DUMPSARENA

Nokia Segment Routing Exam

Nokia 4A0-116

Version Demo

Total Demo Questions: 6

Total Premium Questions: 64

Buy Premium PDF

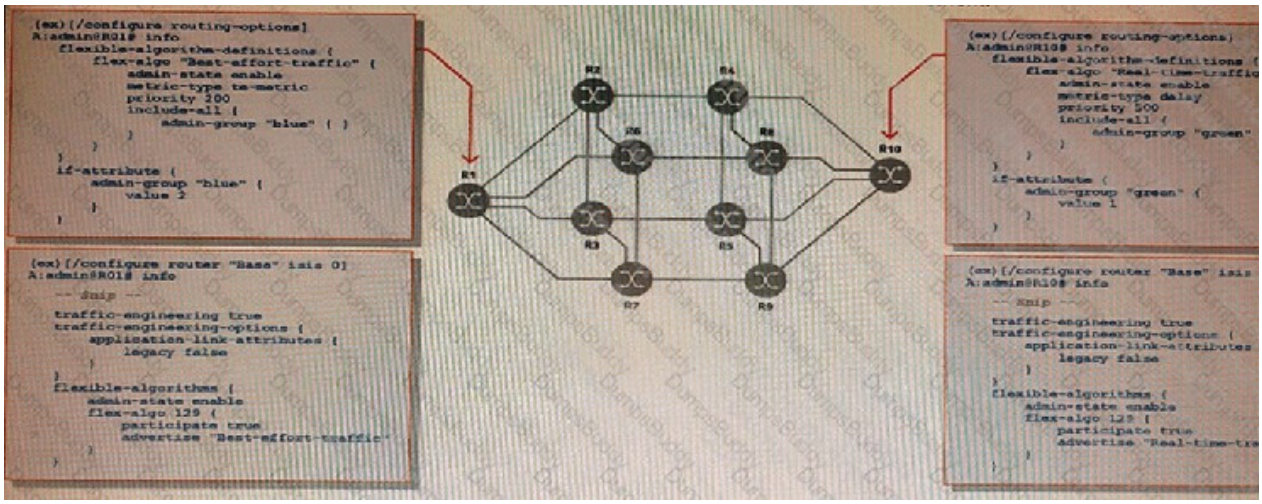
https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

Based on the configuration shown for routers R1 and R10, what valid flex-algo definitions exist in the network?



A. No valid flex-algo definitions, because of the conflicting parameters.

B. One valid flex-algo definition, using delay as the link metric and including green links.

C. One valid flex-algo definition, using te-metric as the link metric and including blue links.

D. Two valid flex-algo definitions, one uses te-metric as the link metric and includes blue links and the other uses delay as the link metric and includes green links.

ANSWER: D

Explanation:

Two valid flex-algo definitions, one uses te-metric as the link metric and includes blue links and the other uses delay as the link metric and includes green links. A Flex-Algo definition is identified by its algorithm identifier and is formed from consistently advertised calculation constraints, including the metric type and administrative-group constraints. Therefore, the network can carry separate valid algorithm definitions when each definition has its own consistent set of attributes. In the displayed configuration, the definition associated with blue links uses the TE metric, while the separate definition associated with green links uses delay as its metric. These are independent Flex-Algo computations rather than mutually exclusive settings within one definition. The resulting topology and shortest-path calculation for each algorithm are determined from links that meet that algorithm's administrative-group inclusion rule, using the specified metric for path calculation. This behavior allows the control plane to advertise and calculate distinct constrained topologies concurrently, then use the corresponding algorithm-specific Segment Routing identifiers for forwarding. The Flex-Algo framework and its metric and affinity constraints are specified in RFC 9350. Nokia's implementation and configuration model are described in the Nokia SR OS Flex-Algo documentation.

QUESTION NO: 2

Which of the following is an advantage of enabling label stack reduction for a CSPF-calculated path, as compared to using plain CSPF?

A. The number of hops included in the end-to-end path is reduced.

B. The path calculation can be delegated to an external path computation element (PCE).

C. There is a larger list of traffic-engineering constraint types that can be imposed on the LSP path.

D. The routers in the calculated path can take advantage of ECMP to better distribute the traffic load.

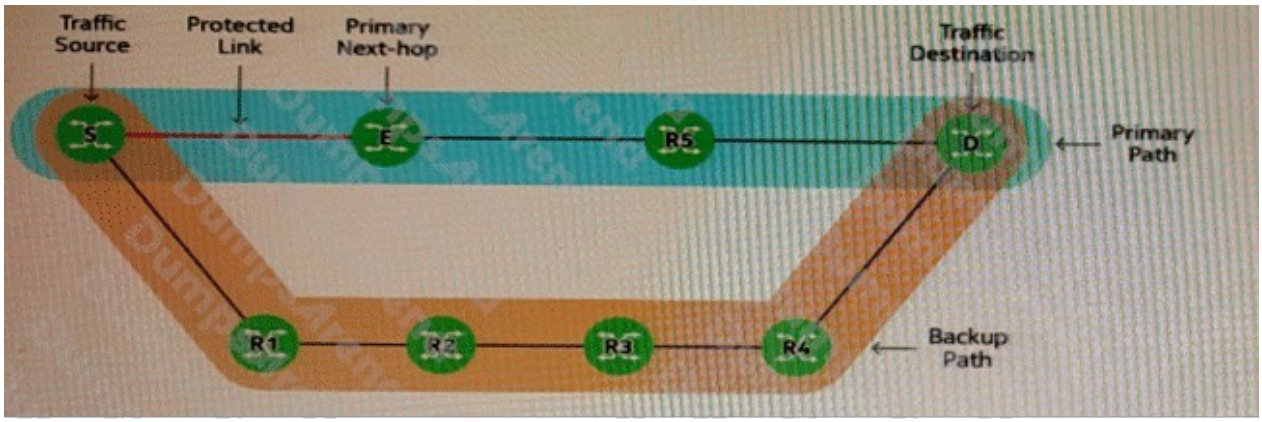
ANSWER: D

Explanation:

The routers in the calculated path can take advantage of ECMP to better distribute the traffic load. Label stack reduction optimizes how a CSPF-computed SR-TE path is encoded in the segment list. Rather than retaining labels that strictly prescribe every individual hop, it can replace an eligible portion of the computed path with a node segment that directs traffic toward a downstream node using the IGP shortest-path topology. This reduces the imposed label stack while retaining the intended traffic-engineering result where the shortened segment is valid. A key operational benefit is that, when multiple equal-cost shortest paths exist toward that downstream node, the forwarding routers can use their normal ECMP behavior across those equal-cost paths. Traffic can therefore be spread over available equal-cost next hops instead of being constrained to one explicitly encoded adjacency sequence. This combines a smaller MPLS label stack with better use of the underlay's equal-cost capacity. Nokia's segment-routing documentation describes SR path encoding and the use of node segments with IGP forwarding behavior; the architectural basis for node segments and ECMP follows the SR architecture defined in Nokia SR OS Segment Routing documentation and RFC 8402.

QUESTION NO: 3

For any of the LFA methods to work properly, there must be a router in the backup path that will forward traffic addressed to the destination without sending it back to the source. Which of the following statements describes the way R-LFA selects such a router?



A. It selects an immediate neighbor.

B. It selects a remote router, reachable through a tunnel defined by a single Node-SID.

C. It selects a remote router, reachable through a tunnel defined by a Node-SID and one or two Adjacency-SIDs.

D. It can select an immediate neighbor or a remote router, as long as the backup path coincides with the path that will become active after IGP reconvergence.

ANSWER: D

Explanation:

It can select an immediate neighbor or a remote router, as long as the backup path coincides with the path that will become active after IGP reconvergence. This captures the purpose of Remote Loop-Free Alternate: it extends conventional LFA protection beyond directly connected neighbors by finding a loop-free repair path to a suitable repair endpoint, commonly called a PQ node. The selected endpoint must be reachable from the point of local repair through the appropriate protected topology and must itself have a loop-free path toward the destination that does not return traffic to the repairing router. In an SR-enabled network, the repairing router can steer packets over the calculated repair path with a segment-routing policy or segment list, allowing the repair route to reach a remote endpoint where normal shortest-path forwarding toward the destination can continue. If a directly connected neighbor satisfies the same loop-free and post-convergence path requirements, it can also serve as the repair next hop. The essential selection criterion is therefore not whether the endpoint is necessarily local or remote, nor a fixed SID composition, but whether the precomputed repair path is loop-free and aligns with the forwarding path expected after IGP convergence. RFC 7490 defines remote LFA and its PQ-node-based repair model; RFC 8670 describes the segment-routing mechanisms used to encode steering instructions. See RFC 7490 and RFC 8670.

QUESTION NO: 4

Which of the following steps is NOT required when configuring IS-IS to support Segment Routing?

A. MPLS label range reserved for Segment Routing.

B. Enable interfaces used for Segment Routing.

C. The flooding scope of Segment Routing information.

D. The Segment Routing Global Block range.

ANSWER: B

Explanation:

Enable interfaces used for Segment Routing is the step that is not required as a distinct Segment Routing configuration action under IS-IS. Segment Routing for MPLS extends IS-IS by advertising Segment Routing capabilities, the Segment Routing Global Block, and prefix or adjacency Segment Identifiers through IS-IS TLVs. Once these extensions are enabled and advertised, IS-IS performs its normal link-state database synchronization and flooding processes to distribute the information throughout the applicable IS-IS level. There is no separate requirement to enable each IS-IS interface specifically for Segment Routing simply to make the Segment Routing information flood; interfaces already participating in the IS-IS adjacency and level carry the applicable Link State PDUs.

The operational configuration must still establish the local MPLS/SR label resources and the Segment Routing Global Block, because these define the label space from which globally significant Segment Identifiers are allocated and interpreted. The flooding scope must also be appropriate for the intended IS-IS topology, so that nodes which must calculate Segment Routing paths receive the advertised information. RFC 8667 defines the IS-IS extensions used to advertise Segment Routing information, including the SRGB and Segment Identifiers, while RFC 8402 describes the overall Segment Routing architecture and control-plane distribution model. See RFC 8667 and RFC 8402.

QUESTION NO: 5

Which of the following statements about primary and secondary SR-TE LSP paths is FALSE?

A. Only one LSP path forwards the traffic at any time,

B. Up to three paths can be configured for a given SR-TE LSP.

C. The primary path is always preferred over a secondary path.

D. Preference values can be configured for non-standby secondary paths.

ANSWER: B

Explanation:

“Up to three paths can be configured for a given SR-TE LSP” is false because three is not the SR OS maximum or defining limit for SR-TE LSP path configuration. SR OS supports one primary path and can support multiple secondary paths, subject to the applicable platform and software release limits. Consequently, describing three as the maximum number of configurable paths incorrectly imposes a fixed ceiling that does not represent the SR-TE LSP path model. The primary path provides the normal preferred forwarding path, while secondary paths provide protection or alternate forwarding behavior. Secondary-path behavior can also be influenced by attributes such as standby status and preference, allowing the router to select an eligible path according to the configured operational rules. This path architecture provides resiliency without treating a primary-plus-two-secondary arrangement as a universal maximum. Nokia’s Segment Routing documentation describes SR-TE configuration and operational concepts, including SR-TE LSP path behavior, in the Nokia SR OS documentation portal. Further SR OS routing and MPLS configuration material is available through the Nokia 7750 SR Router Configuration Guide.

QUESTION NO: 6

Which of the following is the primary function of the active segment in an SR-MPLS label stack?

A. It identifies the forwarding instruction currently processed by the router.

B. It reserves bandwidth on the next-hop interface.

C. It identifies the bottom label of the packet stack.

D. It specifies the return path toward the ingress router.

ANSWER: A

Explanation:

It identifies the forwarding instruction currently processed by the router. is correct. In an SR-MPLS packet, the top label is the active segment. The router processes that segment to determine the required forwarding action. For example, a Node-SID causes forwarding toward the identified node using the applicable IGP path, while an Adjacency-SID causes the owning router to use the specified outgoing adjacency. When the current segment is completed, the label is swapped or popped as appropriate, exposing the next segment in the stack. This segment-list processing model is defined in RFC 8402 and RFC 8660.