVP associates state with Message_ID values and can send a compact Srefresh (summary refresh) message containing Message_ID information for multiple states. The receiving router uses those identifiers to determine which RSVP state is being refreshed, avoiding the need to receive the complete individual refresh message content for every refresh interval. Message identifiers must reliably distinguish the current version of a state block. Consequently, when the state for an LSP changes, the sender advances the Message_ID so peers can recognize that the state has been updated and process the corresponding refreshed information correctly. This mechanism preserves RSVP soft-state reliability while making refresh processing and bandwidth use substantially more efficient, particularly where a router maintains many RSVP-TE LSPs. See RFC 2961: RSVP Refresh Overhead Reduction Extensions and IETF RFC 2961.

QUESTION NO: 69

Which of the following features are supported by IS-IS? Choose three answers.

A. PDU authentication.

B. The ability to customize the link cost metric.

C. Use of Bellman-Ford routing algorithm.

D. Classless routing.

E. Updates sent as layer 3 multicast.

ANSWER: A B D

Explanation:

IS-IS supports PDU authentication, allowing routers to validate received IS-IS protocol data units before accepting the routing information they contain. Authentication is applicable to IS-IS PDUs through authentication-related TLVs, and modern deployments can use cryptographic authentication mechanisms defined for IS-IS. This helps protect the routing domain against unauthorized or spoofed control-plane messages.

The ability to customize the link cost metric is a core IS-IS capability. IS-IS is a link-state protocol that calculates paths using configured interface metrics; an operator can assign a metric to each adjacency, influencing the shortest-path calculation and enabling deliberate traffic-engineering-oriented path selection. The wide-metric extensions provide a larger metric range for contemporary IS-IS networks.

Classless routing is also supported because IS-IS carries IP reachability information in TLVs that include both network prefixes and their prefix lengths. Consequently, it can advertise and calculate routes for variable-length subnet masks and CIDR prefixes rather than being constrained by historical classful network boundaries. For protocol specifications, see RFC 1195, which defines integrated IS-IS for IP, and RFC 5304, which specifies IS-IS cryptographic authentication.

QUESTION NO: 70

Which of the following about CSPF is TRUE?

A. The tail-end router performs the CSPF calculation for an LSP.

B. CSPF calculations are always performed on RSVP-TE based LSPs.

C. The head-end router performs the CSPF calculation for an LSP.

D. Each downstream router performs its own CSPF calculation for an LSP.

ANSWER: C

Explanation:

The head-end router performs the CSPF calculation for an LSP. Constrained Shortest Path First (CSPF) is an ingress, or head-end, path-computation process used with MPLS traffic engineering. Before signaling the LSP, the head-end uses its traffic-engineering database, populated through IGP TE advertisements, to calculate an explicit route that satisfies both the configured destination and applicable constraints. Such constraints can include required bandwidth, administrative groups (link colors), metrics, and other TE attributes. The resulting route is then carried in RSVP-TE signaling so that the LSP can be established hop by hop along the selected path. This model lets the router that originates the LSP apply the service's traffic-engineering policy consistently while using a network-wide view of TE topology and available resources. CSPF is therefore associated with path selection at the LSP ingress rather than being a calculation distributed among transit routers or performed by the LSP egress. Nokia's MPLS documentation describes CSPF as the mechanism used by the LSP head-end to compute constrained TE paths; the RSVP-TE signaling model is specified in RFC 3209. See also Nokia's MPLS documentation.

QUESTION NO: 71

Which of the following can be used for label signaling and distribution on the Alcatel-Lucent 7750 SR? Choose two answers

A. RSVP-TE

B. TDP

C. LDP

D. ADP

E. MPLS-TE

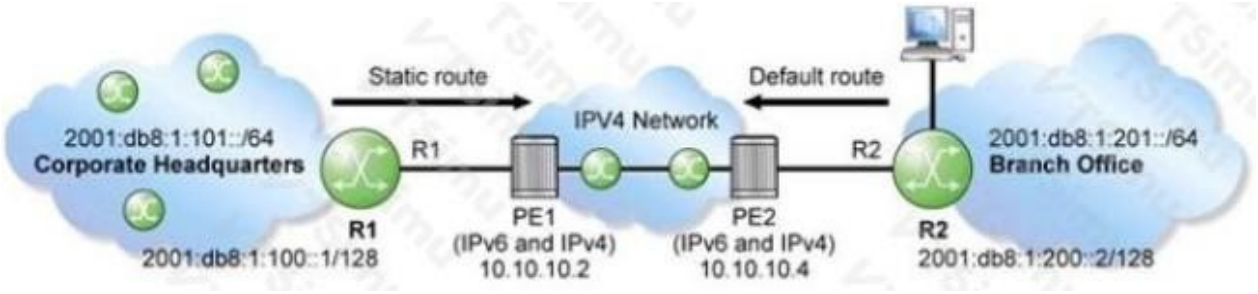
ANSWER: A C

Explanation:

RSVP-TE and LDP are both supported MPLS label-signaling mechanisms on the Alcatel-Lucent 7750 SR. LDP distributes label bindings for MPLS forwarding equivalence classes, commonly supporting transport LSPs that follow the IGP shortest path. Routers establish LDP sessions with peers and exchange label mappings, allowing an MPLS packet to be forwarded along an LSP by label swapping at each hop. RSVP-TE signals explicitly routed, traffic-engineered MPLS LSPs. It carries RSVP Path and Resv signaling messages that can request labels while also reserving resources and applying traffic-engineering constraints such as bandwidth, administrative groups, and explicit hop selection. The 7750 SR MPLS implementation supports both approaches, allowing an operator to use LDP-based transport where ordinary IGP-derived paths are appropriate and RSVP-TE where controlled path placement and resource-aware engineering are required. This distinction is important operationally: both protocols distribute or signal MPLS labels, but RSVP-TE additionally establishes the state needed for traffic-engineered tunnels. Nokia's SR OS MPLS documentation describes MPLS, LDP, and RSVP-TE operation on SR routers; the protocol specifications are also defined in the IETF standards for LDP and RSVP-TE.

QUESTION NO: 72

Click on the exhibit.



A service provider is deploying a 6 over 4 tunnel to connect a customer's corporate IPv6 network to all devices at the branch office as shown in the exhibit.

Which command is used to create the route to the branch office network on PE1?

A. Configure router static-route 2001:DB8:1:201::/64 indirect 10.10.10.4

B. Configure router static-route 2001:DB8:1:201:764 indirect 2001:DB8:1:200::2

C. Configure router static-route 2001:DB8:1:200::/56 indirect 10.10.10.4

D. Configure router static-route 2001:DB8:1:200:756 indirect 2001:DB8:1:200::2

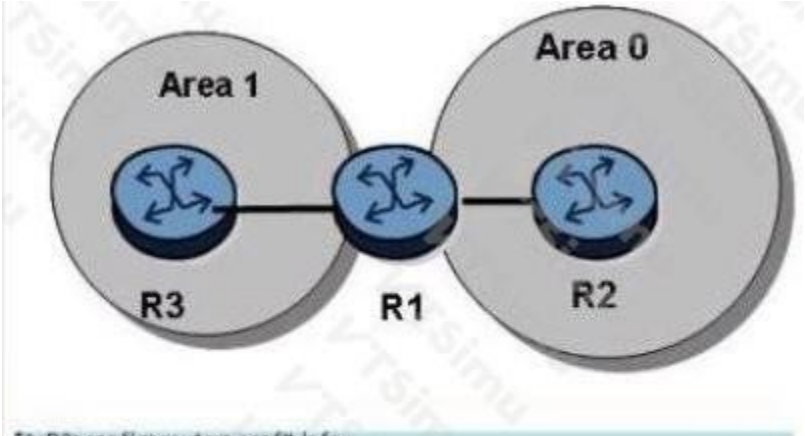
ANSWER: A

Explanation:

`Configure router static-route 2001:DB8:1:201::/64 indirect 10.10.10.4` is the required configuration because the route on PE1 must identify the IPv6 prefix located at the branch office and forward traffic to the IPv4 address of the remote 6-over-4 tunnel endpoint. The branch LAN is the `2001:DB8:1:201::/64` IPv6 subnet, while `10.10.10.4` is reachable through the provider IPv4 network and acts as the indirect next hop for the tunnel destination. This lets PE1 resolve the IPv4 underlay path, encapsulate IPv6 traffic in IPv4 as required for the tunnel, and deliver packets toward the branch router. An indirect static route is appropriate because the IPv4 endpoint is not necessarily a directly connected next hop from PE1; normal IPv4 route resolution determines how to reach it. The tunnel addressing prefix is used for the tunnel interfaces themselves, whereas the route destination must be the actual branch-office IPv6 LAN prefix. Nokia SR OS static routes support indirect next-hop resolution, including recursive resolution through the routing table. See Nokia's Static Routes documentation and the IETF's RFC 4213 IPv6 transition mechanisms for IPv6-over-IPv4 tunneling behavior.

QUESTION NO: 73

Click the exhibit button.




```
*A:R3>config>router>ospf# info
----------------------------------------------
        area 0.0.0.1
            stub
            exit
            interface "system"
            exit
            interface "toR1"
                interface-type point-to-point
                authentication-type password
                authentication-key
"i/XVK6X/L7YM44ZaNqVrJKWb5q244dGbfRvF4eS5KH." hash2
            exit
        exit
----------------------------------------------
*A:R3>config>router>ospf#
```

```
*A:R1>config>router>ospf# info
----------------------------------------------
        area 0.0.0.0
            interface "toR2"
                interface-type point-to-point
            exit
        exit
        area 0.0.0.1
            stub
            exit
            interface "system"
            exit
            interface "toR3"
                interface-type point-to-point
                authentication-type password
                message-digest-key 1 md5
"d7oaaU5pMrRRj7Xl3w7bZxF9h6rw3D0oYBUOKD.mtFs" hash2
            exit
        exit
----------------------------------------------
*A:R1>config>router>ospf#
```

The OSPF adjacency between routers R3 and R1 is down. What is the problem?

A. Area 1 of router R1 should not be configured as a stub.

B. Router R3 has the incorrect authentication key type.

C. Router R1 has the incorrect authentication key type.

D. There is not enough information to tell.

ANSWER: C

Explanation:

Router R1 has the incorrect authentication key type. OSPF neighbors must use compatible authentication settings on the shared link before they can successfully exchange and accept OSPF control packets. The authentication type is carried in the OSPF packet header, and a router validates received Hellos against the authentication mechanism configured for its OSPF interface or area. When R1 is configured with a different key type from R3—for example, simple password authentication on one side and message-digest authentication on the other—the routers reject the received OSPF packets rather than progressing through neighbor-state negotiation. As a result, the neighbor relationship cannot reach the full adjacency state and is reported as down. The authentication type, along with the associated credentials and any message-digest key identifiers where applicable, must be consistent on both ends of the link. This behavior follows the OSPF protocol authentication processing specified in RFC 2328. Nokia SR OS documentation also describes OSPF interface authentication configuration and the need to apply matching authentication parameters to communicating OSPF peers; see the Nokia SR OS documentation portal.

QUESTION NO: 74

Which of the following statements regarding link state protocols are true? (Choose two)

A. When a router receives updates from its neighbors, it adds them to its link state database, performs an SPF computation, and sends the results to its neighbors.

B. An SPF computation is done by each router to determine the best path to destination networks. All the best paths determined by the SPF calculation will be seen in the route table.

C. Each router constructs its own link state database with updates received from neighbors.

D. An SPF computation is done by each router to determine the best path to destination networks. All the best paths determined by the SPF calculation will be offered to the route table manager.

ANSWER: C D

Explanation:

Each router constructs its own link state database with updates received from neighbors. In a link-state protocol, routers originate advertisements describing their directly connected links and flood those advertisements throughout the routing domain or area. Every router receives the topology information and builds a local link-state database representing the network topology. This common topology view enables each router to calculate paths independently rather than relying on a neighbor to provide a completed route calculation.

An SPF computation is done by each router to determine the best path to destination networks. All the best paths determined by the SPF calculation will be offered to the route table manager. The router runs the Shortest Path First algorithm against its local link-state database, producing candidate best routes. On Nokia SR OS, protocol route processing provides these routes to the route table manager, which selects and installs routes according to route preference and other route-selection rules. This separation allows routes learned from multiple protocols to be evaluated consistently before they become active forwarding routes. The OSPF specification describes maintaining a link-state database and using the SPF calculation to construct the shortest-path tree; the IS-IS specification similarly defines distribution of link-state information and shortest-path computation. See RFC 2328 and RFC 1195.

QUESTION NO: 75

When IS-IS elects a DIS, what is used as the tie breaker on an Alcatel-Lucent 7750 SR, if the priorities are the same?

A. The system ID.

B. The loopback address.

C. The sequence number of the Hello packet.

D. The device that first initiated communication becomes the DIS

E. The highest interface MAC address.

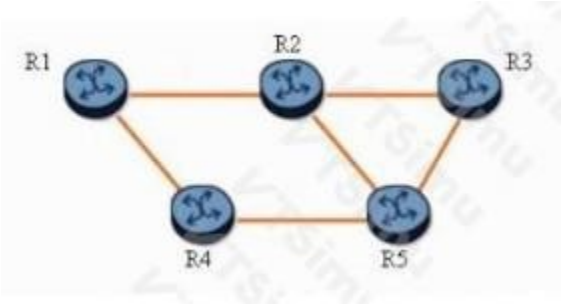
ANSWER: E

Explanation:

The highest interface MAC address is used as the tie breaker. On a broadcast multi-access IS-IS LAN, routers advertise their DIS priority in IS-IS Hello PDUs. The router with the highest priority is elected the Designated Intermediate System (DIS) for that LAN level. When two or more eligible routers advertise the same priority, IS-IS compares the Subnetwork Point of Attachment (SNPA) values associated with the interfaces on that LAN. For an Ethernet interface, the SNPA is the interface MAC address. The router with the numerically highest MAC address therefore becomes the DIS. This behavior is significant operationally because it makes the election deterministic even when administrators leave equal default priorities on multiple routers. On a 7750 SR, an operator can instead influence the preferred DIS directly by configuring a higher IS-IS interface priority on the intended router; the MAC-address comparison is only reached when those priorities are equal. The DIS role is elected independently for each LAN and IS-IS level as applicable. The ISO IS-IS election procedure, including use of the highest SNPA after priority comparison, is specified in RFC 1195, section 7.2.2. Nokia's SR OS documentation is available through the Nokia SR OS documentation portal.

QUESTION NO: 76

Click on the exhibit button below.



With fast reroute enabled, and given the following:

Which of the following statements is true? (Choose two)

A. Router R1 is the LSP originator and is also a PLR.

B. With one-to-one link or node protection, the detour LSP path is R1-R4-R5-R2-R3.

C. With one-to-one link or node protection, the detour LSP path is R1-R4-R5-R3.

D. If the R2-R3 link fails, router R3 becomes a PLR

E. None of the nodes can be protected with FRR.

ANSWER: A C

Explanation:

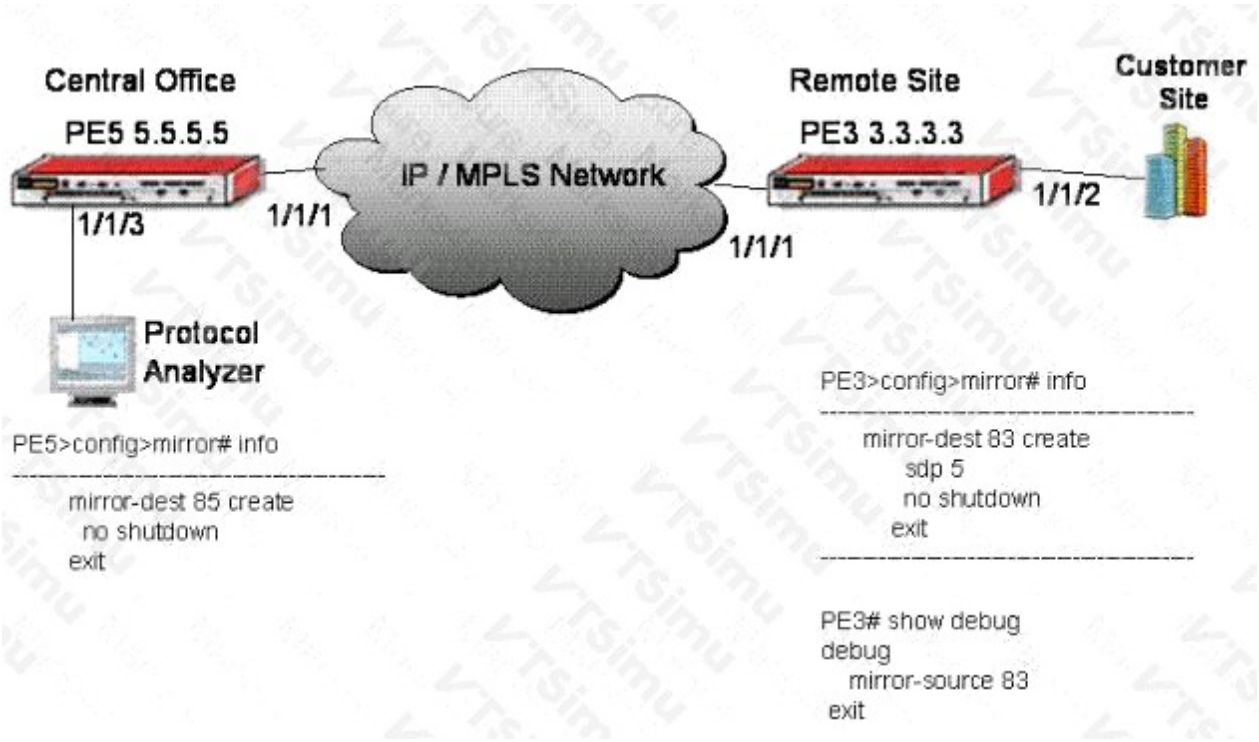
Router R1 is the LSP originator and is also a PLR. The primary RSVP-TE LSP originates at R1 and proceeds toward R3 through R2. Fast Reroute (FRR) permits the router immediately upstream of the protected resource to act as the Point of Local Repair (PLR). Consequently, R1 is both the ingress/originating router for this LSP and the PLR responsible for locally diverting traffic if its protected next-hop resource becomes unavailable.

With one-to-one link or node protection, the detour LSP path is R1-R4-R5-R3. R1 can signal a facility-bypassing detour through R4 and R5 that merges back into the protected LSP at R3, the merge point. This route avoids the R1-R2-R3 primary segment, allowing it to protect either the R1-R2 adjacency or the downstream node R2, while maintaining forwarding toward

the egress at R3. One-to-one backup protection uses a detour LSP associated with the protected primary LSP and is established by the PLR toward an appropriate merge point. This behavior follows RSVP-TE local repair procedures defined in RFC 4090; RSVP-TE LSP signaling fundamentals are specified in RFC 3209.

QUESTION NO: 77

Click on the exhibit button below.



Assuming the customer only needs to monitor traffic from PE3 to the customer site, what three items are missing from PE3 ' s configuration? (Assume that the IP/MPLS network structure between the local and remote PEsis in place.)

A. The mirror destination configuration on PE5 is missing the " remote-source " command,

B. The mirror destination on PE3 needs a " remote-source " command.

C. The mirror service IDs on the local and remote destinations must match.

D. The mirror source on PE3 needs a " remote-source " command.

E. The mirror source command on PE3 needs a command that tells it to monitor port 1/1/2.

ANSWER: A C E

Explanation:

Remote mirroring requires coordinated configuration at both ends of the mirror path. The mirror destination configuration on PE5 is missing the "remote-source" command because the destination PE must identify that it will receive remotely mirrored packets across the IP/MPLS network. This enables the remote endpoint to accept and process the mirrored traffic delivered from the source PE.

The mirror service IDs on the local and remote destinations must match so that the remote mirror stream is associated with the intended mirror service at each endpoint. Matching identifiers provides the service-level correlation needed for the remote mirroring relationship.

The mirror source on PE3 needs a command that tells it to monitor port 1/1/2 because a mirror source must explicitly select the traffic-bearing object being copied. Applying the monitored port under the mirror-source context causes traffic entering and/or leaving that customer-facing port, according to the configured direction, to be replicated toward the remote destination. Together, these settings establish the monitored interface, the remote-receiving destination role, and the

common mirror-service association. Nokia's SR OS documentation portal is available at Nokia SR OS documentation; additional SR OS product and operational information is available from Nokia Service Router Operating System.

QUESTION NO: 78

VPWS allows carriers to offer services such as point-to-point frame-relay, ATM, and Ethernet on a single platform without the need for multiple overlay networks.

A. TRUE

B. FALSE

ANSWER: A

Explanation:

TRUE is correct. A Virtual Private Wire Service (VPWS) is a Layer 2 VPN service that emulates a dedicated point-to-point circuit, often called a pseudowire, across a packet-switched MPLS/IP provider network. This architecture lets a provider use one common transport and service platform to carry traffic from several legacy and native Layer 2 technologies, including Frame Relay, ATM, and Ethernet. Customer-facing attachment circuits retain the appropriate service characteristics while the provider transports the service over MPLS tunnels and pseudowires rather than building distinct, technology-specific overlay networks for each service type.

This is a key operational benefit of VPWS: it supports migration and service consolidation. A carrier can continue delivering point-to-point services to customers using different access technologies while standardizing the provider core on a scalable packet network. The IETF's VPWS framework explicitly describes the provision of point-to-point Layer 2 services over a packet-switched network, and Nokia SR OS documentation covers Layer 2 VPN services and pseudowire-based service delivery. See RFC 4664: Framework for Layer 2 Virtual Private Networks and Nokia SR OS documentation.

QUESTION NO: 79

Which of the following OSPF area types are supported by the Alcatel- Lucent 7750 SR? Choose two answers.

A. Not so stubby areas

B. Level1 area WC

C. Stub areas

D. Partially stubby areas

E. Level 2 areas

ANSWER: A C

Explanation:

The supported area types in this question are **Not so stubby areas** and **Stub areas**. The 7750 SR OSPF implementation supports configuring an area as a stub area, which limits external OSPF route information within that area and uses a default route to reach destinations outside it. This is useful where routers in an area do not need the full external-routing database.

The platform also supports a not-so-stubby area (NSSA). An NSSA retains the stub-area approach to external routes received from other OSPF areas while permitting an autonomous system boundary router within the NSSA to inject external routing information. Those routes are advertised as Type 7 LSAs in the NSSA and are translated for distribution into the normal OSPF domain when appropriate. Nokia's 7750 SR OSPF configuration documentation includes stub and NSSA area configuration under OSPF area parameters. The underlying NSSA behavior is defined by the IETF NSSA specification.

References: Nokia 7750 SR OSPF documentation and RFC 3101: The OSPF Not-So-Stubby Area.

QUESTION NO: 80

When configuring RSVP-TE for transport tunnels, which command enables RSVP by default?

A. configure router tunnel

B. configure router gre

C. configure router rsvp-te

D. configure router transport-tunnel

E. configure router mpls

ANSWER: E

Explanation:

`configure router mpls` is correct because RSVP is part of the SR OS MPLS control-plane feature set used to signal RSVP-TE label-switched paths, including transport tunnels. Entering the router MPLS configuration context is the starting point for enabling and configuring MPLS transport functionality. In the standard SR OS behavior described for MPLS/RSVP operation, RSVP support is enabled by default with MPLS; administrators then configure the RSVP parameters, interfaces, and tunnel or LSP behavior required by the design. This default relationship avoids requiring a separate base RSVP activation step before RSVP-TE transport-tunnel configuration can be applied. The MPLS configuration hierarchy also provides the location from which RSVP-related settings are maintained, ensuring that RSVP signaling is associated with the router's MPLS instance and its participating interfaces. Operationally, MPLS must be administratively enabled with the applicable `no shutdown` configuration before the protocol can operate, but `configure router mpls` is the command context identified by the question for the default RSVP enablement behavior. Nokia's SR OS documentation portal contains the MPLS and RSVP-TE configuration and operational material: Nokia SR OS documentation. The relevant MPLS technology overview is also available through Nokia Service Router Operating System.

QUESTION NO: 81

Click on the exhibit.

```
*A:R2# show router ospf database 1.1.1.1 detail

===============================================================================
OSPF Link State Database (Type : All)(Detailed)
===============================================================================
-------------------------------------------------------------------------------
Network LSA for Area0.0.0.0
-------------------------------------------------------------------------------
Area Id          : 0.0.0.0              Adv Router Id     : 10.10.10.3
Link State Id    : 1.1.1.1 (16843009)
LSA Type         : Network
Sequence No      : 0x80000002           Checksum          : 0x3991
Age              : 97                   Length            : 36
Options          : E
Network Mask     : 255.255.255.0        No of Adj Rtrs    : 3
Router Id (1)    : 10.10.10.3           Router Id (2)     : 10.10.10.2
Router Id (3)    : 10.10.10.6
===============================================================================
```

Which router is the designated router for the broadcast network?

A. 10.10.10.2

B. 10.10.10.3

C. 10.10.10.6

D. There is not enough information given to determine the designated router.

ANSWER: B

Explanation:

10.10.10.3 is the designated router for the broadcast network. OSPF performs a designated-router election on broadcast multi-access networks so that routers do not need to establish a full adjacency with every other router on the segment. The election first considers the configured OSPF interface priority; among eligible routers with the applicable highest priority, the router ID is used as the tie-breaker. A router with an OSPF priority of zero is ineligible to become either the designated router or backup designated router. Based on the election attributes displayed in the exhibit, the router identified as 10.10.10.3 wins the designated-router election for this shared broadcast segment. The elected designated router originates the network-LSA describing the multi-access network and serves as the central adjacency point for the other routers on that segment. This behavior follows the OSPF designated-router election defined in RFC 2328, section 9.4 and is implemented by Nokia SR OS OSPF as described in the Nokia SR OS OSPF documentation.

QUESTION NO: 82

At what point is the LSP SRM flag cleared when an LSP is sent on an IS-IS point-to-point link?

Choose two answers.

A. As soon as the LSP is transmitted.

B. When an acknowledgment is received in the form of PSNP.

C. When an older LSP is received.

D. When a newer or same aged LSP is received.

E. Only when a newer LSP is received

ANSWER: B D

Explanation:

The SRM (Send Routing Message) flag records that an LSP still needs reliable flooding toward a particular adjacent IS-IS neighbor. On a point-to-point link, the flag is cleared when the originating router receives positive evidence that the neighbor has the advertised LSP. A received PSNP provides explicit acknowledgment: its LSP entry and remaining-lifetime information tell the sender that the neighbor has received the relevant version, so retransmission is no longer necessary. The flag is also cleared when a newer or same-aged LSP is received from that neighbor. Receipt of an equal or more recent version demonstrates that the neighbor's database already contains the advertised information, or has progressed beyond it; consequently, retaining the LSP in the retransmission set would serve no purpose. These behaviors are part of IS-IS reliable flooding and database synchronization procedures, preventing needless repeated transmission while ensuring that each adjacency converges on the current LSP version. The IS-IS protocol specification describes the SRM flag and its clearing conditions in its update-process rules; see RFC 1195, LSP transmission process. Point-to-point IS-IS operation is further defined in RFC 5303.

QUESTION NO: 83

What happens by default on an Alcatel-Lucent 7750 SR when a VPLS accepts a tagged frame at the ingress SAP 1/1/1:100?

A. The FCS is verified and kept in the customer frame. The VLAN tag is kept and transported over the network.

B. The FCS is verified and kept in the customer frame. The VLAN tag is removed for transport over the network.

C. The FCS is verified and removed from the frame. The VLAN tag is kept for transport over the network.

D. The FCS is verified and removed from the frame. The VLAN tag is removed for transport over the network.

ANSWER: D

Explanation:

The FCS is verified and removed from the frame. The VLAN tag is removed for transport over the network. On a 7750 SR Ethernet SAP using a VLAN-qualified encapsulation such as the SAP identified by `1/1/1:100`, the received IEEE 802.1Q tag is the service-delimiting tag: it identifies which SAP and therefore which service accepts the frame. After ingress classification has succeeded, SR OS removes that service-delimiting tag from the customer frame as part of SAP ingress processing. The tag is therefore not carried as part of the VPLS payload across the MPLS service core; appropriate egress SAP processing supplies the required service tag when the frame leaves the VPLS.

The Ethernet FCS is also checked by the receiving Ethernet interface and is not retained in the service payload. This is consistent with Ethernet switching and pseudowire operation: the physical-link FCS protects a received link frame, while a new FCS is generated for transmission on the next physical Ethernet link. Thus, VPLS forwards the applicable customer Ethernet content after the ingress service tag and received-link FCS have been handled. Nokia's SR OS Layer 2 service documentation describes SAP encapsulation and VPLS forwarding behavior; see the Nokia SR OS VPLS documentation and the Nokia SR OS SAP documentation.

QUESTION NO: 84

In OSPFv3, which of the following LSA types carries prefixes and is only flooded in the area where it is generated?

A. Link

B. AS External

C. Router

D. Intra-Area Prefix

ANSWER: D

Explanation:

Intra-Area Prefix is correct. OSPFv3 separates topology information from IPv6 address-prefix information. An Intra-Area-Prefix LSA advertises IPv6 prefixes associated with a router or a transit network within an area. It references the relevant Router or Network LSA so that SPF topology calculation can be performed independently of prefix advertisement, then the advertised prefixes can be attached to the appropriate routing information.

This LSA has area flooding scope: it is originated and flooded throughout the OSPF area in which the advertised intra-area prefixes exist, rather than being propagated throughout the entire autonomous system. This design is one of the notable differences between OSPFv3 and OSPFv2, where prefix and topology details are more closely combined in Router and Network LSAs. RFC 5340 defines the Intra-Area-Prefix LSA format, its association with Router and Network LSAs, and its area-scoped flooding behavior. Nokia SR OS OSPF documentation likewise describes OSPFv3 prefix advertisements as separate LSA information used for IPv6 route calculation.

References: RFC 5340, Intra-Area-Prefix-LSA; Nokia SR OS OSPF documentation.

QUESTION NO: 85

You wish to advertise LDP labels for all local networks; which is the most effective policy statement to use?

A. from prefix-list direct

B. match local

C. from protocol direct

D. match protocol direct

E. from protocol local

ANSWER: C

Explanation:

`from protocol direct` is correct because an LDP export policy selects the IP routes for which the router allocates and advertises label bindings. In SR OS routing policy terminology, directly connected network routes are installed with the `direct` protocol. Matching that protocol therefore covers all local, directly connected networks without requiring every connected prefix to be individually listed in a prefix list. This is the most effective approach when the intent is to advertise LDP labels consistently for the complete set of local networks, including networks added later through interface configuration changes. The policy is applied in the LDP context as an export policy, where the matching routes determine the FECs whose labels may be advertised to LDP peers. It also expresses the operational intent clearly: select routes based on their route source, rather than on a manually maintained address list. Nokia's SR OS documentation describes routing-policy match criteria, including protocol-based matching, and its MPLS/LDP documentation describes using policies to control LDP label advertisement. See the Nokia SR OS documentation portal and the SR OS routing-policy documentation.

QUESTION NO: 86

Click the exhibit button.

The following command sequence is executed on router R2:

A: R2# configure router ospf router-id 10.10.10.99
A: R2# configure router router-id 10.10.10.66

On router R1, what router ID appears for router R2 directly after these commands are executed?

A. 10.10.10.99

B. 10.10.10.66

C. 10.10.10.2

D. 10.10.10.1

ANSWER: C

Explanation:

10.10.10.2 is correct because an OSPF neighbor identifies itself to adjacent routers using the router ID that was active when its OSPF instance and neighbor relationship were established. Entering a new OSPF-specific router ID or a router-wide router ID in the configuration does not immediately replace the router ID being advertised in already-operational OSPF hello packets, the neighbor database, or the existing adjacency. Consequently, immediately after these configuration commands, R1 continues to display R2 using its current operational OSPF router ID, 10.10.10.2.

For a changed OSPF router ID to become operational, the OSPF process must be restarted or otherwise reset so that adjacencies are rebuilt using the new identity. This behavior prevents an active neighbor from changing identity in place, which would disrupt the link-state database and adjacency state. Nokia SR OS documentation describes the OSPF router-ID as the identifier used by the OSPF instance and documents the router-level and OSPF configuration hierarchy. See the Nokia SR OS OSPF documentation and the Nokia SR OS router configuration documentation.

QUESTION NO: 87

What modes of authentication are supported for RIPv2 on the Alcatel-Lucent 7750 SR?

A. RIPv2 supports basic and MD5 authentication only.

B. RIPv2 supports password and MD5 authentication only.

C. RIPv2 does not support authentication, password and MD5 authentication.

D. RIPv2 does not support authentication.

ANSWER: C

Explanation:

RIPv2 does not support authentication, password and MD5 authentication is correct because the 7750 SR provides three operational choices for RIP version 2 authentication on an interface: no authentication, plain-text password authentication, or MD5 authentication. “Does not support authentication” in this wording identifies the unauthenticated mode; it does not mean that the RIP implementation lacks authentication capability entirely. Password authentication places a simple password in the RIPv2 authentication information, while MD5 uses keyed message-digest authentication to validate RIP updates. These choices allow an operator to retain interoperability where no authentication is configured, deploy basic RIPv2 password protection where required, or use MD5 authentication when stronger integrity protection is needed. The behavior follows the RIPv2 authentication framework, which defines the simple-password authentication mechanism in the RIPv2 specification, and the keyed-MD5 extension defined for RIP. Nokia’s SR OS routing-protocol documentation describes RIP configuration and its interface-level authentication controls. See RFC 2453: RIP Version 2, especially its authentication packet format, and RFC 2082: RIP-2 MD5 Authentication for the MD5 method.

QUESTION NO: 88

How can a router ID be created on an Alcatel-Lucent 7750 SR router? Choose three answers

A. From the least significant 32 bits of the chassis MAC address.

B. From the most significant 32 bits of the chassis MAC address.

C. From the loopback IP address.

D. By using the command "config router router-id X.X.X.X"

E. From the system IP address.

F. From the highest interface IP address.

ANSWER: A D E

Explanation:

On an SR OS router, the router ID is a 32-bit value represented in IPv4 dotted-decimal notation and is used to identify the router in routing-protocol operations. It can be explicitly set with `config router router-id X.X.X.X`, which is the preferred approach when an operator needs a stable, intentional identifier independent of interface addressing. When no explicit router ID is configured, SR OS can derive the ID from the configured system IP address. The system interface is the router’s logical management and control-plane endpoint, so its address is designed to remain reachable and stable across physical-interface changes. If neither an explicit router ID nor a system IP address provides the value, the platform can derive a router ID using the least significant 32 bits of the chassis MAC address. This provides an automatically generated identifier without requiring interface-IP selection. These mechanisms establish a deterministic router ID while allowing operators to override the automatic selection through configuration. Nokia’s SR OS router configuration documentation describes router-level identity and system-interface behavior, and the command reference documents router configuration commands: Nokia SR OS documentation and Nokia 7750 SR Router Configuration Guide.

QUESTION NO: 89

What are the two major differences between configuring an IES service and configuring VPWS or VPLS services? (Choose 2)

A. IES does not support QoS.

B. IES has a configurable IP-MTU

C. A virtual route table must be configured for an IES service.

D. An IES interface has an IP address assigned to it.

ANSWER: B D

Explanation:

An IES service provides routed IP connectivity using the system's base routing context. Consequently, an IES interface is a Layer 3 interface and is assigned an IP address. That address provides the IP endpoint for attached customer equipment and enables the router to route traffic to and from the IES interface using standard IP forwarding behavior. This is a fundamental distinction from VPWS and VPLS, which are Layer 2 service constructs that transport customer Ethernet frames rather than terminating customer IP addressing on the service interface.

IES configuration also includes an IP MTU, which controls the maximum IP packet size accepted for forwarding on the IP interface. The configured value must account for the payload size supported by the underlying SAP or network transport so that routed packets can be handled without unintended fragmentation or drops. Nokia's service documentation describes IES as an IP service with IP interfaces and associated interface parameters, including MTU behavior. See the Nokia IES service overview and the Nokia router configuration documentation for the routed-interface model and IP configuration context.

QUESTION NO: 90

Which of the two following statements are TRUE?

A. When an Alcatel-Lucent Service Router is rebooted, the mirror destination and mirror source must be re-configured

B. A mirror destination is shut down by default on an Alcatel-Lucent Service Router.

C. A mirror source is shut down by default on an Alcatel-Lucent Service Router.

D. When an Alcatel-Lucent Service Router is rebooted, the mirror destination must be re-configured.

E. When an Alcatel-Lucent Service Router is rebooted the mirror source must be re-configured.

ANSWER: B E

Explanation:

A mirror destination is shut down by default on an Alcatel-Lucent Service Router, and a mirror source must be re-configured after an Alcatel-Lucent Service Router reboot. In SR OS, a mirror destination is a configured object that identifies where mirrored traffic is sent. Its administrative state defaults to shutdown, so an operator must explicitly enable it before it forwards copies of packets. This prevents traffic from being mirrored unintentionally when a destination is created or restored from configuration. Mirror-source associations, by contrast, are operational monitoring settings that are deliberately not retained across a reboot. This behavior ensures that packet replication does not automatically resume after a restart without an operator intentionally restoring the required monitoring session. The destination configuration remains available in the saved router configuration; the source configuration is the item that must be reapplied. These behaviors are documented in Nokia SR OS mirror-service configuration and operational guidance. See the Nokia SR OS documentation portal and the Nokia 7750 SR OS information center for mirror-service configuration documentation.

QUESTION NO: 91

Which of the following SAP encapsulations support multiple services on a port? (Choose 3)

A. dot1q

B. null

C. qinq

D. atm

E. ipcp

F. bcp-null

ANSWER: A C D

Explanation:

In SR OS, **dot1q**, **qinq**, and **atm** SAP encapsulations can identify individual traffic instances on the same physical port, allowing that port to carry SAPs for multiple services. With dot1q encapsulation, the customer VLAN identifier provides the service-delimiting value, so different VLAN tags received on a port can be mapped to separate service access points. QinQ similarly uses stacked VLAN-tag information, supporting service differentiation where an outer and inner VLAN context is required. ATM supports multiplexing through virtual-path and virtual-channel identification, enabling separate ATM PVCs or VCs on the same ATM interface to be associated with different services. These encapsulations therefore provide an identifier beyond the physical port itself, which is the essential capability needed to support multiple services on one port. Nokia's Services Overview documentation describes SAPs as the binding between a service and an access interface, including the use of encapsulation values to distinguish traffic on shared access ports. See the Nokia SR OS Services Overview Guide and the Nokia SR OS MD-CLI Command Reference.

QUESTION NO: 92

Which of the following statements regarding IPv4 fragmentation are true? Choose three answers.

A. A router can fragment a packet when its size exceeds the outgoing interface MTU and the DF bit is clear.

B. Intermediate routers reassemble fragments before forwarding them.

C. The destination host reassembles the received fragments.

D. A packet with the DF bit set is fragmented regardless of the outgoing interface MTU.

E. The fragment offset field is used during reassembly.

ANSWER: A C E

Explanation:

An IPv4 router can fragment a packet when the packet is larger than the outgoing interface MTU and the Don't Fragment (DF) bit is not set. Each fragment receives its own IPv4 header, and the fragment offset field is used by the destination host to reassemble the original datagram. Fragment reassembly is performed by the destination host, not by intermediate routers.

If the DF bit is set and the packet exceeds the outgoing interface MTU, the router discards the packet and normally sends an ICMP Destination Unreachable message with the fragmentation-needed code. IPv4 fragmentation behavior is specified in RFC 791 and ICMP reporting behavior is described in RFC 1191.

QUESTION NO: 93

Click the exhibit.

```
R1>config>service# info
----------------------------------------------
        customer 1 create
            description "Default customer"
        exit
        customer 100 create
            description "Cust 100"
        exit
        sdp 2 mpls create
            far-end 10.10.10.2
            ldp
            keep-alive
                shutdown
            exit
            no shutdown
        exit
        epipe 12 customer 100 create
            sap 1/1/1:30 create
            exit
            spoke-sdp 2:44 create
                no shutdown
            exit
            no shutdown
        exit
----------------------------------------------
```

Which of the following statements about the E-pipe service is TRUE?

A. The service is a local service.

B. The service requires a second spoke SDP on this router.

C. The service's VC ID is 44.

D. The service will accept VLAN tag 100.

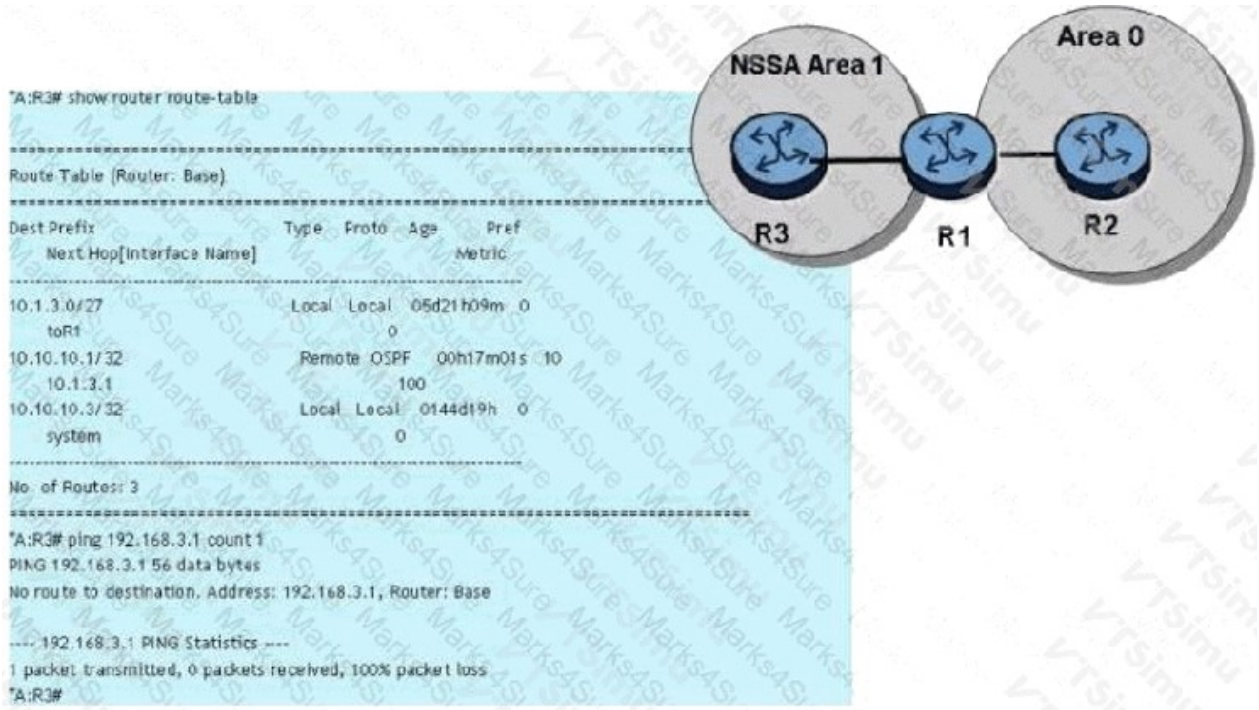
ANSWER: C

Explanation:

The service's VC ID is 44. In an SR OS Epipe configuration, the `vc-id` configured under a spoke SDP identifies the pseudowire, also called a virtual circuit, used for that point-to-point service. The exhibit shows a VC ID value of 44 for the service's spoke SDP; therefore, 44 is the value that must be used to identify and match this pseudowire with its far-end counterpart. This identifier is significant because a Layer 2 VPN pseudowire is established only when the relevant signaling and service parameters, including the VC identifier where applicable, correspond between the provider-edge routers. Epipe is Nokia SR OS's point-to-point Ethernet service type, and its transport-side association is made through an SDP or spoke SDP rather than by the customer-facing SAP alone. Nokia's SR OS service documentation describes Epipe as a point-to-point service and documents the role of SDPs and spoke SDPs in carrying service traffic. The pseudowire signaling model and use of VC identifiers are also defined in the IETF LDP pseudowire specification. See the Nokia SR OS Services Overview and RFC 4447: Pseudowire Setup and Maintenance Using LDP.

QUESTION NO: 94

Click on the exhibit.



192.168.3.1 is a loopback interface on router R2 and is distributed to OSPF area 0, but the ping fails from router R3.

Which of the following is a possible solution to the problem?

A. Add the " originate-default-route " option in the NSSA area configuration on router R3.

B. Add the " summaries " option to the NSSA area configuration on router R1.

C. Configure a static route on router R1 to 192.168.3.1.

D. Configure a default route on router R2.

ANSWER: B

Explanation:

Add the " summaries " option to the NSSA area configuration on router R1. R1 is the area border router between the backbone and the NSSA containing R3. The loopback address 192.168.3.1 is learned in area 0, so R3 requires an inter-area route advertised by R1 to reach it. In an NSSA configuration where inter-area summaries are suppressed, the ABR does not inject the Type 3 summary LSA for the backbone destination into the NSSA. Consequently, R3 has no OSPF route for 192.168.3.1 and cannot forward the ping toward R1. Enabling summaries on R1's NSSA configuration permits R1 to advertise the relevant inter-area summary information into the NSSA. R3 can then install a route to 192.168.3.1 with R1 as the next hop, while R1 retains its backbone reachability toward R2. This is the appropriate control point because the ABR governs exchange of inter-area OSPF information between area 0 and the NSSA. Nokia documents NSSA behavior and ABR route advertisement under its OSPF configuration guidance: Nokia SR OS OSPF documentation. The underlying NSSA specification, including handling of summary LSAs, is defined in RFC 3101.

QUESTION NO: 95

Which of the following statements regarding OSPF designated router election are true? Choose two answers.

A. An interface priority of 0 prevents a router from becoming DR or BDR.

B. The router with the highest eligible interface priority is preferred as DR.

C. The router with the lowest router ID is selected when priorities are equal.

D. A new router with a higher priority always preempts the current DR.

E. A DR election is required on every point-to-point interface.

ANSWER: A B

Explanation:

OSPF performs a DR and BDR election on broadcast multi-access networks. An interface priority of zero makes a router ineligible to become either DR or BDR. Among eligible routers, the highest interface priority is preferred. If priorities are equal, the highest router ID is used as the tie-breaker. The election is non-preemptive, so a newly arriving router with a higher priority does not normally replace an existing DR.

Point-to-point interfaces do not require a DR election because there are only two routers on the link. The OSPF designated-router election procedure is defined in RFC 2328, section 9.4.

QUESTION NO: 96

Which of the following commands may only be used on the head end router to verify the status of an LSP named "LSP 1"? (Choose two)

A. show router mpls lsp "LSP 1" originate

B. show router mpls lsp "LSP 1" path detail

C. show router rsvp session originate

D. show router rsvp session "LSP 1"

ANSWER: B C

Explanation:

The head-end-only commands are `show router mpls lsp "LSP 1" path detail` (represented in the question with "lsp," which is evidently intended to be `lsp`) and `show router rsvp session originate`. An RSVP-TE LSP is originated at its head-end router, where the router maintains the locally configured LSP object, its configured paths, path preference, active/standby path state, signaling status, and operational path details. Consequently, the detailed MPLS LSP path display is associated with that locally originated LSP and is used at the ingress/head-end to inspect its path operation. The RSVP command with the `originate` qualifier similarly restricts output to RSVP sessions created locally by that router. Since RSVP signaling for an LSP is initiated by the ingress router, this is specifically a head-end view and can be used to locate the session for "LSP 1" among locally originated sessions. These commands complement one another: the MPLS LSP command presents LSP and path-oriented operational information, while the RSVP originated-session view confirms the RSVP signaling session initiated by the router. Nokia's SR OS documentation provides MPLS and RSVP operational-command context in its SR OS documentation portal and its router configuration documentation.

QUESTION NO: 97

Which of the following statements regarding per-platform label space are true? Choose two answers.

A. A separate label is used for each interface that the FEC is advertised on.

B. It uses fewer label resources than a per-interface label space.

C. It is typically used when the device has ATM or Frame Relay interfaces.

D. It uses more label resources than a per-interface label space.

E. A single label may be assigned to an FEC and used across all interfaces of the same router.

F. It is only valid when the LDP peers are directly connected over an interface, and the label is only going to be used for traffic sent over that interface.

ANSWER: B E

Explanation:

“It uses fewer label resources than a per-interface label space” is correct because a per-platform label space is shared by all interfaces on the same LSR. The router can bind one locally significant label to a given FEC and reuse that label regardless of the outgoing interface. This avoids maintaining separate label bindings for the same FEC on every interface and therefore conserves the router’s available label values and forwarding-label resources.

“A single label may be assigned to an FEC and used across all interfaces of the same router” correctly describes the defining behavior of a platform-wide label space. In LDP terminology, the label is allocated from the platform label space rather than from an interface-specific space, so its scope is the entire LSR. The receiving router can interpret the label consistently no matter which interface receives the labeled packet, subject to its MPLS forwarding configuration. RFC 5036 defines the platform-wide label-space concept and its use by LDP in RFC 5036, section 2.2.1. Nokia SR OS MPLS documentation is available through the Nokia Documentation Center.

QUESTION NO: 98

Which of the following are characteristics of ECMP LSPs in LDP? Choose three answers

A. Liberal label retention mode

B. T-LDP enabled with peers to which load balancing is desired.

C. ECMP enabled in the config>router context.

D. Hashing used to select the egress label.

E. Export policy applied to advertise multiple labels for a given FEC.

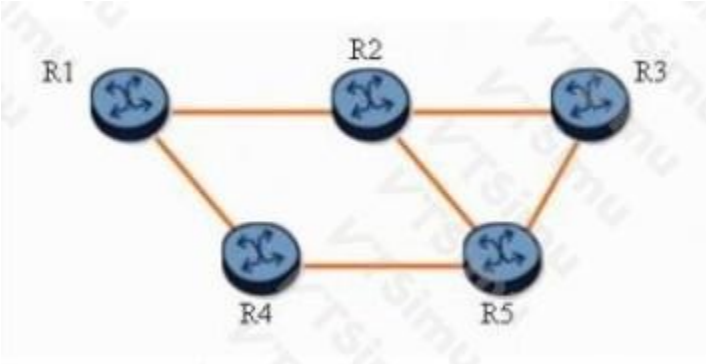
ANSWER: A C E

Explanation:

ECMP operation for LDP relies on retaining and using label mappings learned over multiple eligible equal-cost paths. **Liberal label retention mode** supports this behavior because the router keeps label bindings learned from LDP peers even when they are not currently the selected path; this allows alternate equal-cost LSPs to remain available without waiting for a new label exchange. **ECMP enabled in the config>router context** is required because the routing table must install and expose multiple equal-cost next hops before LDP can resolve a FEC across those equal-cost paths and create usable multipath transport. **Export policy applied to advertise multiple labels for a given FEC.** is also characteristic of LDP ECMP operation: LDP policy controls which FECs and associated label bindings are advertised, enabling downstream label information to be distributed for the paths participating in the LDP topology. Together, route-level ECMP, retained alternate label bindings, and appropriate LDP label advertisement policy provide the control-plane foundation for multiple LDP LSPs toward the same destination. Nokia’s LDP configuration and operational behavior are documented in its LDP configuration guide; the underlying LDP label-distribution procedures are standardized in RFC 5036.

QUESTION NO: 99

Click on the exhibit button below.



Given the following fast reroute protected LSPs:

-LSP1.R1-R4-R5-R3

-LSP2.R1-R4-R5-R2

Which routers become bypass tunnel PLRs? Choose two answers

A. Router R1

B. Router R2

C. Router R3

D. Router R4

E. Router R5

ANSWER: A E

Explanation:

Router R1 and **Router R5** become bypass-tunnel PLRs in the illustrated FRR design. A point of local repair is the router that locally detects a protected failure and switches affected RSVP-TE LSP traffic onto a pre-established bypass tunnel, avoiding the delay of waiting for end-to-end RSVP-TE signaling and path recomputation. R1 is a PLR for the protected resource immediately downstream on the common initial portion of both listed LSPs. R5 is also a PLR because it is upstream of the separately protected downstream resources used by the two LSPs as they proceed toward R3 and R2. Thus, these routers provide local repair at the protection points represented in the topology while preserving traffic forwarding for LSP1 and LSP2 after the associated protected failure. RSVP-TE fast reroute defines a bypass tunnel as a detour that a PLR can use to protect one or more LSPs over a protected link or node. This behavior and the PLR role are specified in RFC 4090: Fast Reroute Extensions to RSVP-TE; RSVP-TE signaling fundamentals are specified in RFC 3209.

QUESTION NO: 100

Which of the following statements regarding LSP path configuration are true? (Choose three)

A. The path must include at least one hop.

B. The path may be used multiple times in a single LSP.

C. The path may comprise loose hop entries only.

D. The path must define the LSP's tail end.

E. The path may be used for multiple LSPs.

ANSWER: A C E

Explanation:

In SR OS MPLS RSVP-TE configuration, a named path is an explicit-path definition made up of configured hop entries, so it must contain at least one hop to provide a usable path definition. Each hop identifies a routing constraint used when RSVP signaling establishes the LSP. The configured hops may all be loose hops. A loose hop directs the signaling process toward a specified downstream node or address while allowing normal routing or CSPF to select the intervening route, which is useful when only selected transit points need to be constrained rather than every link along the route.

A named path is also a reusable configuration object. It can be referenced by multiple LSPs, enabling a common traffic-engineering policy to be consistently applied without duplicating the hop list under every LSP. The LSP itself supplies its configured endpoint, while the path supplies the explicit routing constraints that guide signaling toward that endpoint. Nokia's MPLS documentation describes configured paths and their use with RSVP-TE LSPs; RSVP explicit-route behavior is standardized by the IETF. See the Nokia SR OS documentation portal and RFC 3209: RSVP-TE.

QUESTION NO: 101

If routers R1 and R2 have a Level 1 and Level 2 IS-IS adjacency, which of the following statements are true? Choose two answers.

A. Routers R1 and R2 are in different areas,

B. Routers R1 and R2 are in the same area.

C. Routers R1 and R2 will send a Level 1 Hello PDU and Level 2 Hello PDU, if the interface type between them is point-to-point.

D. Routers R1 and R2 will send a Level 1 Hello PDU and Level 2 Hello PDU, if the interface type between them is broadcast.

E. Routers R1 and R2 will send a Level 1/2 Hello PDU, if the interface type between them is broadcast.

ANSWER: B D

Explanation:

Routers R1 and R2 are in the same area. A Level 1 IS-IS adjacency is established only between routers that share the same Level 1 area address. Since the routers have a Level 1 adjacency in addition to a Level 2 adjacency, they necessarily have a common area and can exchange Level 1 link-state information for that area. Level 2 adjacency operation allows inter-area connectivity, but the simultaneous presence of Level 1 adjacency establishes that these two routers meet the same-area requirement.

Routers R1 and R2 will send a Level 1 Hello PDU and Level 2 Hello PDU, if the interface type between them is broadcast. IS-IS uses LAN Hello PDUs on broadcast media. A router supporting both levels sends distinct Level 1 and Level 2 LAN Hello PDUs, allowing neighbors to form the corresponding adjacencies and participate in the appropriate level-specific LAN election and database synchronization processes. This behavior follows the IS-IS protocol definition for Level 1 and Level 2 LAN operation. See RFC 1195 and the ISO IS-IS protocol overview in RFC 1142.

QUESTION NO: 102

Which of the following conditions are necessary in order for service labels to be signaled between two PE routers? (Choose 3)

A. LDP must be enabled on both PE routers.

B. RSVP must be enabled on both PE routers.

C. Traffic engineering must be enabled in the IGP.

D. The VC-ID must match between the two PE routers in both directions.

E. The transport tunnel must be operationally up between the two PE routers in both directions.

F. The SDP ID must match between the two PE routers in both directions.

ANSWER: A D E

Explanation:

Service-label signaling for an MPLS Layer 2 VPN service is performed using the Label Distribution Protocol between the PE routers, normally through a targeted LDP session associated with the service distribution path. Therefore, LDP must be enabled on both PE routers so that the peers can establish the signaling relationship and advertise the pseudowire/service-label bindings. The VC-ID must match between the two PE routers in both directions because it identifies the same pseudowire endpoint/service to each PE during LDP label mapping; matching identifiers allow each side to associate the received label binding with the intended service. In addition, the transport tunnel must be operationally up between the two PE routers in both directions. This provides the MPLS transport path over which labeled service traffic can be carried once the remote service label is learned and installed. These requirements reflect the separation of service-label signaling from the underlying MPLS transport: the signaling peers need a working bidirectional transport reachability relationship, while the service itself needs a common VC identifier. RFC 4447 defines LDP-based pseudowire setup and label mapping procedures, including use of the VC identifier: RFC 4447. Nokia SR OS Layer 2 service documentation is available at Nokia SR OS documentation.

QUESTION NO: 103

On which object is ADSPEC enabled?

A. SDP

B. Network port

C. LSP

D. Service

E. RSVP

ANSWER: C

Explanation:

LSP is correct. In SR OS RSVP-TE configuration, ADSPEC is an RSVP attribute associated with an MPLS label-switched path. The ADSPEC (advertising specification) carries end-to-end service information used by RSVP-capable applications and receivers when establishing or evaluating a reservation. Because this information relates to the signaling and characteristics of a particular RSVP-TE path, SR OS enables it in the RSVP context of the configured LSP, rather than on an underlying physical interface or a service construct. An LSP defines the traffic-engineered path through the MPLS network and is the appropriate scope for RSVP signaling behaviors such as ADSPEC advertisement. This aligns with the RSVP integrated-services model, in which ADSPEC is carried in RSVP PATH messages to describe the service available along the path. The RSVP specification defines ADSPEC semantics and its role in PATH processing in RFC 2210. Nokia's MPLS documentation describes RSVP-TE LSP configuration as part of the MPLS feature set; see the Nokia SR OS MPLS Guide.

QUESTION NO: 104

A router has two admin groups defined (GREEN=1 and RED=3). An interface is associated to both admin groups. What value will be shown in the Admin Group sub-TLV for this interface?

A. 0x4

B. 0x5

C. 0x10

D. 0xa

ANSWER: D

Explanation:

0xa is correct because an IS-IS administrative group is advertised as a bit mask, with the configured administrative-group value identifying the bit position rather than a numeric value to be summed directly. GREEN=1 sets bit 1, whose mask value is 2 (0x2). RED=3 sets bit 3, whose mask value is 8 (0x8). Since the interface belongs to both groups, the advertised Admin Group sub-TLV combines those independently set bits with a logical OR: 0x2 OR 0x8 = 0xA. This encoding lets a link advertise membership in multiple administrative groups simultaneously, allowing constraint-based SPF calculations to include or exclude links based on the relevant group bits. The IS-IS Traffic Engineering extensions define the Administrative Group sub-TLV as a 32-bit field in which each bit represents membership in an administrative group. The extended administrative-group mechanism retains the same bit-mask membership model when more than 32 group bits are needed. See RFC 5305, IS-IS Extensions for Traffic Engineering and RFC 7308, Extended Administrative Groups in IS-IS.