PassQueen
a.k.a dumpsqueen.com

Alcatel-Lucent Services Architecture

Alcatel-Lucent 4A0-104

Version Demo

Total Demo Questions: 44

Total Premium Questions: 444

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Volume A	76
Topic 2, Volume B	77
Topic 3, Volume C	83
Topic 4, Volume D	89
Topic 5, Volume E	99
Total	444

QUESTION NO: 1

Which of the following statements are TRUE about configuring a distributed E-pipe service?

(Choose two.)

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 E-pipe, the spoke-SDP binding carries the pseudowire identifier used to associate the remote spoke endpoint with the intended E-pipe. Consequently, the VC ID on each spoke-SDP must be explicitly entered as part of the service configuration. This ensures that the same pseudowire value is signaled and recognized at the relevant service endpoints, rather than relying on an implicit local service value.

The explicitly configured spoke-SDP VC ID may use the same numerical value as the service ID. Using matching values is a common administrative convention because it makes the relationship between the local E-pipe service and its spoke pseudowire easy to identify during provisioning and troubleshooting. Matching the service ID does not remove the requirement to configure the spoke-SDP VC ID explicitly; it simply selects the service ID as the configured VC-ID value.

Nokia's SR OS Layer 2 Services documentation describes E-pipe service constructs and SDP bindings, including the role of pseudowire/VC identifiers in service signaling. See the Nokia SR OS documentation portal and the SR OS Layer 2 Services Guide for E-pipe and SDP-binding configuration details.

QUESTION NO: 3

Which of the following about Hierarchical VPLS is FALSE?

A. They allow a VPLS to scale more easily.

B. They use mesh SDPs to connect smaller VPLSes together.

C. They allow a service to span multiple metro networks.

D. They simplify the configuration required when adding new PEs.

ANSWER: B

Explanation:

"They use mesh SDPs to connect smaller VPLSes together." is false because Hierarchical VPLS (H-VPLS) is built by separating the VPLS service into hierarchical components and using spoke connectivity to attach an access or edge VPLS to an aggregation/core VPLS. In an H-VPLS design, the spoke pseudowire/SDP provides the hierarchical attachment between these VPLS instances. Mesh pseudowires or mesh SDPs instead provide the full-mesh relationship among peer provider-edge nodes within a given VPLS forwarding domain; they are not the mechanism used to hierarchically join smaller VPLS instances.

This distinction is fundamental to the scaling model of H-VPLS. By aggregating multiple edge domains through spoke connections into a core domain, the design limits the number of full-mesh relationships required in the network. The IETF's VPLS specification describes the distinction between mesh and spoke pseudowires and explains that an H-VPLS

arrangement uses spoke pseudowires to interconnect VPLS instances while preserving appropriate split-horizon forwarding behavior. See RFC 4762: Virtual Private LAN Service Using LDP Signaling, especially its hierarchical VPLS discussion.

QUESTION NO: 4

Which of the following must be configured when configuring an Epipe service? (Choose 4)

A. Configuring an Epipe service with a customer ID

B. Defining a SAP or SAPs.

C. Binding to an SDP if it is a distributed service.

D. Enabling the service using the no shutdown command.

E. Creating and assigning an IP filter to the SAP.

F. Configuring QoS parameters

ANSWER: A B C D

Explanation:

An Epipe is a point-to-point Layer 2 service in Nokia SR OS and needs the basic service hierarchy, attachment point, and operational activation to carry customer traffic. **Configuring an Epipe service with a customer ID** associates the service with an existing customer context; the customer identifier is part of the service creation syntax and provides the administrative ownership for the service. **Defining a SAP or SAPs.** supplies the local service access point endpoint or endpoints through which frames enter and leave the Epipe. For an Epipe whose endpoints are on separate routers, **Binding to an SDP if it is a distributed service.** provides the service-distribution path to the remote PE; the Epipe uses an SDP binding to extend the service across the provider network. Finally, **Enabling the service using the no shutdown command.** is required because newly created service objects are administratively shut down by default. The service must be enabled after its required endpoint configuration is in place before it can become operational and forward traffic. Nokia's SR OS service documentation describes Epipe as a point-to-point service and documents the required SAP and SDP service components: Nokia SR OS Epipe service overview and Nokia SR OS services overview.

QUESTION NO: 5

Which of the following statements best describes a route target?

A. A route target is a BGP extended community used to identify the VRF table for a prefix at the receiving PE

B. A route target is a mechanism from which VPRN controls the distribution of VPN routing information.

C. Route target attributes are earned in a MP-BGP as attributes of the route.

D. A route target is used by the PE router to identify the VRF that a VPN-IPv4 prefix is associated with.

E. All of the above statements describe a route target

F. None of the above statements describe a route target

ANSWER: E

Explanation:

All of the above statements describe a route target. In an RFC 4364 BGP/MPLS IP VPN, a route target is a BGP extended-community attribute attached to VPN-IPv4 routes advertised through Multiprotocol BGP. It provides the VPN membership and route-distribution policy that determines which VRFs import a received route. A PE examines the route-target communities on a received VPN route and imports the route into the local VRF or VRFs configured with matching import route targets. This lets a VPRN control distribution of customer VPN routing information independently of the route distinguisher, which makes otherwise overlapping customer prefixes unique in VPN-IPv4 address space. Route targets are

carried as BGP path attributes, so they travel with the MP-BGP VPN route between PEs. Their import/export use also associates a VPN-IPv4 prefix with the appropriate receiving VRF context, including deployments where route leaking or hub-and-spoke connectivity requires routes to be imported by more than one VRF. The defining standards description is available in RFC 4364, BGP/MPLS IP Virtual Private Networks; the extended-community encoding used by route targets is specified in RFC 4360, BGP Extended Communities Attribute.

QUESTION NO: 6

Which of the following best describes the flooding of traffic on a PE when unknown 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 an Alcatel-Lucent/Nokia VPLS service, a mesh SDP represents the meshed provider-network portion of the service. When a PE receives an unknown destination frame from a mesh SDP, it performs data-plane flooding toward local attachment circuits, represented by SAPs, and toward spoke SDPs. This behavior ensures that the frame can reach customer sites connected locally or through hierarchical VPLS spoke connections while retaining the split-horizon rule used between mesh SDPs.

Mesh SDPs form the full-mesh connectivity between PE routers. The service forwarding model treats traffic arriving from that mesh domain as ineligible for replication to other mesh SDPs on the same PE. Instead, flooding is directed to the service's access and spoke-facing destinations. This prevents unnecessary replication and looping within the provider mesh while allowing unknown unicast, broadcast, or other floodable traffic to reach eligible service endpoints. Nokia's service architecture documentation describes the VPLS forwarding and split-horizon behavior for mesh and spoke SDP bindings; see the Nokia 7750 SR Services Overview and the Nokia Layer 2 Services Guide.

QUESTION NO: 7

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

A. config>service# epipe service-id enable

B. config>service# vpws service-id enable

C. config>service# service-id no shutdown

D. config>port# port-id no shutdown

E. config>service>epipe# no shutdown

ANSWER: E

Explanation:

An E-pipe service is administratively enabled with the `no shutdown` command in the E-pipe service configuration context. An E-pipe is a Nokia SR OS point-to-point Layer 2 service; after the service has been created and its required service access points or spoke SDPs have been configured, `no shutdown` changes its administrative state to enabled. The command is applied beneath the specific E-pipe instance, represented by the CLI prompt `config>service>epipe#`. Service IDs identify the individual E-pipe instance when entering or creating its configuration, but the administrative enable

action itself is performed by `no shutdown` at that service level. This follows the standard SR OS service model, in which services are created in an administratively disabled state until explicitly enabled. For Nokia SR OS service concepts and E-pipe configuration guidance, see the Nokia SR OS documentation portal and the SR OS Services Overview documentation.

QUESTION NO: 8

Which of the following about VPRN VRF tables is FALSE?

A. A PE uses the same VRF table for all the VPRN services on it.

B. A PE uses a VRF to maintain forwarding information.

C. Information for VRFs can be advertised between PEs by MP-BGP.

D. A PE uses a VRF to route traffic to CEs.

ANSWER: A

Explanation:

A PE uses the same VRF table for all the VPRN services on it. is the false statement. In a VPRN implementation, each customer VPRN service requires its own logically separate VPN routing and forwarding context, commonly called a VRF. This separation is fundamental: the VRF holds the routing information, interfaces, forwarding state, and route-import/export policy applicable to that particular customer VPN. Consequently, a PE can support many VPRN services while keeping their customer address spaces and forwarding decisions isolated, including cases where multiple customers use overlapping IP prefixes.

The PE associates CE-facing interfaces with the appropriate service VRF and uses that VRF's route table when forwarding customer traffic. VPN routes may be distributed between PEs using Multiprotocol BGP with VPN-specific route attributes, allowing the receiving PE to install routes in the intended VRF according to the configured import policy. Sharing one VRF table among all VPRN services would remove this required per-customer routing isolation. The VPN/VRF model is defined in RFC 4364; Nokia's service documentation is available through the Nokia Documentation Center.

QUESTION NO: 9

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

(Choose two.)

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 on which sap or sdp traffic should be replicated

ANSWER: B C

Explanation:

A mirror service requires a mirror source to define the traffic that is to be copied and a mirror destination to define where those copies are sent. "A mirror source to indicate which ports to monitor" is correct because the source configuration identifies the physical port or service traffic to be observed, along with applicable direction and filtering criteria. The mirror function creates copies of packets matching that source definition; it does not alter the original forwarding path.

"A mirror destination to indicate on which sap or sdp traffic should be replicated" is also correct because the destination identifies the egress delivery point for the copied packets. In SR OS, a mirror destination can use a SAP or an SDP to carry mirrored traffic toward an analyzer, collector, or remote monitoring location. This separation between source selection and destination delivery is fundamental to configuring local or remote port/service mirroring. Nokia's SR OS service documentation describes mirror services as mechanisms for copying selected traffic to a defined destination, while the SR

OS CLI documentation provides the mirror source and destination configuration hierarchy. See the Nokia SR OS Mirror Services documentation and the Nokia SR OS Services Overview.

QUESTION NO: 10

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, a spoke-SDP represents an endpoint-facing pseudowire binding and requires an explicitly provisioned `vc-id`. Explicitly setting the value is important because the VC identifier is the pseudowire identifier used to associate the two ends of the service across the SDP. The `vc-id` is a service-specific value; configuring it directly avoids reliance on an implicit value and ensures that the remote pseudowire endpoint is provisioned with the corresponding identifier.

The explicitly configured spoke-SDP `vc-id` may be set equal to the Epipe service ID. This is a valid and common administrative convention that simplifies service deployment and troubleshooting, because the local service identifier and the transport pseudowire identifier are then easy to correlate. Equality is not an inherent protocol requirement: the essential requirement is that the explicitly provisioned VC identifiers correspond at the pseudowire endpoints. This behavior is consistent with the SR OS service model for Epipe and spoke-SDP bindings, and with the generalized pseudowire control-word and identifier framework described by the IETF. See the Nokia SR OS documentation portal and RFC 4447, Pseudowire Setup and Maintenance Using LDP.

QUESTION NO: 11

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. In SR OS, creating a mirror destination does not automatically enable packet transmission to its configured egress port; it remains administratively disabled until the operator explicitly applies the no shutdown command. This default is intentional, since it prevents mirrored traffic from being sent before the destination, encapsulation settings where applicable, and receiving monitoring equipment have been fully validated.

When an Alcatel-Lucent Service Router is rebooted the mirror source must be re-configured. Mirror-source entries are operational monitoring definitions rather than persistent configuration objects. They are therefore not retained across a system reboot, even though the mirror destination configuration itself is retained in the saved router configuration. An operator must recreate the required source associations after the router has restarted before traffic is mirrored again. This distinction is important for operational procedures: saving the configuration preserves the destination setup, but does not eliminate the need to restore mirror-source monitoring after a reboot. Nokia's SR OS documentation is available through the SR OS documentation portal; platform and product context is provided on the Nokia 7750 Service Router page.

QUESTION NO: 12

Consider an E-pipe service configured with SAP 1/1/1:5 and SAP 1/1/1:6. If the SAP 1/1/1:5 accepts an ingress frame, which of the following is TRUE?

A. The frame egresses from 1/1/1:6 with no VLAN tag.

B. The frame egresses from 1/1/1:6 with VLAN tag 5.

C. The frame egresses from 1/1/1:6 with VLAN tag 6.

D. The frame egresses from 1/1/1:6 with an inner VLAN tag of 5 and an outer VLAN tag of 6.

ANSWER: C

Explanation:

The frame egresses from 1/1/1:6 with VLAN tag 6. In an SR OS E-pipe, the service provides point-to-point Layer 2 forwarding between its two Service Access Points (SAPs). The VLAN identifier in a dot1q SAP definition identifies the customer traffic that the SAP accepts on ingress and the VLAN encapsulation that applies at that SAP. Therefore, SAP 1/1/1:5 accepts traffic associated with VLAN 5 and maps it into the E-pipe service. Once the frame is forwarded across the service to the opposite SAP, egress processing applies the encapsulation associated with SAP 1/1/1:6. The outgoing Ethernet frame is consequently transmitted on port 1/1/1 carrying VLAN ID 6. This behavior allows an E-pipe to connect customer attachment circuits that use different VLAN IDs while retaining a simple point-to-point service model. The service forwards the Ethernet payload through the service, while each SAP defines the appropriate local attachment encapsulation at its respective edge. Nokia's SR OS Layer 2 service documentation describes E-pipe services and SAP-based service access in the Layer 2 Services guide; SAP encapsulation concepts are also covered in the Services Overview guide.

QUESTION NO: 13

What needs to be configured before a SAP appears? (Choose two.)

A. Configuring the physical port in access mode.

B. Configuring the physical port in network mode.

C. Enabling the port using the no shutdown command.

D. Ports are enabled by default, thus no need to use the no shutdown command.

ANSWER: A C

Explanation:

Configuring the physical port in access mode and enabling the port using the `no shutdown` command are required before a Service Access Point (SAP) can become available operationally. In SR OS, a SAP is the service-facing attachment to an access interface or access encapsulation, so the underlying port must first be designated as an access port. This establishes that the port is available for customer/service access rather than for network-facing routing functions.

The port must also be administratively enabled with `no shutdown`. SR OS uses an administrative shutdown state on ports; removing that shutdown allows the interface to progress toward an operational state, subject to physical link conditions. A SAP depends on the state of its underlying port, meaning a port that remains administratively disabled cannot provide an

active service attachment. Once the port is configured for access use and enabled, the SAP configured under the applicable service can be brought up when its remaining service and link requirements are satisfied. Nokia's SR OS documentation describes port administrative configuration and the access-versus-network port role in its interface and services material: SR OS port configuration and SR OS service model.

QUESTION NO: 14

Which of the following MUST be identical for a service between two PEs?

A. Customer ID

B. SDP ID

C. Service ID

D. VC ID

ANSWER: D

Explanation:

VC ID must be identical at the two provider-edge endpoints for the same pseudowire/service connection. In an SR OS Layer 2 VPN service, the VC ID is the value used to identify the virtual circuit over the service distribution path between the participating PEs. Each endpoint uses this common identifier to associate the remotely signaled or configured pseudowire with the intended service instance. Consistent VC-ID configuration therefore ensures that both PEs bind the same customer service across the packet network.

This requirement follows the pseudowire architecture: a pseudowire identifier provides an endpoint-to-endpoint identification of the emulated service, allowing the control plane and forwarding plane to establish the correct service relationship. For VPLS and related Layer 2 VPN deployments, the virtual-circuit identifier is carried or matched as part of pseudowire setup and must correspond at both ends. Nokia SR OS documentation describes service distribution paths and pseudowire-based service interconnection in its Layer 2 Services documentation. The underlying PWE3 architecture and pseudowire identifiers are also defined by the IETF in RFC 3985.

QUESTION NO: 15

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 is a Layer 3 Internet Enhanced Service, so its service interfaces participate in IP routing. Consequently, **IES has a configurable IP-MTU**: the IP MTU is an interface-level Layer 3 parameter that controls the maximum IP packet size accepted or transmitted by the routed interface. This is distinct from the service MTU considerations used for Layer 2 VPWS and VPLS forwarding.

Also, **An IES interface has an IP address assigned to it**. Assigning an IPv4 or IPv6 address enables the interface to act as a routed gateway and allows the router to forward traffic using its routing table. IES interfaces use the router's base routing context, whereas Layer 2 services such as VPWS and VPLS provide Ethernet transport/bridging behavior rather than an IP-routed interface with an assigned gateway address. Nokia's service documentation describes IES as a routed IP service and

documents IP interface attributes, including addressing and MTU configuration. See the Nokia IES overview and the Nokia IES configuration guide.

QUESTION NO: 16

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 E-pipe service on Nokia/Alcatel-Lucent SR OS. An E-pipe is a point-to-point Layer 2 VPN service, created in the service configuration context and identified by its service ID. Like other SR OS services, a newly created E-pipe is administratively shut down by default until the `no shutdown` command is applied within that service's configuration context. Entering this command changes the service's administrative state to up, allowing the router to activate the service when its required service access points, SDP bindings, and operational resources are available. The command is therefore entered as part of the E-pipe creation/configuration hierarchy: first select or create the E-pipe using `epipe <service-id>`, then issue `no shutdown`. Nokia's SR OS service documentation describes E-pipe as an Ethernet point-to-point service and documents the administrative `shutdown`/`no shutdown` control used for service activation. See the Nokia SR OS Epipe service documentation and the Nokia SR OS service command reference.

QUESTION NO: 17

What is the minimum Ethernet physical network MTU required for a Layer 2 service on an MPLS encapsulated SDP with a service MTU of 1500?

A. 1500

B. 1508

C. 1514

D. 1522

ANSWER: D

Explanation:

The minimum Ethernet physical network MTU is **1522**. A Layer 2 service MTU of 1500 represents the customer payload size supported by the service. When that service is carried through an MPLS-encapsulated SDP, the provider network must also accommodate the Ethernet framing header and the MPLS label stack used to transport the service. For the basic MPLS pseudowire forwarding case, this requires 14 additional bytes for the Ethernet header and 8 bytes for two 4-byte MPLS label-stack entries: one transport/tunnel label and one service (VC) label. Therefore, the physical Ethernet frame size that the network interface must support is 1500 + 14 + 8 = 1522 bytes. Configuring the network MTU at this value prevents frames at the configured service MTU from being discarded because the MPLS encapsulation causes their on-wire size to exceed the underlying Ethernet capability. MPLS label-stack entries are defined as 4 octets each by RFC 3032. Nokia's SR OS service documentation also describes the need to account for service and transport encapsulation when planning MTU values; see the Nokia SR OS Services Overview.

QUESTION NO: 19

Click the exhibit.

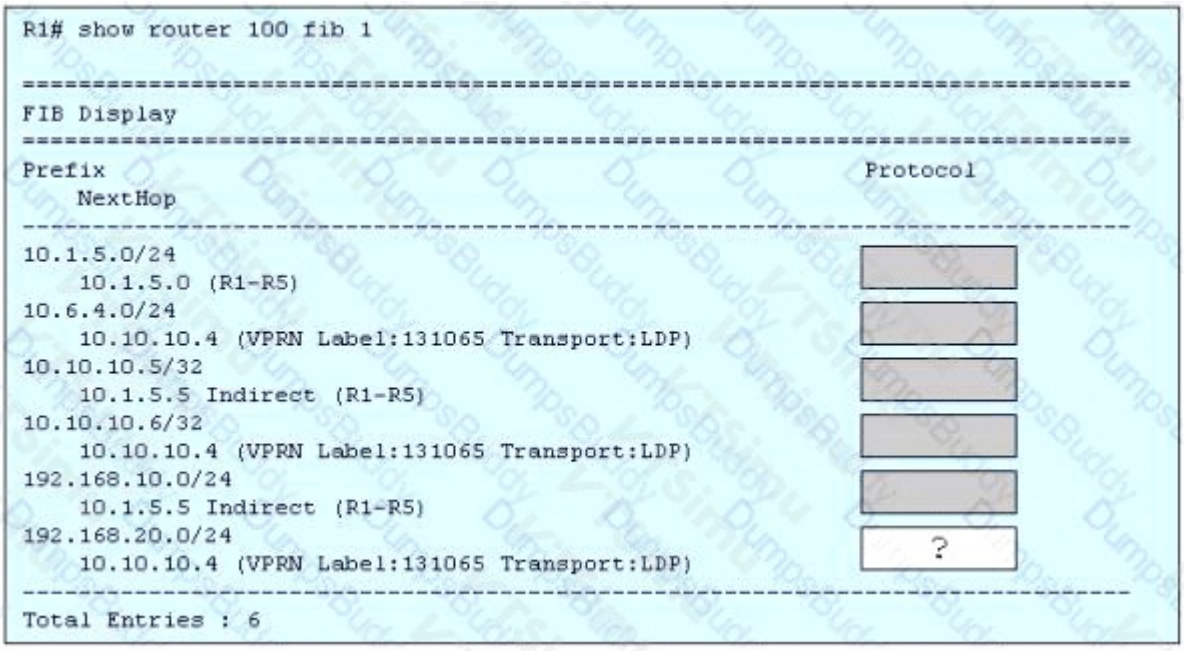

The output is from PE router R1about VPRN service 100.From which protocol did router R1learn the prefix 192.168.20.0/24?

A. LDP

B. BGP

C. MP-BGP

D. MPLS

ANSWER: C

Explanation:

MP-BGP is correct because VPRN routes learned from a remote PE are distributed as VPN-IPv4 routes through Multiprotocol BGP. In an RFC 4364-based Layer 3 VPN, the originating PE combines the customer IPv4 prefix with its route distinguisher and advertises that VPN route using the BGP multiprotocol extensions. The route is also associated with one or more route-target extended communities, which control which VPRN instances import it. On the receiving PE, the route-target policy for VPRN service 100 causes the matching VPN route to be imported into that service's routing table, where it appears as the customer prefix 192.168.20.0/24. The MPLS transport label used to reach the remote PE and the VPN service label are installed with the route for forwarding, but the VPN route advertisement and control-plane learning mechanism is MP-BGP. This separation of MP-BGP VPN route distribution from MPLS packet forwarding is fundamental to Nokia SR OS VPRN operation. See Nokia's VPRN overview and the VPN route-distribution model in RFC 4364.

QUESTION NO: 20

Which of the following about VPRNs is FALSE?

A. A VPRN is a Layer 3 service.

B. A PE maintains a separate forwarding table for each VPRN.

C. The service provider manages the customers IP addressing.

D. Multiple VPRNs can be deployed over the same IP/MPLS network infrastructure.

ANSWER: C

Explanation:

The statement *"The service provider manages the customers IP addressing."* is false because a VPRN is designed to provide customer traffic separation and private Layer 3 routing over a shared provider MPLS/IP backbone without requiring the provider to own or administer the customer's addressing plan. Each customer retains control of its own IP address space, including the use of private addresses, overlapping addresses with other customers, and its internal addressing policies. The provider's responsibility is to configure the VPRN service, attach customer-facing interfaces, and maintain the service-specific routing and forwarding context on the provider edge router. Customer routes are distributed between relevant sites through the VPN architecture, but that route transport does not transfer ownership or management of the customer's IP numbering scheme to the provider. This separation is a central VPN characteristic: multiple independently administered customer networks can operate securely across the same service-provider infrastructure. RFC 4364 describes how VPNs use distinct routing and forwarding contexts to keep customer networks separate while carrying their traffic through the provider network. See RFC 4364: BGP/MPLS IP Virtual Private Networks and the Nokia documentation portal for SR OS service documentation.

QUESTION NO: 21

Which of the following commands is used to configure the encapsulation type of an Ethernet port?

A. configure ethernetencap-typedot1q|null|qinq port port-id

B. configure service port port-id ethemetencap-typedot1q|null|qinq create

C. configure port port-id ethernet encap-type dot1q|null|qinq

D. configure service service-id Ethernet encap-type dot1q|null|qinq create

ANSWER: C

Explanation:

The command `configure port port-id ethernet encap-type dot1q|null|qinq` is used from the port configuration context to define the Ethernet encapsulation accepted on the physical port. In SR OS, the `ethernet` branch under `configure port` contains Ethernet-specific port attributes, and `encap-type` selects the framing model. The available values identify whether the port uses standard IEEE 802.1Q VLAN tagging, no VLAN encapsulation, or stacked VLAN tags using QinQ. This is a port-level setting, so it must be configured before services and SAPs relying on the corresponding encapsulation can be provisioned consistently. The port identifier in the command is replaced with the actual physical port name, such as `1/1/1`. For example, configuring `configure port 1/1/1 ethernet encap-type qinq` prepares that Ethernet port for double-tagged QinQ traffic. Nokia's SR OS Layer 2 services documentation describes Ethernet port encapsulation and its relationship to service access points, while the CLI documentation provides the configuration hierarchy and command syntax. See Nokia SR OS Ethernet encapsulation documentation and Nokia SR OS port configuration documentation.

QUESTION NO: 22

Which of the following comparisons between E-pipes and VPLSes is FALSE?

A. They both function as a Layer 2 switch from the service provider's perspective.

B. They both support SAP encapsulations of null, dot1Q and Q-in-Q.

C. They both use SAPs as the demarcation point between the customer and the provider.

D. MAC addresses are learned by VPLSes, but not by E-pipes.

ANSWER: A

Explanation:

“They both function as a Layer 2 switch from the service provider's perspective.” is the false comparison. An E-pipe is a point-to-point Ethernet service: it binds two service access points into a logical Ethernet connection and forwards customer traffic between those endpoints. It is therefore comparable to a virtual Ethernet wire rather than a multipoint switching domain. The service does not need to learn customer MAC addresses to choose among multiple customer-facing destinations, because the E-pipe has only one remote endpoint for the service.

A VPLS, in contrast, provides multipoint Layer 2 connectivity. It creates a virtual LAN across provider transport infrastructure and performs Ethernet bridging behavior, including MAC-address learning and forwarding decisions among local and remote attachment points. Thus, VPLS is the service that acts as a Layer 2 switch/bridge from the provider-service perspective, while an E-pipe is a point-to-point service. Nokia’s Layer 2 service documentation describes Epipe as point-to-point and VPLS as multipoint Ethernet service constructs; see the Nokia SR OS Layer 2 Services Overview and the Nokia SR OS VPLS documentation.

QUESTION NO: 23

Which of the following about VPRN labels is TRUE?

A. The core P routers will swap the service label at each hop.

B. The CE pushes the VPRN service label on to the packet and then forwards it to the ingress PE.

C. The PE swaps the transport label from the CE and then forwards the packet.

D. At the egress PE, the transport and service labels are popped.

ANSWER: D

Explanation:

At the egress PE, the transport and service labels are popped.

is correct because a VPRN packet crossing an MPLS provider network normally carries a two-label stack. The outer transport (or tunnel) label provides reachability across the provider core to the egress PE. Once the packet reaches that PE, the transport-label forwarding operation has completed; depending on the configured behavior, that label may already have been removed by penultimate-hop popping or is removed at the egress PE. The remaining inner VPN/service label identifies the destination VPRN context and is used by the egress PE to select the appropriate VRF and customer-facing forwarding treatment. Before the packet is sent toward the CE as a native customer IP packet, the service label is removed as well. Thus, the egress PE terminates the MPLS encapsulation and forwards the packet in the correct VPN routing context. This follows the BGP/MPLS IP VPN label-stack model specified in RFC 4364, which describes the inner VPN label and outer tunnel label. Nokia’s MPLS documentation also describes the use and disposition of transport and service labels in VPN forwarding: Nokia 7750 SR MPLS overview.

QUESTION NO: 24

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 the MPLS transport/signaling method through either `ldp` or `lsp lsp-name`. The `ldp` command associates the SDP with Label Distribution Protocol signaling. LDP distributes MPLS label bindings

dynamically between LSRs, allowing the SDP to use the LDP-established transport path to its far-end router. This is appropriate where the network uses conventional LDP-based MPLS forwarding.

The `lsp lsp-name` command associates the SDP with a named MPLS Label Switched Path. The named LSP is an RSVP-TE signaled path, allowing the service provider to use an explicitly configured traffic-engineered MPLS tunnel for SDP transport. Therefore, these two commands represent the available MPLS signaling selections at the SDP MPLS configuration level: dynamically signaled LDP transport or a specified RSVP-TE LSP. Nokia's SR OS documentation is available through the Nokia SR OS documentation portal. The protocol behavior underlying LDP is defined in RFC 5036.

QUESTION NO: 25

Fragmentation is supported on a Layer 2 service

A. TRUE

B. FALSE

ANSWER: B

Explanation:

FALSE is correct. A Layer 2 service transports customer Ethernet frames transparently within the configured service MTU; it does not act as an IP layer endpoint that can divide an oversized IP packet into smaller IP fragments. Fragmentation is a Layer 3 IP function, performed only when the packet is an IP packet and the applicable IP rules allow it. A Layer 2 Ethernet service instead has to accommodate the entire customer frame plus the service's encapsulation overhead, such as MPLS transport labels and pseudowire encapsulation, within the available path MTU.

Consequently, proper service design requires configuring a customer/service MTU that fits the underlying transport MTU after all required overhead is considered. Frames that exceed the supported Layer 2 service or transport MTU cannot be made acceptable through normal IP fragmentation by the Layer 2 service itself. This preserves the transparent, frame-oriented nature of Ethernet and other Layer 2 VPN services. The Ethernet pseudowire specification defines an explicit pseudowire MTU parameter for this purpose, while IP fragmentation behavior is specified separately by the IP protocol. See RFC 4448, Encapsulation Methods for Transport of Ethernet over MPLS Networks and RFC 791, Internet Protocol.

QUESTION NO: 26

Click on the exhibit button below.

Assuming the customer only needs to monitor traffic from PE3 to the customer site, what two items are missing in the configuration of PE5, the Alcatel-Lucent 7750 Service Router in the Head Office? (Assume that the IP/MPLS network structure between the local and remote PEs is in place.)

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

B. There needs to be a mirror source configuration on PE5 with the "remote-source" command.

C. This is not possible since the protocol analyzer is located at the Central Office. It needs to be at the customer site.

D. There needs to be a configuration command to put sap 1/1/3 on the mirror destination of PE5.

E. There needs to be a configuration command to put sap 1/1/3 on the mirror source of PE5

ANSWER: A D

Explanation:

Remote mirroring on the 7750 SR separates packet capture from the point at which the monitored traffic enters or leaves the service. PE3 can identify and send the required customer-bound traffic across the IP/MPLS network, while PE5 acts as the remote mirror destination at the Head Office, where the protocol analyzer is connected. For PE5 to accept mirrored packets delivered from the remote router, its mirror-destination context must be explicitly enabled as a remote source by using the

`remote-source` command. This identifies the destination as receiving encapsulated remote-mirror traffic rather than locally replicated traffic only. In addition, the destination mirror must specify the local egress attachment toward the analyzer. Adding SAP `1/1/3` under the mirror destination provides the physical/service access point through which PE5 transmits the copied packets to the monitoring device. Together, these elements complete the receiving side of the remote-mirroring design without requiring the analyzer to be located at the customer site. Nokia's SR OS documentation library includes the 7750 SR mirroring configuration material and command references: Nokia SR OS documentation. Product context for the platform is available at Nokia 7750 Service Router.

QUESTION NO: 29

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

A. VPRN is an IP service that connects multiple sites in a single routed domain over the providermanaged 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 an IP/MPLS Layer 3 VPN in which customer sites can attach to different provider-edge routers throughout the provider network. The PE routers use the provider MPLS/IP backbone and VPN route distribution mechanisms to make the geographically separate sites part of the same customer routing domain; there is no requirement for all sites to terminate on one PE. At the service edge, a customer-edge router normally forms its routing adjacency with its directly connected PE router, using a supported PE–CE routing method such as static routing, BGP, OSPF, or RIP. Customer routers at separate sites do not need to establish direct routing-peer relationships with one another. Instead, the provider network exchanges and imports the customer VPN routes into the appropriate VPRN routing instances, while the MPLS core transports traffic between the relevant PE routers. This architecture provides the multi-site routed VPN behavior while keeping core MPLS details hidden from the customer. See Nokia's VPRN service overview and RFC 4364: BGP/MPLS IP Virtual Private Networks.

QUESTION NO: 30

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

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 on a Nokia 7750 SR identifies customer traffic using the VLAN identifier in the IEEE 802.1Q header, or by the absence of that header for untagged traffic. The SAP `1/1/1:400` specifically matches customer frames carrying VLAN ID 400, so it can forward that VLAN's traffic into the service to which the SAP belongs. This is the normal use of an explicitly specified dot1q SAP value: the configured VLAN value is the customer encapsulation that identifies frames for the service.

The SAP `1/1/1:0` is the null-encapsulation SAP value for dot1q operation. It handles untagged customer frames; Nokia SR OS also treats frames bearing VLAN ID 0 as untagged for this SAP-matching purpose. VLAN 0 is commonly used for priority-tagged Ethernet frames, where the priority information is present but there is no customer VLAN membership indicated by a nonzero VLAN ID. Consequently, both the explicit VLAN 400 SAP and the null-VLAN SAP behavior stated in the question are valid. Nokia's SR OS service documentation describes SAP encapsulation and VLAN matching behavior in the 7750 SR Services Overview; the complete product documentation library is available from Nokia SR OS documentation.

QUESTION NO: 33

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

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 LDP-signaled service relies on an LDP control-plane relationship between the provider-edge routers. Therefore, *LDP must be enabled on both PE routers*; this permits the peers to establish the targeted LDP session used to exchange pseudowire label-mapping messages. The service endpoints must also identify the same pseudowire, which is why *The VC-ID must match between the two PE routers in both directions* is required. The VC-ID is carried in the pseudowire FEC and allows each provider-edge router to associate the received label mapping with the intended service binding.

In addition, *The transport tunnel must be operationally up between the two PE routers in both directions*. The service distribution path depends on usable transport connectivity between the endpoints; when the SDP transport is operational, the service can establish its forwarding relationship and use the exchanged service labels. These conditions collectively provide the LDP signaling mechanism, a common pseudowire identity, and an available transport path required for an operational LDP-signaled service. Nokia's Layer 2 Services documentation describes SDPs and LDP-signaled service operation: Nokia 7750 SR Layer 2 Services Guide. General MPLS LDP label-distribution behavior is specified in RFC 5036.

QUESTION NO: 34

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

A. dot1q

B. null

C. qinq

D. atm

E. ipcp

F. bcp-null

ANSWER: A C D

Explanation:

dot1q, **qinq**, and **atm** support multiple services on a physical port because each provides a service-demultiplexing identifier that can distinguish traffic belonging to separate SAPs. With dot1q encapsulation, the VLAN ID in the IEEE 802.1Q tag identifies the individual service, allowing numerous VLAN-based SAPs to share one Ethernet port. QinQ extends this service identification model by using an outer and inner VLAN-tag context, supporting service separation for stacked-tag traffic. For ATM encapsulation, the virtual-path and virtual-channel identifiers provide the distinction required to associate traffic with separate services on the same ATM interface. In each case, the port can therefore host more than one SAP as long as the configured encapsulation identifiers are unique and do not overlap. This is central to SR OS service architecture: a SAP binds a service to an access interface and uses the encapsulation value to classify ingress traffic into the proper service instance. Nokia's SR OS documentation describes SAPs and their encapsulation-based traffic demultiplexing in its service and interface configuration material. See the Nokia SR OS documentation portal and the Nokia 7750 SR OS Information Center.

QUESTION NO: 35

How many service labels must be signaled to bring up a fully-meshed VPLS among five PE routers?

A. 1

B. 5

C. 10

D. 20

ANSWER: D

Explanation:

20 is correct because a fully meshed VPLS with five PE routers requires a logical, directed pseudowire relationship from each PE to every other PE. Each PE therefore needs to advertise one service (VC) label for each of its four remote PE peers. This produces 5 × 4 = 20 signaled service labels. Put another way, there are ten unique PE-to-PE pairings, but each pairing has two directions and each direction requires a label advertisement from the transmitting PE, resulting in 10 × 2 = 20 labels. In LDP-signaled VPLS, the Label Mapping message carries the VC label that the receiving PE uses to identify traffic for the appropriate VPLS service over the pseudowire. The full mesh ensures that any PE can forward customer frames directly to every other participating PE while maintaining the required service separation. RFC 4762 describes this VPLS model and the use of VC labels in LDP-based VPLS signaling: RFC 4762: Virtual Private LAN Service Using LDP Signaling. The general pseudowire label-signaling framework is also specified in RFC 4447: Pseudowire Setup and Maintenance Using LDP.

QUESTION NO: 36

Which of the following best describes a local Epipe Service?

A. An Epipe service that has 1 SAP and 1 SDP configured.

B. An Epipe service that has 3 or more SAPs configured.

C. An Epipe service that has 2 SAPs on the same router and has no SDPs configured.

D. An Epipe service that is configured between two routers, each having 1 SAP and 1 SDP configured

ANSWER: C

Explanation:

An Epipe service that has 2 SAPs on the same router and has no SDPs configured is a local Epipe service. Epipe is Nokia's point-to-point Ethernet service type, designed to associate exactly two service endpoints. A local Epipe is entirely terminated within one service router: each endpoint is represented by a Service Access Point (SAP), and both SAPs belong to the same Epipe service instance on that router. Ethernet frames received at either SAP are switched directly to the other local SAP according to the service configuration. Because no traffic must cross an MPLS or IP transport network to reach a remote

service router, the configuration does not require a Service Distribution Point (SDP) or an SDP binding. This is useful where two locally attached customer circuits, ports, or VLAN encapsulations need a private, point-to-point Layer 2 connection. Nokia's service documentation describes Epipe as a point-to-point service and distinguishes local service connectivity from connectivity that uses SDPs to extend a service to another router. See the Nokia SR OS Epipe service documentation and the Nokia SR OS services overview.

QUESTION NO: 37

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 operational troubleshooting tool for validating an SDP's data-plane reachability. It sends test traffic in the same SDP encapsulation context used by the service, rather than merely checking ordinary IP reachability outside that transport context. Consequently, "Tests ability of reaching the far-end IP address of an SDP ID within the SDP encapsulation." correctly describes its key purpose: proving that the configured far-end endpoint can be reached through the selected SDP tunnel or transport mechanism.

The tool supports both directions of service-path validation. "Provides in-band uni-directional connectivity tests." is correct because an operator can verify forwarding toward the remote SDP endpoint in-band. "Provides in-band round-trip connectivity tests." is also correct because the remote endpoint can return a response, allowing validation of bidirectional connectivity and measurement of the request/response result. These functions make sdp-ping particularly useful when isolating whether a service issue lies in SDP transport connectivity rather than in customer traffic or control-plane routing. Nokia's SR OS documentation portal contains the CLI and OAM command references for SDP operational tools: Nokia SR OS documentation. Additional protocol context for IP echo request/reply behavior is provided by RFC 792.

QUESTION NO: 38

Which of the following tasks must be carried out before deleting an E-pipe service? (Choose four.)

A. Shut down all SAPs associated with the service.

B. Shut down all SDPs associated with the service.

C. Delete all SAPs and SDPs associated with the service.

D. Shut down all physical ports that are bound to a SAP.

E. Delete the Customer ID.

F. Shut down the service.

ANSWER: A B C F

Explanation:

Before an E-pipe service can be deleted, the service must first be administratively shut down and cleared of the service objects that depend on it. "Shut down the service" places the service in the required administrative state for removal. Each attached SAP must also be shut down before it can be removed, so "Shut down all SAPs associated with the service" is

required. Likewise, the service's SDP associations must be shut down before their removal; this is represented by "Shut down all SDPs associated with the service." After the dependent attachment and transport objects are shut down, "Delete all SAPs and SDPs associated with the service" removes those bindings from the E-pipe. The service can then be deleted because no SAP or SDP resources remain configured under it. This ordered approach preserves configuration integrity and follows the SR OS service model, in which service access points and service-distribution-path bindings are subordinate service components. Nokia's SR OS documentation provides the service configuration model and operational procedures in its Service Router documentation portal; the associated service and E-pipe configuration material is available through the SR OS Services Overview.

QUESTION NO: 39

Which of the following best describes a MAC FDB used by a VPLS service?

A. A VPLS service does not use a MAC FDB.

B. A single MAC FDB is shared by all VPLS services on a PE.

C. A separate MAC FDB is used on the PE for each VPLS service.

D. Each VPLS service uses a separate MAC FDBs for unicast and multicast traffic.

ANSWER: C

Explanation:

A separate MAC FDB is used on the PE for each VPLS service. VPLS emulates an individual Ethernet LAN across a packet-switched provider network, so each VPLS instance constitutes a distinct bridging domain. The PE learns source MAC addresses from frames received on that service's customer-facing SAPs and pseudowires, then records the learned MAC-to-forwarding-location association in the forwarding database for that same service. When a destination MAC address is known, the PE uses this service-specific FDB to forward the frame only to the applicable local attachment circuit or remote pseudowire. When it is unknown, the frame is flooded within that VPLS instance according to its forwarding rules.

Keeping the FDB per service is fundamental to customer and service separation: identical customer MAC addresses may legitimately appear in different VPLS instances without creating an ambiguity, and MAC learning in one service must not influence forwarding in another. This behavior follows the VPLS model defined as an emulated LAN with its own forwarding and learning context. RFC 4762 describes VPLS forwarding based on MAC-address learning and forwarding within a VPLS instance, while RFC 4664 defines the broader framework for L2 VPN service instances. See RFC 4762 and RFC 4664.

QUESTION NO: 40

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

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 in-band service-distribution-path diagnostic feature used to validate forwarding over the configured SDP rather than merely checking ordinary IP reachability. It supports a unidirectional connectivity test, enabling validation that test traffic sent through the SDP reaches the remote service termination. This is useful where the operational objective is to establish successful forwarding in the direction being tested. It also supports round-trip connectivity testing: a probe is sent through the

SDP to the far-end system and a response is returned, allowing the originating router to verify end-to-end SDP connectivity from its own operational perspective.

The test is encapsulated using the SDP's transport context, so it specifically checks the ability to reach the far-end IP address associated with the selected SDP ID through that SDP encapsulation. Consequently, the result validates the service transport path and its far-end endpoint, not just the availability of a generic routed path. These functions align with the broader MPLS/OAM model of using in-band probe and reply mechanisms to verify a provisioned forwarding path. Nokia SR OS OAM documentation is available through the Nokia SR OS Documentation portal; the underlying MPLS echo-request/echo-reply diagnostic model is described in RFC 4379.

QUESTION NO: 41

An operator needs to provision 6VPE on a 7750 SR that is also used to provide existing IPv4 VPRN services. IPv6 VPRN customers use BGP as a PE-CE protocol. Which of the following statements is FALSE regarding the configuration steps required?

A. MP-BGP sessions need to be configured in the service provider network to support a VPN-IPv6 address family.

B. PE routers must be configured to run MP-BGP and IPv6

C. PE-to-CE BGP sessions need to be configured to support IPv6.

D. The provider core must be configured to support IPv6

E. None of the above statements are false.

F. All of the above statements are false

ANSWER: D

Explanation:

The correct false statement is **"The provider core must be configured to support IPv6"**. The defining purpose of 6VPE is to carry IPv6 VPN customer traffic across an existing IPv4/MPLS provider backbone without requiring IPv6 forwarding or IPv6 IGP deployment in the P-router core. IPv6 routes from customer sites are exchanged with the PE using an IPv6-capable PE-CE protocol, and the PE advertises those VPN routes through MP-BGP using the VPN-IPv6 address family. For transport across the provider network, the ingress PE imposes MPLS labels and the core forwards the labeled packets using its established IPv4/MPLS transport infrastructure. The egress PE removes the transport labels and forwards the native IPv6 packet toward the destination customer site. Consequently, IPv6 capability is required where IPv6 VPN routes and customer-facing IPv6 connectivity are handled—particularly at the PEs—but it is not a prerequisite for transit P routers in the MPLS core. This model permits an operator to introduce IPv6 VPRN services while preserving existing IPv4 VPRN services and avoiding a core-wide IPv6 migration. RFC 4798 defines this IPv6 Provider Edge approach: RFC 4798. Nokia's SR documentation is available through the Nokia Service Router documentation portal.

QUESTION NO: 42

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 MPLS OAM diagnostic tool used to trace the forwarding path of a label-switched path (LSP). It operates by sending MPLS echo requests with incrementing time-to-live values so that successive LSRs along the LSP return diagnostic replies. The resulting output identifies the transit and egress nodes encountered by the LSP and can include MPLS forwarding information such as labels, interfaces, or next-hop details, depending on the router's configuration and software release. This makes lsp-trace appropriate when the exhibit shows a hop-by-hop MPLS LSP path rather than ordinary IP routing reachability or a service-specific probe. The tool implements MPLS LSP traceroute behavior standardized through MPLS echo request/reply processing. RFC 4379 defines the MPLS LSP ping and traceroute mechanisms, including TTL-expiry processing that enables path discovery: RFC 4379: Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures. Nokia's Service Router documentation also covers MPLS and service OAM commands, including LSP diagnostic operations: Nokia Service Router documentation.

QUESTION NO: 43

Click on the exhibit button below.

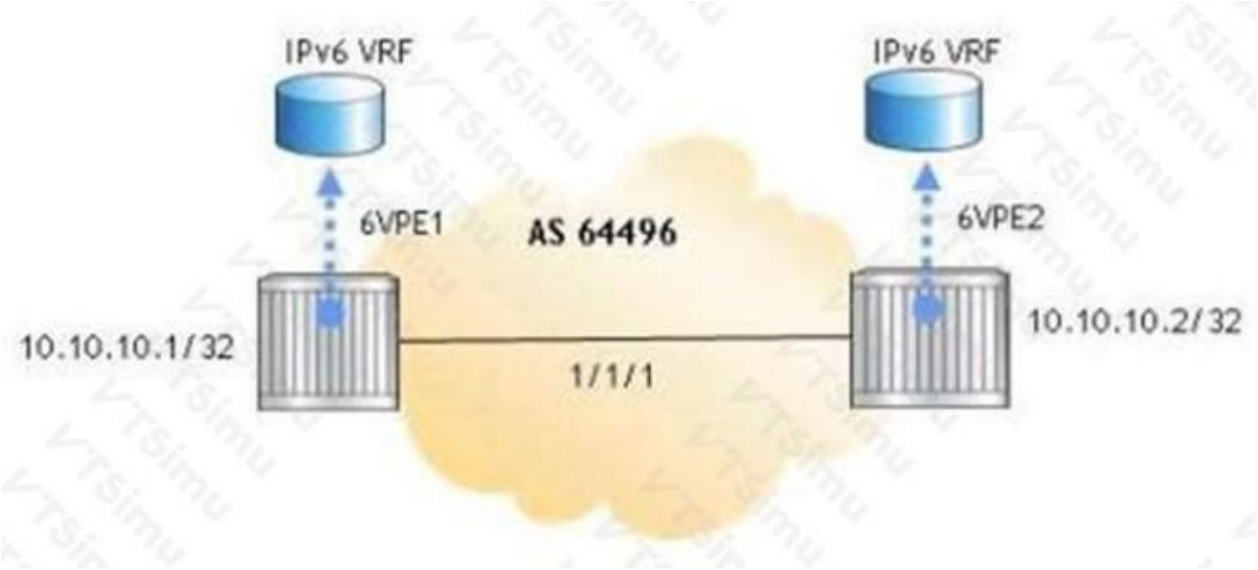

Which of the following is the correct configuration for an MP-BGP session on 6VPE1?

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 correct configuration is **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**. A 6VPE provider edge router uses Multiprotocol BGP to advertise and receive IPv6 VPN routes while the provider core can continue using IPv4 transport and IPv4 LDP/RSVP tunnels. Therefore, the BGP address family required for this PE-to-PE session is `vpn-ipv6`, which carries VPN-IPv6 Network Layer Reachability Information, including the route distinguisher and MPLS VPN label needed to identify and forward customer IPv6 prefixes across the MPLS backbone. The peer and local addresses shown are IPv4 addresses, which is appropriate for an IPv4-based provider transport/control-plane peer relationship. In SR OS, the BGP neighbor is configured with its IPv4 address in the BGP group, alongside the local IPv4 source address and the remote autonomous-system number. This combination enables the PE routers to establish the MP-BGP session and exchange labeled IPv6 VPN reachability for 6VPE service operation. Nokia's SR OS routing documentation describes BGP family configuration and VPN-IPv6 route exchange: Nokia SR OS BGP documentation. Additional 6VPE concepts are covered in Nokia SR OS IPv6 VPN documentation.

QUESTION NO: 44

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. By default, a 7750 SR 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:

On a 7750 SR VPRN service, route advertisement has distinct PE-to-CE and PE-to-PE behaviors. By default, the PE does not advertise every route present in the VRF to a directly attached CE. The CE-facing routing protocol configuration and any applicable policies determine the routes that the CE receives. This permits the service provider to control customer routing information and to avoid advertising routes that are not intended for that CE site.

For VPN reachability between PE routers, eligible VRF routes are exported from the local VPRN routing table into MP-BGP VPN routes. The PE advertises these routes to remote PEs, where route targets govern which VRFs can import them. This PE-to-PE exchange is the fundamental RFC 4364 BGP/MPLS IP VPN model: VPN-IPv4 routes carry a route distinguisher and route-target extended communities so remote PEs can preserve customer address separation and distribute routes to the appropriate VPN instances. Filters and explicit import/export policies can alter the default behavior when required. See RFC 4364, BGP/MPLS IP Virtual Private Networks and the Nokia SR OS documentation portal.