PassQueen
a.k.a dumpsqueen.com

Alcatel-Lucent Multi Protocol Label Switching

Alcatel-Lucent 4A0-103

Version Demo

Total Demo Questions: 35

Total Premium Questions: 352

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Volume A	59
Topic 2, Volume B	99
Topic 3, Volume C	88
Topic 4, Volume D	86
Total	352

QUESTION NO: 1

Click on the exhibit button below.

Assuming that LSP R1-R2-R3 requests the default fast reroute configuration, which router becomes a DMP for the detour tunnels?

A. Router R1

B. Router R2

C. Router R3

D. Router R4

E. Router R5

ANSWER: E

Explanation:

Router R5 becomes the detour merge point (DMP) for the detour tunnel created to protect the RSVP-TE LSP segment. In a one-to-one fast-reroute implementation, the point of local repair computes and signals a detour LSP around the protected resource. The detour must rejoin the original primary LSP at a router downstream from the protected resource; that downstream rejoin router is the DMP. With the default fast-reroute behavior and the topology shown in the exhibit, the eligible bypass path from the repair point reaches the original LSP again at Router R5. Router R5 therefore terminates the detour portion and forwards traffic back onto the primary LSP path toward its egress. This preserves forwarding continuity during a protected-link or protected-node failure while retaining the original LSP forwarding state after the repair path merges. RFC 4090 defines this RSVP-TE local-repair model, including detour LSPs and their merge behavior: RFC 4090, Fast Reroute Extensions to RSVP-TE. Nokia's MPLS fast-reroute documentation also describes RSVP-TE detours and local protection operation: Nokia MPLS Fast Reroute documentation.

QUESTION NO: 3

Click on the exhibit.

```
*A:SRC_R3# oam lsp-ping prefix 192.10.1.2/32
LSP-PING 192.10.1.2/32: 80 bytes MPLS payload
Seq=1, send from intf toR1, reply from 10.10.10.2
       udp-data-len=32 ttl=255 rtt=2.45ms rc=3 (EgressRtr)

---- LSP 192.10.1.2/32 PING Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min = 2.45ms, avg = 2.45ms, max = 2.45ms, stddev = 0.000ms
*A:SRC_R3#
```

Which of the following is FALSE?

A. The router can reach prefix 192.10.1.2/32 through a RSVP-TE tunnel.

B. The router can reach prefix 192.10.1.2/32 through an LDP tunnel.

C. The router 10.10.10.2 is the egress router of the prefix 192.10.1.2/32.

D. The lsp-ping tool is used to perform a unidirectional LSP test.

ANSWER: A

Explanation:

The statement “The router can reach prefix 192.10.1.2/32 through a RSVP-TE tunnel.” is false because the forwarding information shown in the exhibit associates reachability to the destination with the LDP-signaled transport LSP, not an RSVP-TE LSP. In SR OS, the transport tunnel selected for an IP prefix is determined by the route resolution and tunnel-table-manager resolution process. A configured or available RSVP-TE LSP does not by itself establish that a particular destination prefix is forwarded through that RSVP-TE tunnel; the active resolved transport tunnel must be the RSVP-TE LSP. Here, the displayed LDP label-switched path is the transport path that provides the MPLS forwarding reachability toward the destination’s egress router. Therefore, describing the prefix as reachable through an RSVP-TE tunnel does not match the forwarding state depicted in the exhibit. Nokia’s MPLS documentation describes how SR OS supports LDP and RSVP-TE transport LSPs and how tunnel selection is used for route resolution. The ISP ping function then validates forwarding over the selected MPLS data-plane LSP toward the egress, rather than changing the transport tunnel selected for the route. See the Nokia SR OS MPLS Guide and the Nokia SR OS MPLS OAM Guide.

QUESTION NO: 4

If the destination IP address of a received unlabeled packet is successfully matched in the FIB, which of the following statements are true? Choose two answers.

A. A PUSH label stack operation must be performed.

B. The outgoing interface and next-hop address will be determined from the FIB.

C. The packet can be forwarded as a labeled or unlabeled packet.

D. A POP label stack operation may be performed.

ANSWER: B C

Explanation:

The outgoing interface and next-hop address will be determined from the FIB. A successful destination-prefix lookup selects the forwarding entry associated with the best matching route. That entry supplies the information needed to send the packet onward, including the egress interface and the appropriate next-hop resolution. This is the fundamental IP forwarding function performed for an arriving packet that does not yet carry an MPLS label.

The packet can be forwarded as a labeled or unlabeled packet. The selected forwarding entry can describe ordinary IP forwarding, in which the packet leaves without an MPLS label stack, or it can resolve to an MPLS forwarding action at an ingress LER. Where the route is associated with an MPLS forwarding equivalence class and a usable label forwarding entry, the router imposes the required label stack before transmission. Thus, label imposition depends on the forwarding information selected for that destination rather than being an unconditional consequence of every successful IP FIB lookup.

This behavior follows the MPLS architecture’s separation of IP routing lookups from label-based forwarding and its definition of label-stack imposition for packets entering an MPLS domain. See RFC 3031, Multiprotocol Label Switching Architecture and RFC 3032, MPLS Label Stack Encoding.

QUESTION NO: 5

Which of the following prerequisites must you configure on a router to enable LDP-over-RSVP support? (Choose two)

A. Configure the T-LDP peers for tunneling.

B. Configure tunneled T-LDP sessions between PE routers.

C. Enable ldp-over-rsvp in the IGP configuration on the PEs and ABRs

D. Enable LSP stitching on the RSVP LSPs

ANSWER: A C

Explanation:

LDP-over-RSVP requires the router to use targeted LDP peers for tunneling and to enable LDP-over-RSVP in the IGP configuration on the PEs and ABRs. Configuring the T-LDP peer for tunneling makes the targeted LDP session use an RSVP-TE LSP as its transport path rather than relying on ordinary IP forwarding. This is essential when the LDP session must remain associated with, and protected by, the RSVP signaled tunnel across the MPLS core.

Enabling LDP-over-RSVP in the relevant IGP instances on PE and ABR routers provides the control-plane support needed for the network to resolve and use RSVP LSP transport appropriately as the LDP-over-RSVP path traverses IGP areas. The option's "ldp" spelling is understood as LDP. Together, these settings establish the targeted LDP tunnel behavior and the IGP-level LDP-over-RSVP capability required by the feature. Nokia SR OS MPLS documentation is available from the Nokia SR OS documentation portal. The underlying LDP session and label-distribution behavior is standardized in RFC 5036.

QUESTION NO: 6

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:

RSVP refresh reduction uses Message_ID objects to identify the current state of RSVP PATH and RESV messages. Instead of repeatedly transmitting a complete individual refresh message for every maintained state, a router can send a Summary Refresh message containing the relevant message identifiers. The receiving router uses those identifiers to confirm that its corresponding RSVP state remains current; if it is missing state or has an identifier mismatch, it can request the full message through the reliable-message mechanisms. Thus, Message-IDs provide the basis for replacing routine individual refresh traffic with substantially smaller summary-refresh processing. Each maintained RSVP state also has an associated message identifier that represents a particular version of that state. When an LSP state change modifies the information carried in the RSVP message, the sender assigns an incremented Message-ID to identify the new version. This lets the peer distinguish the updated state from the previously acknowledged state and request reliable retransmission if necessary. These mechanisms reduce periodic RSVP control-plane bandwidth while preserving state synchronization and recovery behavior. The RSVP refresh-reduction procedures, including Message_ID objects, Summary Refresh messages, acknowledgements, and retransmission handling, are specified in RFC 2961. RSVP signaling fundamentals and PATH/RESV state behavior are specified in RFC 2205.

QUESTION NO: 7

Assume that, on the Alcatel-Lucent 7750 SR, an LSP is configured with one-to-one FRR and node protection. Which of the following statements are FALSE? (Choose two)

A. Each PLR signals a detour LSP that avoids the downstream node. If the downstream node cannot be avoided, a detour LSP that avoids the downstream link is signaled.

B. The detour LSP can take one or more hops before merging back on the protected LSP.

C. The merge point of the protected LSP must be the next-next-hop node

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 detour LSPs.

E. The label recording desired flag must be set in the session attribute object in order for the PLR to know what label to swap when using the detour LSP.

ANSWER: C E

Explanation:

The false statements are “The merge point of the protected LSP must be the next-next-hop node” and “The label recording desired flag must be set in the session attribute object in order for the PLR to know what label to swap when using the detour LSP.” For node-protecting one-to-one fast reroute, the detour must bypass the protected downstream node, but its merge point is not restricted to the next-next-hop. It may merge at any suitable downstream router on the protected LSP, provided the selected detour satisfies the required protection and routing constraints. This flexibility is fundamental to finding a node-protecting repair path when the topology does not offer a suitable short repair path. The RSVP-TE label-recording request causes label information to be included in the Record Route Object; it is useful for visibility and certain processing purposes, but it is not the forwarding prerequisite that determines whether the PLR can swap traffic to its already established detour. Upon a failure, the PLR uses its local protected-LSP and detour forwarding state to redirect traffic rapidly. RSVP fast-reroute behavior, including detours, PLRs, and merge points, is specified in RFC 4090, section 3. The RSVP-TE Session Attribute and label-recording mechanisms are defined in RFC 3209, section 4.7.

QUESTION NO: 8

Click on the exhibit button below.

LSP 1's link between routers R3 and R6 has failed, and the RESV state at router R3 has timed out. Which of the following actions will router R3 initiate?

A. Send a PathTear message towards router R6.

B. Send a PathTear message towards router R1.

C. Send a ResvTear message toward router R6.

D. Send a ResvTear message toward router R1.

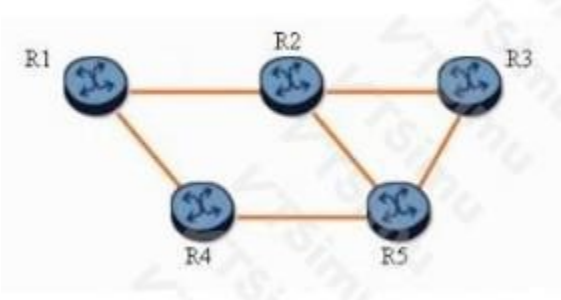
ANSWER: D

Explanation:

Send a ResvTear message toward router R1. is correct. RSVP maintains PATH state in the downstream direction, from the LSP ingress toward the egress, while RESV state is established and refreshed in the reverse, upstream direction—from the receiver/egress side back toward the sender/ingress side. In the stated LSP, R6 is downstream of R3 and R1 is upstream of R3. When R3’s RESV state expires because the downstream adjacency to R6 has failed and RESV refreshes are no longer received, R3 must remove that reservation state and signal its removal upstream. It does this by generating a ResvTear toward R1, allowing upstream routers to promptly release the reservation and associated MPLS RSVP-TE resources rather than waiting for their own refresh timers to expire. This behavior is consistent with RSVP’s soft-state model: state is refreshed periodically, but an explicit teardown propagates in the direction appropriate to the state being removed. The RSVP specification defines ResvTear processing and its upstream propagation in its reservation teardown procedures. See RFC 2205, Reservation Teardown and RFC 2205, Reservation State.

QUESTION NO: 9

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 protected primary LSP originates at R1 and follows the path through R2 toward R3. In RSVP-TE Fast Reroute, the point of local repair (PLR) is the upstream router that detects a protected resource failure and immediately redirects traffic to a pre-established backup or detour LSP. Therefore, R1 performs both functions in this topology: it originates the primary LSP and locally repairs the protected R1-to-R2 resource.

With one-to-one link or node protection, the detour LSP path is R1-R4-R5-R3. One-to-one FRR creates a detour from the PLR to a merge point downstream of the protected resource. For node protection, the detour must bypass not only the failed R1-R2 link but also router R2 itself. The alternate route from R1 through R4 and R5 reaches R3 directly, making R3 the merge point and preserving forwarding toward the LSP destination without traversing the protected node. This is consistent with RSVP-TE local repair procedures, including facility and one-to-one backup methods, specified in RFC 4090. RSVP-TE signaling fundamentals are also defined in RFC 3209.

QUESTION NO: 10

Click on the exhibit.

Both Path-1 and Path-2 are down. If Path-2 becomes available again, how does the headend router select the secondary LSP path?

A. The router will switch to Path-2 because it is the first secondary path configured in the list.

B. The router will not switch to Path-2 because the current path has the longest uptime.

C. The router will switch to Path-2 because it is a more preferred path.

D. The router will not switch to Path-2 because the router will not revert to another secondary path.

ANSWER: C

Explanation:

The router will switch to Path-2 because it is a more preferred path. In an MPLS RSVP-TE LSP configuration, the headend maintains the configured primary and secondary path information and uses path preference to determine the best eligible path. When no configured path is operational, the LSP cannot use either of the unavailable paths. Once Path-2 becomes available again and successfully meets the signaling and operational requirements, it is eligible for selection. Because the exhibit identifies it as the more preferred secondary path, the headend selects it rather than retaining an unavailable state or relying on the order in which secondary paths were entered.

Path preference provides deterministic recovery behavior: the headend can use a less-preferred eligible path during a failure, then move to a higher-preference path when that path becomes available, subject to the configured retry, revertive, and signaling behavior. This selection is made by the headend, which computes and signals the LSP path. Nokia's MPLS documentation describes RSVP-TE LSP path configuration and operational behavior in the Nokia SR OS MPLS Guide. Additional MPLS configuration and path-management material is available through the Nokia SR OS MPLS overview.

QUESTION NO: 11

Which of the following are features of RSVP-TE? Choose three answers.

A. It supports IGP-based LSPs.

B. It supports constrained path LSPs.

C. It does not support loose explicit path LSPs.

D. It is recommended for use in Traffic Engineering applications.

E. It cannot be used with facility backup fast reroute.

ANSWER: A B D

Explanation:

RSVP-TE supports IGP-based LSPs, constrained path LSPs, and traffic-engineering use cases. An RSVP-TE tunnel is signaled using RSVP extensions for MPLS, and, when an explicit route is not imposed, the normal IP routing path toward the tunnel endpoint can determine the forwarding route. That route is ordinarily derived from the IGP, so RSVP-TE can establish an LSP along an IGP-selected path. RSVP-TE also carries traffic-engineering information and path constraints in its signaling. A head-end can use a constraint-based shortest-path calculation to select a route meeting requirements such as reservable bandwidth, administrative constraints, and link attributes, then signal that selected route as an LSP. This ability to select and reserve resources over a path is a principal reason RSVP-TE is used for MPLS traffic engineering. The RSVP-TE specification defines the extensions that allow MPLS label distribution and explicit routing through RSVP signaling, while operational MPLS documentation describes RSVP-TE LSPs as traffic-engineered tunnels whose paths may be calculated using TE constraints. See RFC 3209: RSVP-TE and Nokia SR OS MPLS documentation.

QUESTION NO: 12

Which of the following statements best defines an LSP?

A. An LSP is a bidirectional path across the MPLS domain.

B. An LSP is a bidirectional path for a particular FEC across the MPLS domain.

C. An LSP is a unidirectional path across the MPLS domain.

D. An LSP is a bidirectional path between adjacent LSRs.

E. An LSP is a unidirectional path between adjacent LSRs.

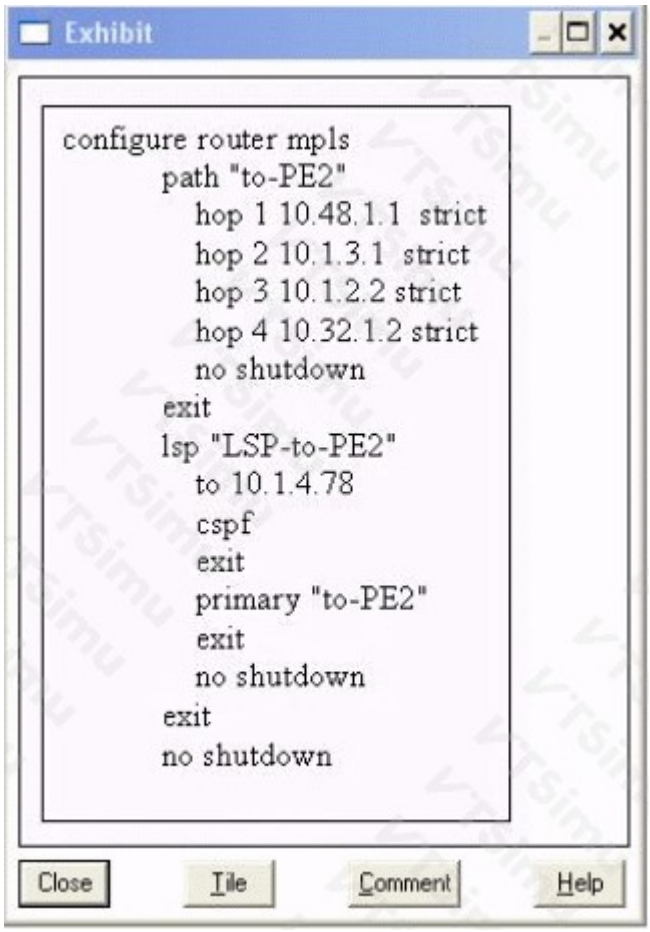
ANSWER: C

Explanation:

An LSP is a unidirectional path across the MPLS domain. In MPLS terminology, a Label Switched Path is the sequence of label-switching routers that packets follow from an ingress LSR toward an egress LSR. At each hop, the forwarding operation uses the incoming MPLS label to select the next hop and the associated label action, such as swapping, popping, or pushing a label. The direction is significant: the forwarding state and labels established for traffic from one ingress to one egress do not inherently provide the reverse forwarding path. If bidirectional service is needed, it is normally implemented using two unidirectional LSPs, one in each direction, which may or may not traverse the same routers. An LSP can be established for traffic belonging to a forwarding equivalence class, but its core definition is the directed path through the MPLS network rather than a bidirectional adjacency or merely a connection between two neighboring routers. RFC 3031 defines an LSP as the path through one or more LSRs followed by a packet for a particular FEC; MPLS forwarding behavior is inherently defined in the packet's direction of travel. See RFC 3031, Multiprotocol Label Switching Architecture and RFC 3036, LDP Specification.

QUESTION NO: 13

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

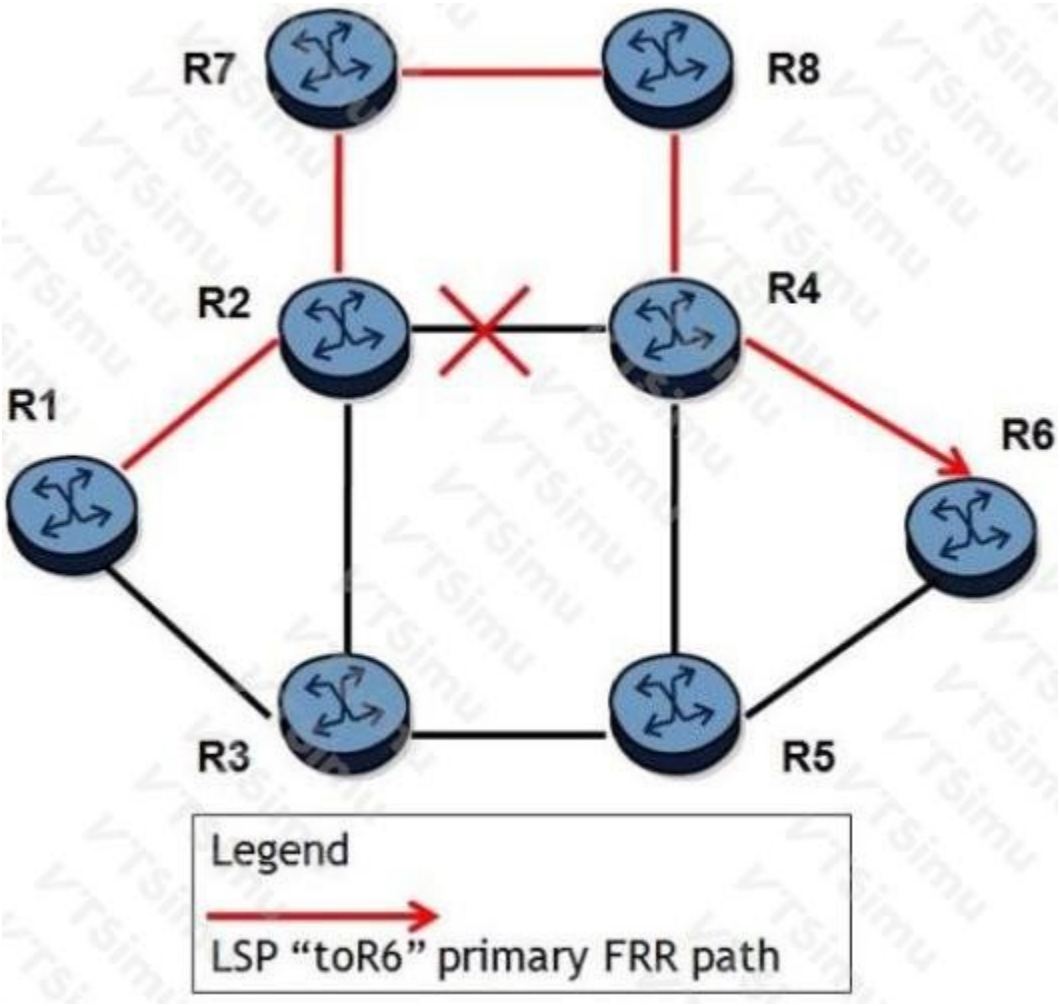
ANSWER: B C

Explanation:

The commands **fast-reroute facility** and **fast-reroute one-to-one** enable MPLS Fast Reroute protection for an LSP by selecting one of the two RSVP-TE local-repair methods. Facility backup creates a bypass tunnel that protects a network-resource segment, such as a link or node, and that bypass can be shared by multiple protected LSPs. This is the facility-based model described in the MPLS Fast Reroute architecture. One-to-one backup instead establishes a detour LSP for the particular protected LSP, providing an individually associated local repair path. In SR OS LSP configuration, either command is entered in the LSP context to request the corresponding FRR behavior; the applicable method is then signaled through RSVP-TE and used by the point of local repair when an eligible protected failure is detected. Both approaches are intended to provide rapid local switchover rather than waiting for the ingress to re-signal an end-to-end replacement path. The MPLS Fast Reroute mechanisms and the distinction between facility and one-to-one protection are specified in RFC 4090. Nokia's SR OS MPLS documentation is available through the Nokia SR OS documentation portal.

QUESTION NO: 14

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 repair at the point of failure, so forwarding can continue over the R2-R7-R8-R4-R6 repair path without waiting for an end-to-end LSP setup. The repair is a temporary protection mechanism; it does not become the preferred CSPF path simply because it is carrying traffic. The ingress router is informed that the protected resource has failed and starts its retry process. When the configured retry interval expires, CSPF runs again using the current traffic-engineering topology, in which the R2-to-R4 link is unavailable. Because the LSP is fully loose, CSPF is free to select a new end-to-end route rather than being constrained to retain R2 as a mandatory hop. With equal link costs, R1-R3-R5-R6 is the available shortest route to R6. The router establishes that replacement LSP using RSVP-TE make-before-break behavior and then moves traffic to it. FRR behavior and the role of local repair are defined in RFC 4090; RSVP-TE signaling and path-error processing are specified in RFC 3209.

QUESTION NO: 15

Which of the following statements regarding the MPLS SWAP label operation is true?

A. The inner label will be exchanged with the MPLS outer label.

B. The inner label will be exchanged for a new label.

C. The outer label will be exchanged for a new label.

D. All labels will be exchanged for one new label.

E. All labels will be exchanged for an equal number of new labels.

ANSWER: C

Explanation:

The outer label will be exchanged for a new label. In MPLS forwarding, a label-switching router examines the topmost entry in the incoming label stack—the outer label—and uses it to look up the forwarding instruction in its label forwarding table. When that instruction is a swap, the router replaces that top label with the locally assigned outgoing label for the next-hop label-switched path, then forwards the packet through the selected interface. The remainder of the label stack is retained unchanged. This permits each MPLS hop to use its own locally significant label value while continuing to identify the same forwarding equivalence class and maintain the established path through the network. The replacement label may have a different numerical value because MPLS labels have local significance between adjacent peers; it is the forwarding state at each router, not a globally identical label number, that preserves path continuity. RFC 3031 defines label swapping as replacing the label at the top of a label stack with a new label before forwarding, and RFC 3032 describes the stack encoding and identifies the top-of-stack forwarding behavior. See RFC 3031, MPLS Architecture and RFC 3032, MPLS Label Stack Encoding.

QUESTION NO: 16

Click on the exhibit button below.

With fast reroute enabled, and given the following:

The protected LSP path is R1-R2-R3.
All links are of equal cost.

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. In RSVP-TE fast reroute, the point of local repair (PLR) is the upstream router that detects a protected-resource failure and immediately switches traffic to a pre-established local repair tunnel. For the protected LSP segment beginning at R1 toward R2, R1 performs that local-repair role while remaining the head-end, or originator, of the end-to-end LSP. The detour LSP path is R1-R4-R5-R3 for the illustrated one-to-one protection case because the alternate route leaves the PLR at R1, traverses the equal-cost alternate connectivity through R4 and R5, and rejoins the protected LSP at R3. This provides a bypass that avoids the protected portion requiring repair and allows traffic to continue toward the merge point without waiting for head-end re-signaling. One-to-one backup uses a detour LSP associated with the protected LSP, and its path is selected to satisfy the applicable link or node protection constraints. These PLR and merge-point behaviors are defined by RSVP-TE fast reroute procedures in RFC 4090. Nokia's MPLS fast-reroute implementation and terminology are also documented in the Nokia 7750 SR MPLS FRR guide.

QUESTION NO: 20

Which of the following regarding the MPLS SWAP label operation is TRUE?

A. The inner label will be exchanged with the MPLS outer label.

B. The inner label will be exchanged for a new label.

C. The outer label will be exchanged for a new label.

D. All labels will be exchanged for one new label.

E. All labels will be exchanged for an equal number of new labels.

ANSWER: C

Explanation:

The outer label will be exchanged for a new label. This accurately describes the MPLS swap operation at a label-switching router. MPLS packets may carry a stack of labels, and forwarding decisions are made using the top, or outermost, label in that stack. After receiving a labeled packet, the router uses this top label to identify the appropriate forwarding entry and next hop. Where the forwarding entry specifies a swap action, the router replaces that incoming outer label with the outgoing label assigned for the next label-switched hop. It then forwards the packet without changing the remaining labels beneath it. This lets each hop use a locally significant label value while maintaining the established label-switched path across the network. The label stack is therefore preserved except for its top entry, which is rewritten to the value required by the downstream router. This behavior is defined by the MPLS architecture and forwarding procedures: label swapping uses the incoming top label as an index and substitutes the corresponding outgoing top label before transmission. See RFC 3031, Multiprotocol Label Switching Architecture and RFC 3032, MPLS Label Stack Encoding.

QUESTION NO: 21

Click on the exhibit button below:

Consider the following:

- Router R2 is an ABR in OSPF Areas 0 and 1.

- Router R2 advertises the aggregate prefix 10.10.10.0/24 into Area 1.

- Router R1 in Area 0 generates LDP labels for FECs 10.10.102/31 and 10.10.10.4/31

Router R3 in Area 1 places these FECs in its LIB, but it does not place them in its LFI

B.
What can you do on router R3 to make it populate these labels in its LFIB?

A. Disable route summarization on router R3

B. Enable LDP aggregate prefix match

C. Create an import policy to accept these FECs

D. Disable LDP exact match on the specified FECs

ANSWER: B

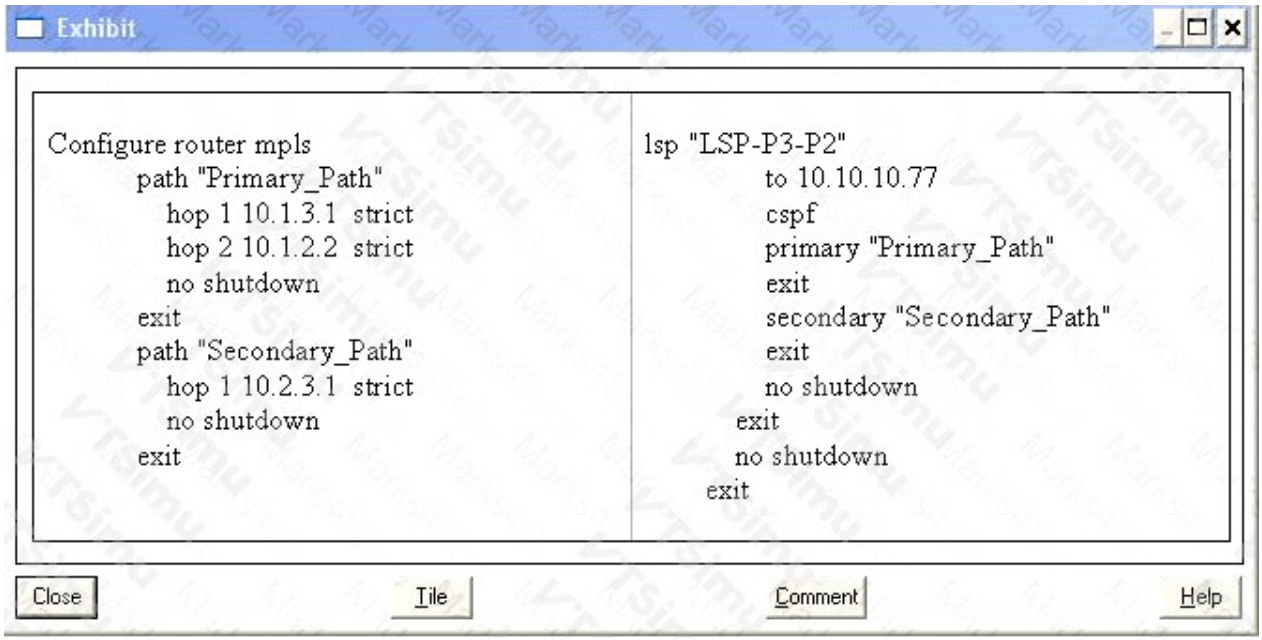
Explanation:

Enable LDP aggregate prefix match. R3 learns an OSPF route only for the ABR-advertised aggregate 10.10.10.0/24, while the LDP mappings received for the relevant forwarding-equivalence classes are more-specific /31 prefixes within that aggregate. For an LDP FEC to be usable for label switching, SR OS must be able to resolve the FEC through the routing table and identify a valid next hop. With the normal exact-prefix behavior, a /31 FEC does not match the installed /24 route,

so its received label mapping can remain in the LIB without producing an active LFIB entry. Aggregate-prefix matching changes this resolution behavior: LDP may use the longest matching covering route, here the advertised /24 aggregate, to resolve the more-specific FEC. R3 can then associate the received label with the resolved next hop and install the necessary forwarding state in its LFIB. This feature is specifically useful when IGP summarization is present between the router advertising labels and the router receiving and using those label bindings. Nokia's SR OS documentation is available through the Nokia SR OS documentation portal; the base LDP protocol behavior and FEC-label mappings are defined in RFC 5036.

QUESTION NO: 22

Click on the exhibit button below.



Given the configuration below, which of the following scenarios is FALSE?

A. Both the primary and secondary paths are explicitly defined.

B. The primary path is active and the secondary path is in hot standby.

C. The LSP defined will use the primary path and will signal the secondary path to become active if the primary path fails.

D. The LSP has the same source and destination for both primary and secondary paths.

ANSWER: B

Explanation:

The statement *"The primary path is active and the secondary path is in hot standby."* is false. In the displayed RSVP-TE LSP configuration, the primary and secondary paths define alternate explicit routes for the same LSP, but a configured secondary path is not automatically a hot-standby, concurrently established forwarding path. The primary path is the active, signaled LSP under normal conditions. If that primary path fails, RSVP-TE attempts to establish and use the configured secondary path, subject to the usual signaling and resource-availability requirements. Thus, the secondary route supplies protection through an alternate path definition, rather than maintaining an already active hot-standby LSP. This behavior lets the router retain a deterministic backup route while avoiding the continuous reservation of resources for a second active path unless a particular protection configuration explicitly requires it. RSVP-TE signaling establishes label-switched paths and reserves resources using RSVP procedures, as specified in RFC 3209. Nokia's MPLS documentation also describes the use of primary and secondary LSP paths for path selection and failover: Nokia SR OS MPLS documentation.

QUESTION NO: 23

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 physical interface has 10 Gbps of bandwidth, but RSVP is configured to make only 30 percent of that capacity available for RSVP-TE LSP reservation. The RSVP reservable bandwidth is therefore 3 Gbps: 10 Gbps × 0.30 = 3 Gbps. When the signaled LSP reserves 1 Gbps, RSVP deducts that reservation from the configured RSVP reservable pool, leaving 2 Gbps of unreserved RSVP bandwidth: 3 Gbps − 1 Gbps = 2 Gbps. This value represents bandwidth still available for additional RSVP-signaled LSPs; it is not the total physical bandwidth remaining for ordinary interface traffic. RSVP-TE carries requested bandwidth as part of LSP signaling and uses available reservable bandwidth when making admission decisions, as described in RFC 3209. Nokia SR OS MPLS and RSVP behavior, including RSVP-TE resource reservation concepts, is documented in the Nokia 7750 SR MPLS Guide.

QUESTION NO: 24

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** are the bypass-tunnel Point of Local Repair (PLR) routers. In MPLS Fast Reroute, a PLR is the router that detects a failure affecting a protected resource and immediately redirects traffic from the protected primary LSP into a pre-established bypass tunnel. This local action provides rapid recovery without waiting for ingress-to-egress signaling or IGP convergence.

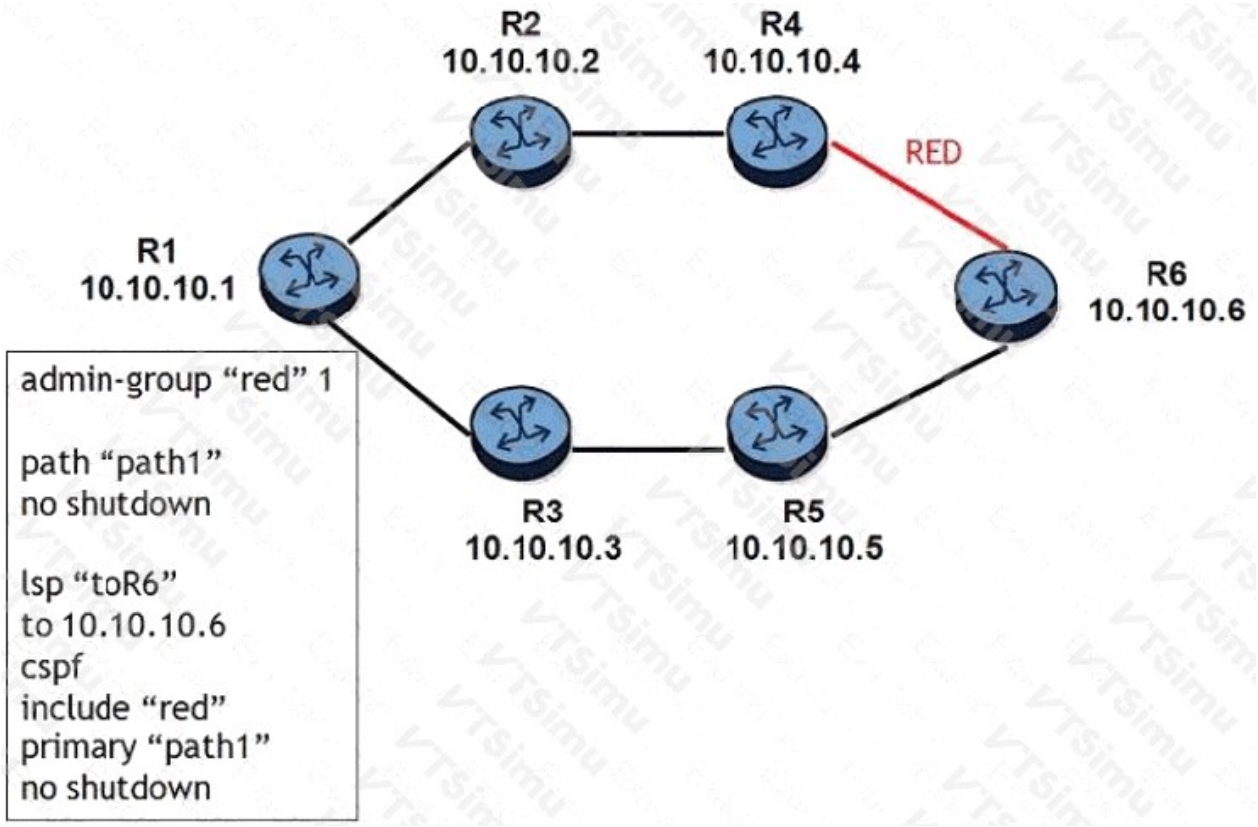
For the displayed protected LSP paths, R1 is the upstream repair point associated with the common protected portion as traffic enters the shared path toward R4 and R5. R5 is also a repair point because it is the last common router before the LSPs continue toward their respective endpoints, R3 and R2; it can locally divert traffic for the protected downstream resource. Thus, the FRR design places bypass-tunnel repair responsibility at R1 and R5 for the two protected LSPs.

This behavior follows the MPLS FRR bypass model defined in RFC 4090, Fast Reroute Extensions to RSVP-TE for LSP Tunnels, which describes the PLR as the router that switches traffic onto a bypass tunnel upon local failure detection.

Nokia's MPLS documentation also describes FRR as a local protection mechanism using precomputed backup paths: Nokia SR OS MPLS Guide.

QUESTION NO: 25

Click on the exhibit.



The link between routers R4 and R6 are configured with an admin group "red". What happens when the LSP is first signaled?

A. The LSP will go through the path R1-R2-R4-R6.

B. The LSP will go through the path R1-R3-R5-R6.

C. The LSP is down because router R1 detects that no path can meet its constraints.

D. The LSP is down because router R4 detects that no path can meet its constraints.

ANSWER: C

Explanation:

The LSP is down because router R1 detects that no path can meet its constraints. RSVP-TE traffic-engineered LSP path computation is performed at the ingress router when the LSP is initially signaled. R1 applies the LSP's configured constraints to the traffic-engineering database, including administrative-group (link-color) information advertised for each TE link. After the R4–R6 link is assigned the *red* administrative group, the available topology no longer contains an end-to-end route that satisfies the complete set of constraints shown in the exhibit. Consequently, R1 cannot produce an eligible explicit route and the LSP cannot be established. The failure is therefore identified at the ingress during constrained shortest-path-first processing, before RSVP signaling can successfully create the LSP along a route. Administrative groups are specifically intended to let TE policies include or exclude links according to attributes such as color, administrative domain, or resource class; these attributes are considered as part of constraint-based path selection. This behavior follows RSVP-TE's use of an ingress-selected explicit route and MPLS traffic-engineering's constraint-based routing model. See RFC 3209: RSVP-TE and RFC 2702: MPLS Traffic Engineering Requirements.

QUESTION NO: 26

Click on the exhibit.

An LSP traverses the path R1-R2-R3. With one-to-one fast reroute enabled, which router becomes a DMP for the detour tunnels?

A. Router R1

B. Router R2

C. Router R3

D. Router R4

E. Router R5

ANSWER: E

Explanation:

Router R5 is the Detour Merge Point (DMP) in the illustrated one-to-one fast-reroute topology. One-to-one FRR creates an RSVP-TE detour LSP for a particular protected LSP. Where multiple detour paths meet at a common router before proceeding toward the original LSP's destination, that router is the DMP. The DMP processes the merging detour state and enables the detour traffic to continue on the appropriate path toward the merge point and, ultimately, the protected LSP's egress. In this case, the detour paths converge at R5, so R5 has the DMP role. This is distinct from the point of local repair, which detects the protected-resource failure and switches traffic onto a repair path. The DMP designation specifically concerns convergence of detour tunnels in the one-to-one backup model. RSVP-TE fast-reroute terminology and detour operation are defined in RFC 4090, Fast Reroute Extensions to RSVP-TE. Nokia's MPLS documentation also describes RSVP-TE protection and fast-reroute behavior in its MPLS Guide.

QUESTION NO: 27

Fill in the blanks: The device at the beginning of an LSP is referred to as the_____, and the device at the end of an LSP is referred to as the_____.

A. ingress LER, egress LER

ANSWER: A

Explanation:

The correct completion is **ingress LER, egress LER**. An MPLS Label Switched Path (LSP) carries labeled packets unidirectionally through an MPLS network. The ingress Label Edge Router (LER) is the edge device where traffic enters the MPLS domain and is classified into a forwarding equivalence class (FEC). It imposes, or pushes, the MPLS label stack that identifies the packet's forwarding treatment and LSP. The egress LER is the edge device at which the LSP terminates and traffic leaves the MPLS domain; it removes the MPLS label information as needed and forwards the resulting packet toward its non-MPLS destination. Thus, these labels describe the edge roles at the start and end of the path, rather than merely the intermediate label-switching function. RFC 3031 defines ingress and egress roles in MPLS architecture and describes label edge routers as devices that operate at the boundary of an MPLS domain. See RFC 3031: Multiprotocol Label Switching Architecture and the Nokia SR OS MPLS overview.

QUESTION NO: 28

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. is correct because this LSP type follows the normal shortest-path reachability selected by the interior gateway protocol. The ingress router uses the routing information available for the destination forward-equivalence class to determine the path; it does not perform traffic-engineering constraint calculations to select an alternate route. This provides a straightforward label-switched path that aligns with the IGP's currently selected path.

The constraint-based routed LSP is calculated at the ingress of the network, based on the router database and user defined constraints. is also correct. The ingress label-switching router performs constraint-based shortest-path-first path computation using traffic-engineering topology information and configured requirements, such as required bandwidth, administrative groups, hop restrictions, and setup or holding priorities. Once the ingress has selected an eligible explicit route, signaling establishes the LSP along that calculated path. This ingress-based computation is the basis for MPLS traffic engineering and enables paths to be selected independently of the ordinary IGP shortest path when network policy or resource requirements demand it.

These principles are described by the RSVP-TE signaling specification in RFC 3209 and the traffic-engineering extensions to OSPF in RFC 3630.

QUESTION NO: 29

Which of the following is NOT a characteristic of MPLS?

A. It provides the ability to establish connection oriented paths over a connectionless IP network.

B. It provides a mechanism to engineer network traffic patterns independently of IP routing tables.

C. It will work over most physical layer protocols.

D. It can use GRE tunnels to transport data from many different protocols over an IP network.

ANSWER: D

Explanation:

"It can use GRE tunnels to transport data from many different protocols over an IP network." is not a characteristic of MPLS. Generic Routing Encapsulation is an IP tunneling protocol: it encapsulates an inner packet within an outer IP packet so that traffic can traverse an IP network between tunnel endpoints. MPLS, by contrast, forwards traffic using short, fixed-format labels and establishes label-switched paths through an MPLS domain. At ingress, packets are classified into forwarding equivalence classes and assigned labels; each label-switching router uses the label to select the next hop and swaps, pops, or pushes labels as required. MPLS is commonly described as protocol-independent because it can carry IP and other network-layer protocols, but that capability comes from MPLS label encapsulation and label distribution/signaling, not from GRE. MPLS also supports traffic engineering by steering traffic onto explicitly selected label-switched paths, which can be calculated separately from ordinary destination-based IP forwarding decisions. The IETF MPLS architecture defines this label-based forwarding model and its relationship to conventional network-layer routing. See RFC 3031: Multiprotocol Label Switching Architecture and Cisco's MPLS overview.

QUESTION NO: 30

Which of the following statements best describes downstream unsolicited label distribution?

A. An LER may answer requests for label mappings immediately, without waiting for a label mapping from the next-hop.

B. An LER may advertise label mappings to all peers for which it might be a next-hop for a given FEC.

C. An LER propagates a label mapping downstream for a FEonly if it has a label mapping for the FEC next-hop.

D. An LER will only distribute a label for a FEC when it has been requested to do so.

ANSWER: B

Explanation:

"An LER may advertise label mappings to all peers for which it might be a next-hop for a given FEC." correctly describes downstream unsolicited label distribution. In this Label Distribution Protocol mode, a label switching router distributes a label mapping proactively to its peers once it has established a mapping for a forwarding equivalence class, rather than first waiting to receive a Label Request. The mapping is sent downstream-to-upstream: the router that is closer to the FEC destination advertises the label that an upstream peer should use when forwarding traffic toward that destination. This proactive behavior allows upstream routers to learn usable MPLS label bindings without issuing individual requests for each FEC and is the normal LDP label-distribution behavior used in many MPLS deployments. The qualification that the advertising router may be a next hop is important because the binding is relevant to peers that could forward traffic to that router for the applicable FEC. RFC 5036 defines downstream unsolicited distribution as the mode in which an LSR distributes label mappings to its peers when it establishes those mappings, without requiring prior requests. Nokia's MPLS documentation likewise describes LDP as advertising label bindings to peers for FECs associated with reachable destinations. See RFC 5036, Section 2.6.3 and Nokia MPLS Label Distribution Protocol documentation.

QUESTION NO: 31

Which of the following statements regarding DiffServ-TE LSPs is true?

A. The Maximum Allocation Model (MAM) provides the most efficient bandwidth utilization model

B. In the Maximum Allocation Model (MAM), each CT receives dedicated bandwidth constraints.

C. In the Russian Doll Model (RDM), a CT2 LSP automatically preempts a CT1 LSP.

D. Using the Russian Doll Mode! (RDM), a CT1 LSP may use unreserved CTO bandwidth.

ANSWER: B

Explanation:

In the Maximum Allocation Model (MAM), each Class-Type (CT) is assigned its own bandwidth constraint on a DiffServ-aware traffic-engineering link. These constraints partition the reservable link bandwidth by class type, so admission control for an LSP evaluates the bandwidth available within that LSP's applicable CT constraint. This provides predictable, class-specific resource allocation: bandwidth reserved for one CT is not ordinarily made available to another CT under the MAM sharing rules. Consequently, stating that each CT receives dedicated bandwidth constraints accurately describes the core MAM behavior.

DiffServ-aware MPLS Traffic Engineering extends RSVP-TE admission control with class types and bandwidth constraints so that a provider can apply differentiated resource-allocation policies to traffic-engineered LSPs. MAM is one of the bandwidth-allocation models defined for this purpose. The IETF's specification describes MAM as using a separate bandwidth constraint for each class type, with no sharing of bandwidth between CTs. See RFC 4124, Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering, especially its bandwidth-allocation-model definitions.

QUESTION NO: 33

Which of the following routing protocols are supported on the Alcatel-Lucent 7750 SR to enable constraint-based LSPs? (Choose two)

A. RIPv2-TE

B. CSPF

C. OSPF-TE

D. ISIS-TE

E. RSVP-TE

ANSWER: C D

Explanation:

OSPF-TE and **ISIS-TE** are the supported traffic-engineering extensions used by the 7750 SR to enable constraint-based MPLS LSP establishment. These link-state routing protocols advertise the additional TE attributes needed for path selection, such as reservable bandwidth, administrative groups, and TE link metrics. The router's constraint-based shortest-path-first computation uses this TE topology information to calculate an LSP path that satisfies the configured constraints before RSVP signaling establishes and maintains the LSP.

OSPF traffic-engineering capability is defined through opaque LSAs that distribute TE link-state information, while IS-IS traffic engineering distributes comparable information using extended IS reachability advertisements. Consequently, either protocol can provide the TE database from which the router derives an eligible constrained path. Nokia's SR OS MPLS documentation describes MPLS TE and CSPF operation for SR routers, and the underlying standards specify the respective TE extensions: Nokia SR OS documentation, RFC 3630: Traffic Engineering Extensions to OSPF, and RFC 5305: IS-IS Extensions for Traffic Engineering.

QUESTION NO: 34

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.

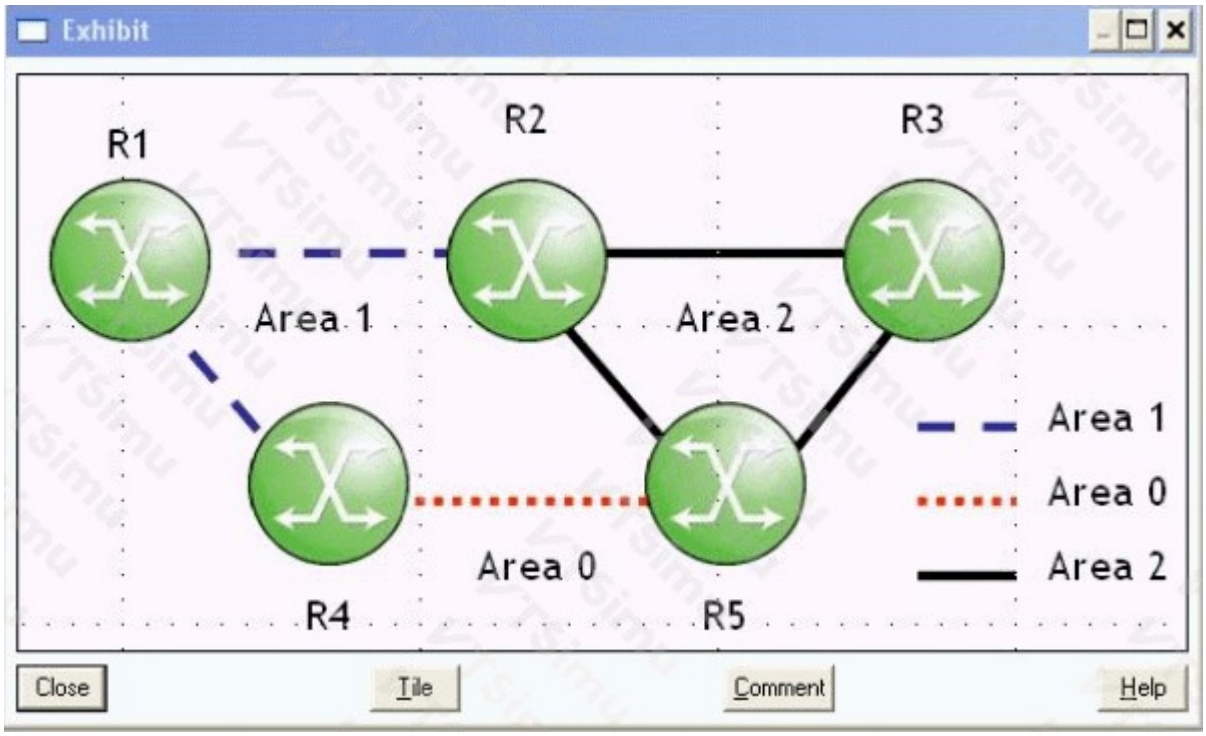
ANSWER: B E

Explanation:

It uses fewer label resources than a per-interface label space because a per-platform label space is shared by the entire LSR. The router can allocate one locally significant label for a given FEC and advertise that same label to all relevant LDP peers, instead of allocating a distinct label for that FEC on every interface. This avoids duplicate label allocation and correspondingly reduces consumption of the router's available label values and LFIB resources. A single label may be assigned to an FEC and used across all interfaces of the same router describes the defining characteristic of the per-platform model: label scope is the platform, rather than an individual interface. The receiving router associates the advertised label with the FEC independently of the outgoing interface used to reach the advertising LSR. LDP explicitly supports both per-platform and per-interface label spaces, and the LDP initialization exchange communicates the label-space identifier used by an LSR. Per-platform label space is the common model for packet-based MPLS forwarding because it provides a single, shared allocation domain. See the LDP specification in RFC 5036, especially its label-space definitions, and the MPLS architecture in RFC 3031.

QUESTION NO: 35

Click on the exhibit button below.



Consider the following:

To support LDP-over-RSVP tunnels between routers R1 and R3 over path R1-R4-R5-R3, with which routers will router R5 require targeted sessions? (Choose two)

A. Router R1

B. Router R4

C. Router R2

D. Router R3

ANSWER: B D

Explanation:

Router R5 requires targeted LDP sessions with **Router R4** and **Router R3**. In the specified RSVP-TE tunnel path, R5 is the transit LSR immediately downstream of R4 and immediately upstream of the tunnel egress, R3. LDP-over-RSVP uses targeted LDP adjacencies to exchange label bindings across RSVP-signaled tunnel segments, allowing the LDP forwarding relationship to follow the RSVP-TE path even where ordinary link-local LDP discovery is not the mechanism used for the service path. Consequently, R5 must establish the targeted LDP relationships needed on each adjacent portion of the stated tunnel path: toward R4 for the upstream segment and toward R3 for the downstream segment. These targeted adjacencies provide the label-distribution connectivity required for R5 to participate in forwarding traffic over the RSVP tunnel from R1 to R3. LDP targeted discovery and session establishment are defined by the LDP protocol specification, including the use of unicast Hellos to form a targeted adjacency; see RFC 5036. RSVP-TE signaling and label-switched path operation are specified in RFC 3209.