DUMPSARENA

Nokia Virtual Private LAN Services

Alcatel-Lucent 4A0-105

Version Demo

Total Demo Questions: 20

Total Premium Questions: 202

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Volume A	62
Topic 2, Volume B	69
Topic 3, Volume C	51
Total	202

QUESTION NO: 3

If the max remote-age timer of a B-VPLS is reached, what happens to the remotely learned MAC addresses in the I-VPLS service?

A. There is no effect. The MAC addresses for the I-VPLS age out independently of the aging of B-VPLS addresses.

B. The MAC addresses for the I-VPLS will age out immediately following the remote-age out of the backbone-macs of the B-VPLS.

C. By default, the remote-age timer of the B-VPLS is configured to 0 and addresses do not age out. The I-VPLS uses its own timer.

D. The same timer is used for B-VPLS and I-VPLS addresses on the node.

ANSWER: B

Explanation:

The MAC addresses for the I-VPLS will age out immediately following the remote-age out of the backbone-macs of the B-VPLS. In an IEEE 802.1ah Provider Backbone Bridging VPLS design, the B-VPLS learns and maintains backbone MAC (B-MAC) reachability, while the I-VPLS/I-VPLS learns customer MAC addresses reached through that backbone association. A remotely learned customer MAC is therefore dependent on the corresponding remote B-MAC forwarding state. The B-VPLS *max remote-age* timer controls how long remote B-MAC information may remain valid without refresh. Once that B-MAC entry reaches its remote-aging limit and is removed, the associated remote customer MAC entries in the I-VPLS are immediately aged out as well. This prevents the system from retaining customer forwarding entries whose transport identity is no longer valid and ensures that subsequent traffic is flooded or relearned as required to restore a current forwarding path. This dependency is a key part of the hierarchical MAC-learning and aging behavior used by PBB-VPLS services. Nokia's service documentation describes the B-VPLS and I-VPLS relationship in its VPLS documentation; the underlying PBB architecture is defined by IEEE 802.1ah.

QUESTION NO: 4

What scaling factors should a network engineer be aware of when designing VPLS networks? (Choose 3)

A. Control plane scaling.

B. Data plane scaling.

C. Network touches.

D. Route table reduction.

E. OAM tools.

ANSWER: A B C

Explanation:

Control plane scaling, data plane scaling, and network touches are the principal VPLS design-scaling considerations. Control plane scaling concerns the growth of service-discovery and MAC-reachability signaling, together with the number of pseudowires and peer relationships that must be maintained as additional provider-edge routers and VPLS instances are deployed. In a full-mesh VPLS design, the pseudowire count grows rapidly with the number of participating nodes, making signaling and control-state capacity essential planning inputs.

Data plane scaling addresses the forwarding work performed by each participating router: learned customer MAC entries, unknown-unicast/broadcast/multicast replication, interface and pseudowire capacity, and the bandwidth consumed when frames are replicated toward remote sites. These factors determine practical service size and traffic-engineering requirements. Network touches represent the operational scaling cost: each new site or service change may require provisioning, validation, and ongoing operational handling across multiple network elements. Designs that reduce manual touch points are substantially easier to expand safely. The VPLS architecture and the effects of full-mesh pseudowire connectivity are described in RFC 4762; the related VPLS service requirements and operational model are covered by RFC 4665.

QUESTION NO: 5

Which failure scenario is NOT supported by IEEE 802.lag?

A. Loss of connectivity.

B. Unidirectional Loss.

C. Loop detection.

D. Merged Services.

E. End-to-end bit error rates.

ANSWER: E

Explanation:

End-to-end bit error rates. is correct because IEEE 802.1ag Connectivity Fault Management (CFM) is designed to establish and supervise Ethernet service connectivity, isolate connectivity faults, and support fault-notification and diagnostic functions within Maintenance Domains. Its principal protocol mechanisms are Continuity Check Messages, Loopback Messages, and Linktrace Messages. These mechanisms allow maintenance entities to monitor continuity and investigate the path of Ethernet frames, but they do not define a method for measuring the number of corrupted bits relative to all transmitted bits across an end-to-end Ethernet service. A bit-error-rate measurement requires performance-monitoring capabilities and an error-measurement methodology beyond the CFM fault-management functions standardized by 802.1ag. The IEEE 802.1 working group identifies Connectivity Fault Management as the subject of the 802.1ag standard; its functionality concerns detection and diagnosis of Ethernet connectivity problems rather than physical-layer or end-to-end bit-error performance measurement. Performance monitoring for Ethernet services is addressed by complementary OAM specifications, such as ITU-T Y.1731. See the IEEE 802.1ag standard page and ITU-T Recommendation Y.1731.

QUESTION NO: 6

Which of the following describes the scope of the IEEE 802.3ah EFM standard? (Choose 3)

A. Monitor and troubleshoot a point-to-point full-duplex link.

B. Monitor and troubleshoot end-to-end services.

C. Can be used at the core to trigger MPLS protection mechanisms when a physical failure detection is not possible.

D. Should not be used at the core to trigger MPLS protection. BFD is the only method that can be used when physical failure detection is not possible.

E. Valuable in the last-mile connection to the Service Provider demarcation device.

ANSWER: A C E

Explanation:

IEEE 802.3ah Ethernet in the First Mile defines Ethernet OAM for a single point-to-point, full-duplex Ethernet link. Its operational scope includes discovery of the directly connected peer, remote fault indication, link monitoring, and remote loopback. Therefore, “Monitor and troubleshoot a point-to-point full-duplex link.” accurately describes the fundamental EFM OAM domain. EFM was designed especially for access and aggregation environments, where an operator needs visibility of the Ethernet segment between customer equipment and the provider edge or demarcation device. This makes “Valuable in the last-mile connection to the Service Provider demarcation device.” an appropriate application of the standard. Link-level fault and event information learned through EFM can also be used by network equipment to initiate local protection actions. Consequently, “Can be used at the core to trigger MPLS protection mechanisms when a physical failure detection is not possible.” is valid when EFM runs over the relevant Ethernet adjacency and its defect indication is integrated with protection policy. The IEEE working-group material describes the First Mile focus and OAM objectives at IEEE 802.3ah Ethernet in the First Mile. The IEEE standards page also identifies 802.3ah as the Ethernet in the First Mile amendment: IEEE 802.3ah.

QUESTION NO: 8

Click on the exhibit below.

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

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

B. R4 and R3 will remain reachable.

C. No devices will be reachable.

ANSWER: B

Explanation:

R4 and R3 will remain reachable. In a VPLS using a full-mesh of service distribution paths, each participating provider-edge router has its own pseudowire or SDP-based forwarding path to the other applicable VPLS participants. The failed SDP removes the VPLS transport adjacency between R2 and R1, but it does not withdraw or disable R2's independent VPLS forwarding adjacencies to R3 and R4. Therefore, R2 can continue to send known unicast traffic, as well as required broadcast, unknown-unicast, and multicast traffic, toward those two still-operational VPLS peers. This behavior is an important resiliency characteristic of a full-mesh VPLS design: a single SDP failure affects the remote endpoint reached through that SDP, while unaffected SDP bindings remain available to the service. VPLS forwarding also applies split-horizon behavior between mesh pseudowires, so service traffic is not expected to use another mesh pseudowire as a transit path to restore the failed direct adjacency. Nokia's Layer 2 service documentation describes VPLS service operation and SDP bindings; the base VPLS architecture and full-mesh pseudowire model are also defined in RFC 4762. See Nokia's Service Router documentation portal for the applicable VPLS configuration and forwarding documentation.

QUESTION NO: 9

The FDB has learned 2 MAC addresses. The first is from a device accessible via a locality defined SAP, the second is from a device accessible via an SDP. Both MAC addresses were learned at the exact same time. If all traffic from these hosts stops at the same time, which address will age out first from the FDB based on default settings?

A. MAC learned via SAP

B. MAC learned via SDP

C. There is not enough information to determine

D. Both will timeout at the same time

ANSWER: A

Explanation:

MAC learned via SAP is correct because the SR OS VPLS forwarding database applies separate default aging behavior to locally and remotely learned MAC addresses. A MAC learned through a Service Access Point is a local FDB entry, whereas a MAC learned through a Service Distribution Point is a remote entry learned from another PE. Under the default VPLS settings, local MAC entries have a shorter aging timer than remote MAC entries. Therefore, when both hosts become silent simultaneously and neither MAC address is refreshed by received traffic, the locally learned SAP MAC is removed first. This distinction helps retain remotely learned reachability information for longer across the provider network while allowing directly attached, local endpoint information to be withdrawn more quickly when it is no longer active. MAC aging is driven by the absence of traffic from the learned source MAC; any qualifying frame received from that source before its timer expires refreshes the corresponding FDB entry. The SR OS Layer 2 Services documentation describes VPLS FDB learning and aging behavior; see the Nokia SR OS documentation portal. The underlying VPLS requirement for MAC address learning and aging is also specified in RFC 4762.

QUESTION NO: 10

What is the default remote-age timer of a B-VPLS?

A. The same timer is used in all VPLS types.

B. Higher than the default timer of an I-VPLS.

C. Lower than the default timer of an I-VPLS.

D. The same timer as an I-VPLS.

ANSWER: C

Explanation:

Lower than the default timer of an I-VPLS is correct. In an IEEE 802.1ah/PBB-based VPLS design, the B-VPLS represents the backbone service layer and maintains forwarding information for backbone-facing connectivity. Its remote MAC-aging behavior is intentionally configured with a shorter default than that used by the I-VPLS service layer. This helps the backbone forwarding database discard stale remotely learned information more quickly after topology, pseudowire, or reachability changes, limiting the duration for which traffic could be directed using outdated forwarding state. The I-VPLS, which provides the customer-instance service layer, uses a longer default remote-age value because its learned customer MAC state has different service-layer aging requirements. The distinction is therefore a deliberate part of the hierarchical PBB-VPLS model rather than a single common timer applied to every VPLS type. Nokia's SR OS documentation describes B-VPLS and I-VPLS as the backbone and customer-instance layers used for PBB services; consult the Nokia SR OS documentation portal for the applicable release's VPLS command defaults and service guide. Background on the hierarchical VPLS architecture is also available in RFC 4762.

QUESTION NO: 12

What is the default service-mtu for a VPLS service?

A. 1500

B. 1514

C. 1518

D. It depends on the access port MTU.

E. It depends on the network port MTU.

ANSWER: B

Explanation:

The default service MTU for a VPLS service on Alcatel-Lucent/Nokia SR OS is **1514** bytes. This value represents the conventional maximum untagged Ethernet frame size as handled at the service layer: a 1500-byte Layer 3 payload together with the 14-byte Ethernet destination MAC address, source MAC address, and EtherType/length fields. The Ethernet FCS is not included in this service-MTU value. VPLS emulates a LAN across the provider network, so SR OS uses 1514 bytes as the default baseline to accommodate standard Ethernet frames consistently between service access points and pseudowire endpoints. The configured service MTU is an important service-level limit: the relevant SAPs and network transport must be able to support frames of at least this size, plus any required encapsulation overhead. Where jumbo-frame support is required, the VPLS service MTU can be explicitly increased, with corresponding MTU planning across the service and transport path. Nokia's SR OS service documentation describes service MTU behavior and its role in determining supported customer-frame sizes; the VPLS configuration documentation also identifies the default service-MTU setting. See the Nokia SR OS Service MTU documentation and the Nokia SR OS service command reference.

QUESTION NO: 13

Why is MAC explosion a factor to be considered when designing a large VPLS network? (Choose 2)

A. The time to relearn thousands of MAC addresses may result in excessive flooding or loss of packets.

B. If the FDB reaches the limit, the router will discard frames.

C. Discarding frames when the FDB limit is reached could deny SLA guarantees.

D. It could take several seconds to relearn thousands of MAC addresses denying QoS/SLA guarantees.

ANSWER: A D

Explanation:

MAC explosion is significant in a large VPLS because every participating provider-edge device may need to learn and maintain a very large forwarding database for customer MAC addresses. Following a topology event, service restart, or aging of learned entries, those addresses must be learned again from customer traffic. During this relearning interval, destinations that are not yet present in the forwarding database are treated as unknown unicast traffic and are flooded across the VPLS service. With thousands of MAC addresses, this process can create a substantial temporary flood load and can delay normal unicast forwarding, resulting in excess traffic or packet loss.

The scale also affects service quality directly: rebuilding thousands of entries can take several seconds, depending on traffic patterns and the rate at which source MAC addresses are observed. Applications carried over the VPLS may therefore experience congestion, loss, or disruption while forwarding converges, which can compromise QoS objectives and SLA commitments. VPLS designs should consequently consider forwarding-database scale, MAC aging, topology stability, and mechanisms that limit unnecessary MAC learning. Nokia's SR OS Layer 2 service documentation describes VPLS forwarding and MAC-learning behavior: Nokia SR OS VPLS documentation. Additional SR OS documentation is available through the Nokia Service Router documentation portal.

QUESTION NO: 14

MAC explosion is best described as a large amount of MAC addresses being withdrawn and relearned?

A. TRUE

B. FALSE

ANSWER: A

Explanation:

TRUE is correct. In a VPLS service, the provider edge devices maintain forwarding databases using learned customer MAC addresses. A MAC explosion describes the operational effect of a major topology or service-path event that causes a very large number of learned MAC entries to be flushed or withdrawn and then learned again. For example, this can follow a failure and recovery event affecting a VPLS attachment circuit, pseudowire, or forwarding path. Once entries have been removed, frames for destinations no longer present in the forwarding database may be flooded within the VPLS until traffic is received from those source MAC addresses and the entries are relearned on their current forwarding locations. The resulting burst of flooding, control-plane learning activity, and forwarding-table churn is what makes the event an "explosion." Thus, the description correctly focuses on the withdrawal and subsequent relearning of a large volume of MAC addresses, rather than merely one ordinary MAC move. Nokia's Layer 2 service documentation discusses VPLS MAC learning and forwarding-database behavior, while its SR OS documentation covers the operational handling of MAC-table entries and service events. See the Nokia SR OS VPLS documentation and the Nokia SR OS MAC learning documentation.

QUESTION NO: 15

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

A. sap 1/1/1

B. sap 1/1/1:0*

C. sap 1/1/1:0

D. sap 1/1/1:*

ANSWER: A D

Explanation:

`sap 1/1/1` identifies a null-encapsulated SAP. A null SAP does not use an outer VLAN identifier to classify customer traffic, so Ethernet frames received on the SAP, including frames containing customer 802.1Q tags, can be carried by the VPLS without requiring a separate SAP for each customer VLAN. This is appropriate where the complete customer Ethernet frame, including its VLAN information, must remain transparent to the service.

`sap 1/1/1:*` is a wildcard SAP on a dot1q-encapsulated access port. The wildcard matches VLAN-tagged frames for all VLAN IDs rather than matching just one configured VLAN. It therefore provides one attachment point for traffic from the full customer VLAN range and supports a VLAN-transparent VPLS design. This avoids provisioning a separate service SAP for every individual customer tag while allowing the VPLS to forward traffic across the service.

Nokia SR OS documents SAP identification and Ethernet service encapsulation behavior in its Layer 2 Services guide. See the Nokia SR OS Layer 2 Services Guide and the Nokia SR OS Services Overview.

QUESTION NO: 16

Which of the following is NOT a feature of the IEEE 802.3ah EFM standard?

A. Discovery and auto-negotiation of OAM capabilities.

B. Link monitoring via event notifications.

C. Remote failure indication (RFI).

D. Remote Loopback

E. Latency and jitter across a service.

ANSWER: E

Explanation:

Latency and jitter across a service is not a feature defined by IEEE 802.3ah Ethernet in the First Mile (EFM) OAM. The standard’s OAM function operates at the Ethernet link level and provides mechanisms that allow directly connected peers to discover OAM support, exchange capability information, monitor link conditions through OAM events, signal faults, and perform remote loopback for troubleshooting. These functions help an operator determine whether a physical Ethernet access link is operating correctly and isolate link-level problems. In contrast, latency is the time taken for traffic to traverse a path, while jitter is variation in that delay. Assessing those characteristics requires performance-measurement mechanisms, typically involving measurement packets, timestamps, and defined measurement intervals across a service or network path. Such service-performance functions are outside the scope of 802.3ah link OAM and are addressed by other Ethernet OAM and performance-monitoring technologies. IEEE identifies 802.3ah as the amendment that introduced Ethernet in the First Mile capabilities: IEEE 802.3ah standard information. Nokia’s documentation also describes 802.3ah OAM as link-level Ethernet OAM supporting discovery, link monitoring, fault notification, and remote loopback: Nokia Ethernet OAM documentation.

QUESTION NO: 17

Multiple VPLS services can be offered over the same set of LSP tunnels?

A. TRUE

B. FALSE

ANSWER: A

Explanation:

TRUE is correct. In an MPLS-based VPLS deployment, the transport LSP provides the shared packet-switched path across the provider network, while service-specific MPLS labels identify the individual VPLS instance carried over that transport. This label stacking allows a single set of LSP tunnels between provider edge routers to transport traffic for multiple independent VPLS services. Each service remains logically separate because the service label and associated service configuration identify the appropriate VPLS forwarding context at the far-end PE router.

On Nokia SR OS, this is commonly implemented through service distribution points (SDPs), including mesh SDPs for VPLS connectivity. An SDP can use MPLS transport tunnels and can support multiple services, avoiding the need to establish a separate end-to-end LSP solely for every VPLS service. This improves scalability and makes efficient use of the MPLS transport infrastructure while preserving the independent MAC learning, flooding, and forwarding behavior required by each VPLS instance. Nokia's Layer 2 Services documentation describes VPLS service transport over SDPs and MPLS tunnels: Nokia SR OS VPLS documentation. The underlying MPLS label-stack model is also defined in RFC 3031.

QUESTION NO: 18

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

A. TRUE

B. FALSE

ANSWER: B

Explanation:

FALSE is correct because the PBB-VPLS service model uses a defined association between an I-VPLS instance and a single B-VPLS instance. The I-VPLS is the customer-facing, backbone-edge service identified by an I-SID, while the B-VPLS is the backbone transport service that carries encapsulated customer traffic across the provider backbone. A B-VPLS may provide the shared backbone forwarding domain for multiple I-VPLS services, allowing traffic from multiple customer-facing services to use the same backbone service. The association in the reverse direction is not many-to-many: an individual I-VPLS selects one B-VPLS for its backbone connectivity. This design preserves the intended PBB separation between customer MAC addressing at the I-component and backbone MAC forwarding within the B-component, while enabling service multiplexing through I-SIDs. Nokia's SR OS documentation describes PBB-VPLS as mapping I-VPLS services into a B-VPLS transport context and documents the corresponding I-component and B-component roles. See the Nokia SR OS PBB-VPLS documentation and the Nokia Layer 2 Services overview.

QUESTION NO: 19

Which of the following statements about a management VPLS are true? (Choose 2)

A. A management VPLS carries customer traffic over the non-blocking spoke-SDP

B. A VPLS service must be defined as a management VPLS upon creation to take on the role of a management VPLS

C. The state of a spoke-SDP within a user VPLS will take on the state of the spoke-sdp within the management VPLS.

D. The management VPLS must be created before the user VPLS.

ANSWER: B C

Explanation:

A VPLS service must be defined as a management VPLS upon creation to take on the role of a management VPLS. In SR OS, the management-VPLS role is a service-creation attribute rather than a role that is applied later to an already-created user VPLS. This ensures that the service is established with the specific control relationship needed to manage the operational state of associated spoke SDPs.

The state of a spoke-SDP within a user VPLS will take on the state of the spoke-sdp within the management VPLS. A management VPLS provides a shared control mechanism for spoke-SDP forwarding state: a user VPLS associated with the management VPLS follows the relevant spoke-SDP state from that management service. This permits consistent activation or blocking of equivalent access paths across multiple user VPLSs, rather than requiring independent state handling in every service. The behavior is part of SR OS VPLS and spoke-SDP service concepts described in Nokia's VPLS documentation and the spoke-SDP documentation.

QUESTION NO: 20

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

A. The loopback function in 802.1ag is an intrusive test that will stop customer data.

B. The loopback function of 802.1ag is a non-intrusive "ping" sent from a MEP to a remote MEP or MIP.

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

D. The loopback function of 802.1ag is an intrusive test that will "loop" all customer data back to the original sender.

ANSWER: B C

Explanation:

The loopback function of 802.1ag is a non-intrusive "ping" sent from a MEP to a remote MEP or MIP. Ethernet Connectivity Fault Management uses Loopback Message and Loopback Reply frames to test continuity and validate reachability within a Maintenance Entity Group. These OAM frames are directed at the selected maintenance point and are returned as replies, allowing the originating MEP to verify the path without placing the customer service into a traffic-looping state. This makes the test analogous to a connectivity probe while normal customer forwarding continues.

The loopback function in 802.3ah is an intrusive test that will loop customer data. In Ethernet in the First Mile OAM, remote loopback is enabled through OAM control so that the far-end device returns received non-OAM traffic toward the sender. It is therefore intended for controlled diagnostic use, because traffic reaching the looped interface is reflected rather than forwarded normally. The distinction is fundamentally between 802.1ag's CFM message-and-reply diagnostic mechanism and 802.3ah's remote data-path loopback mechanism. IEEE identifies 802.1ag as the Ethernet CFM standard at IEEE 802.1ag and identifies 802.3ah as Ethernet in the First Mile OAM at IEEE 802.3ah.