Free Questions for HPE2-W09 by certsdeals

Shared by Briggs on 22-07-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Does this correctly describe how the Virtual Switching Extension (VSX) fabric reacts to various component failure scenarios?

Solution: The ISL and keepalive goes down, and after a few seconds, the keepalive link restores. Switch-l and Switch-2 remains up. The Split-recovery mode is enabled. In this case the secondary switch shutdowns SVls when keepalive is restored.

Options:

A- Yes

B- No

Answer:

A

Explanation:

The ISL and keepalive goes down, and after a few seconds, the keepalive link restores. Switch-1 and Switch-2 remains up. The Split-recovery mode is enabled. In this case the secondary switch shutdowns SVIs when keepalive is restored is a correct description of how the Virtual Switching Extension (VSX) fabric reacts to various component failure scenarios. VSX is a feature that provides active-active forwarding and redundancy for ArubaOS-CX switches. The ISL is the inter-switch link that connects two VSX nodes and carries data traffic. The keepalive link is a separate link that carries control traffic between two VSX nodes. The split-recovery mode is a feature that prevents split-brain scenarios when both VSX nodes lose connectivity with each other but remain up. When the ISL and keepalive goes down, both VSX nodes continue to forward traffic independently.When the keepalive link restores, the secondary switch detects that it has lost synchronization with the primary switch and shuts down its SVIs to prevent traffic loops1.

Question 2

Question Type: **MultipleChoice**

Is this part of a valid strategy for load sharing traffic across the links in an Ethernet Ring Protection Switching (ERPS) ring?

Solution: Implement Virtual Switching Extension (VSX) on pairs of ERPS switches at the same site. Then combine multiple links between two data centers into VSX LAGs (M-LAGs).

Options:

A- Yes

B- No

Answer:

B

Explanation:

Implement Virtual Switching Extension (VSX) on pairs of ERPS switches at the same site. Then combine multiple links between two data centers into VSX LAGs (MC-LAGs) is not part of a valid strategy for load sharing traffic across the links in an Ethernet Ring Protection Switching (ERPS) ring. ERPS is a feature that provides loop prevention and fast convergence for Layer 2 networks that use ring topologies. VSX is a feature that provides active-active forwarding and redundancy for ArubaOS-CX switches. VSX LAGs or MC-LAGs are LAGs that span across two VSX nodes and provide load balancing and resiliency. However, VSX LAGs or MC-LAGs are not supported by ERPS because they can create loops in the ring topology.A better way to load share traffic across the links in an ERPS ring would be to use link aggregation groups (LAGs) between two nodes in a ring as long as they are not multi-chassis LAGs (MC-LAGs)1.

Question 3

Question Type: **MultipleChoice**

Is this part of a valid strategy for load sharing traffic across the links in an Ethernet Ring Protection Switching (ERPS) ring?

Solution: Combine multiple links between two data centers into link aggregations (but not multi-chassis ones).

Options:

A- Yes

B- No

Answer:

A

Explanation:

Combine multiple links between two data centers into link aggregations (but not multi-chassis ones) is part of a valid strategy for load sharing traffic across the links in an Ethernet Ring Protection Switching (ERPS) ring. ERPS is a feature that provides loop prevention and fast convergence for Layer 2 networks that use ring topologies. ERPS can support link aggregation groups (LAGs) between two nodes in a ring as long as they are not multi-chassis LAGs (MC-LAGs). MC-LAGs are not supported by ERPS because they can create loops in the ring topology.

Question 4

Question Type: **MultipleChoice**

Refer to the exhibits.

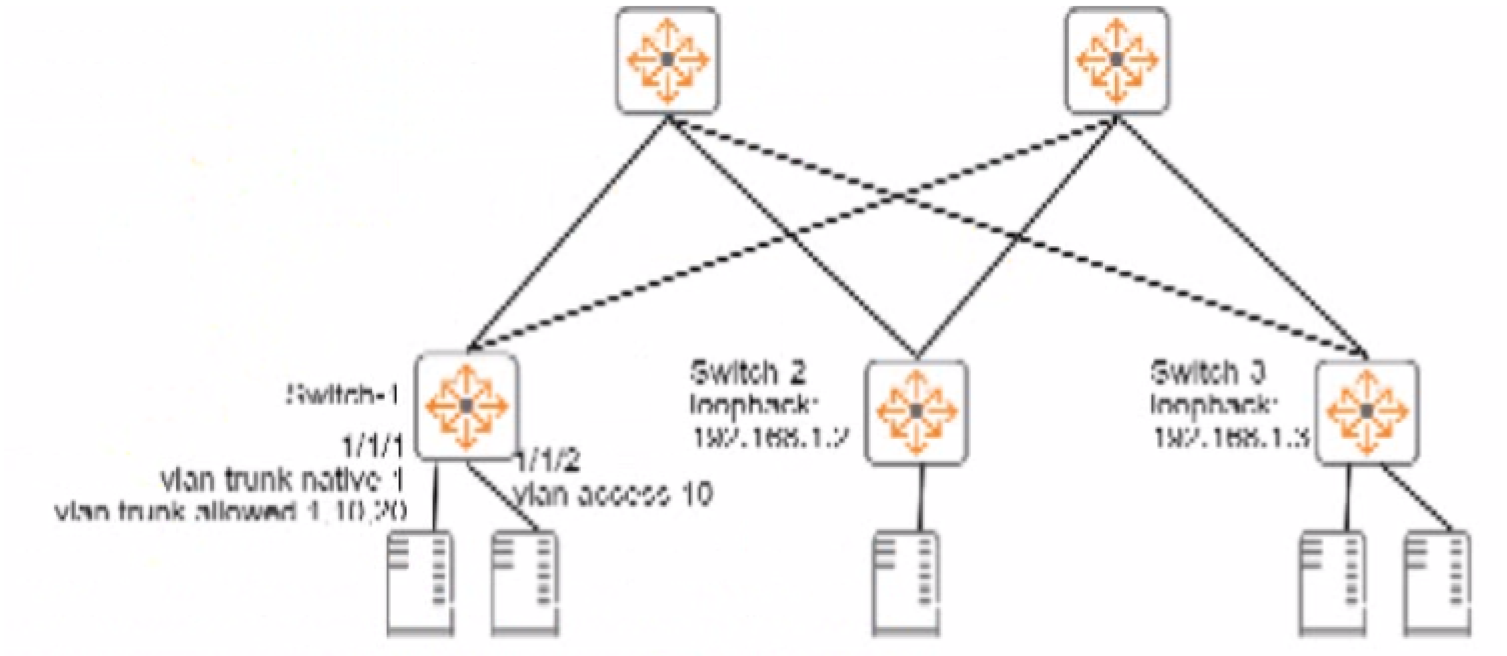

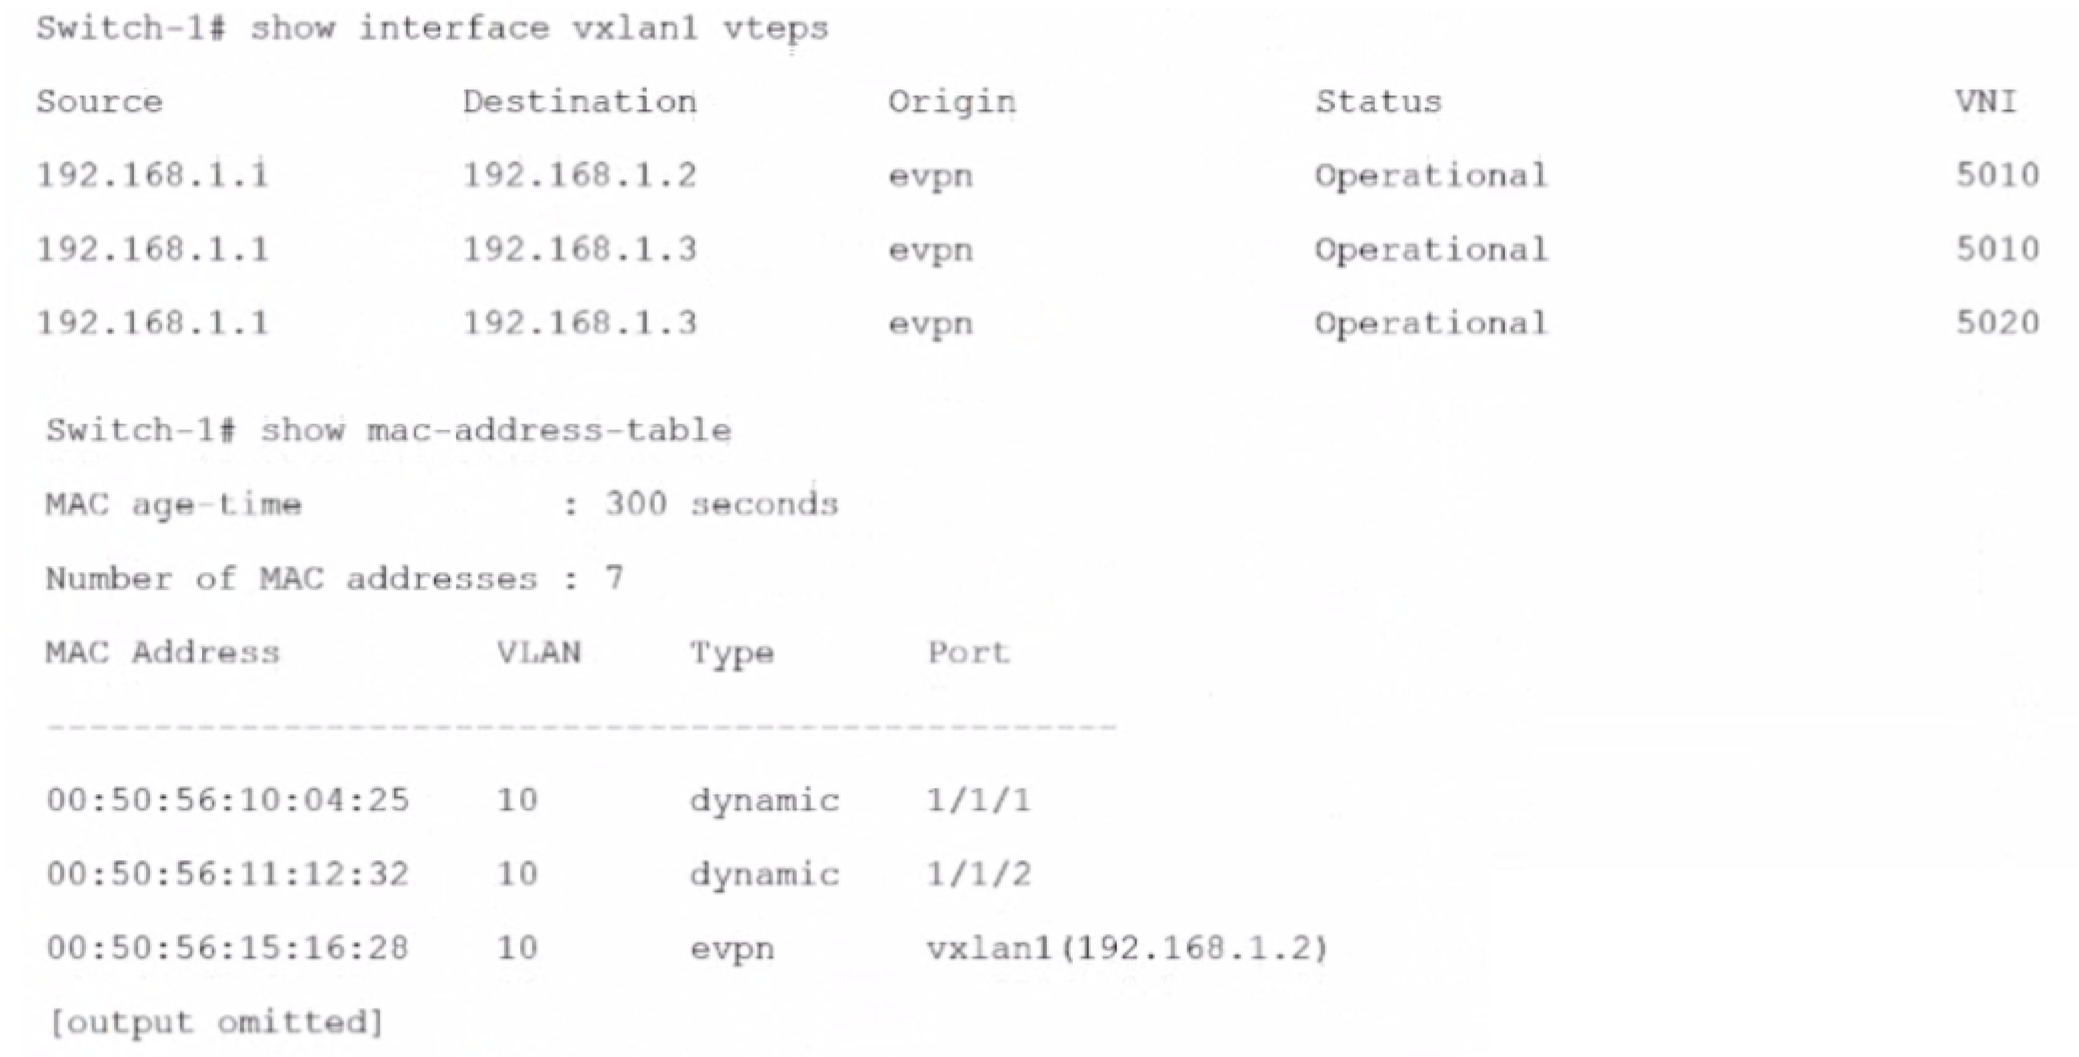

```
Switch-1# show interface vxlan1 vteps

Source          Destination      Origin          Status              VNI

192.168.1.1     192.168.1.2      evpn            Operational         5010

192.168.1.1     192.168.1.3      evpn            Operational         5010

192.168.1.1     192.168.1.3      evpn            Operational         5020

Switch-1# show mac-address-table

MAC age-time            : 300 seconds

Number of MAC addresses : 7

MAC Address          VLAN     Type       Port

--------------------------------------------------

00:50:56:10:04:25    10       dynamic    1/1/1

00:50:56:11:12:32    10       dynamic    1/1/2

00:50:56:15:16:28    10       evpn       vxlan1(192.168.1.2)

[output omitted]
```

Is this how the switch-1 handles the traffic?

Solution: A broadcast arrives in VLAN 10 on Switch-1. Switch 1 forwards the frame on all interfaces assigned to VLAN 10, except the incoming interface. It encapsulates the broadcast with VXIAN and sends it to 192.168.1.3, out not 192.168.1.2.

Options:

A- Yes

B- No

Answer:

B

Explanation:

A broadcast arrives in VLAN 10 on Switch-1. Switch 1 forwards the frame on all interfaces assigned to VLAN 10, except the incoming interface. It encapsulates the broadcast with VXLAN and sends it to 192.168.1.3, but not 192.168.1.2 is not a correct explanation of how the switch handles the traffic. Switch-1, Switch-2, and Switch-3 are ArubaOS-CX switches that use VXLAN and EVPN to provide Layer 2 extension over Layer 3 networks. VXLAN is a feature that uses UDP encapsulation to tunnel Layer 2 frames over Layer 3 networks using VNIs. EVPN is a feature that uses BGP to advertise multicast information for VXLAN networks using IMET routes. Switch-1 receives a broadcast in VLAN 10, which belongs to VNI 5010. Switch-1 forwards the frame on all interfaces assigned to VLAN 10, except the incoming interface, as per normal Layer 2 switching behavior. However, Switch-1 does not encapsulate the broadcast with VXLAN and send it only to 192.168.1.3, which is Switch-2's loopback interface, but rather replicates the broadcast, encapsulates each broadcast with VXLAN, and sends the VXLAN traffic to both 192.168.1.2 and 192.168.1.3, which are Switch-3's and Switch-2's loopback interfaces respectively.

Question 5

Question Type: **MultipleChoice**

Refer to the exhibits.

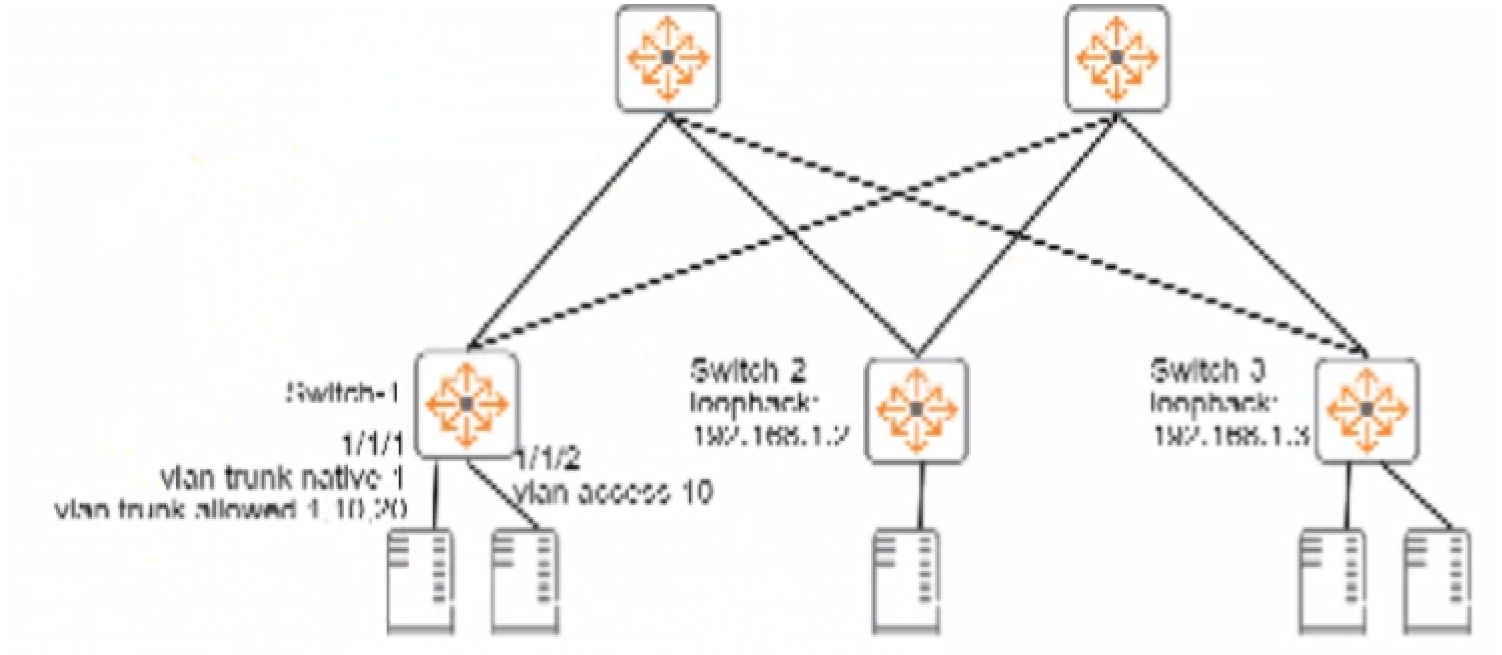

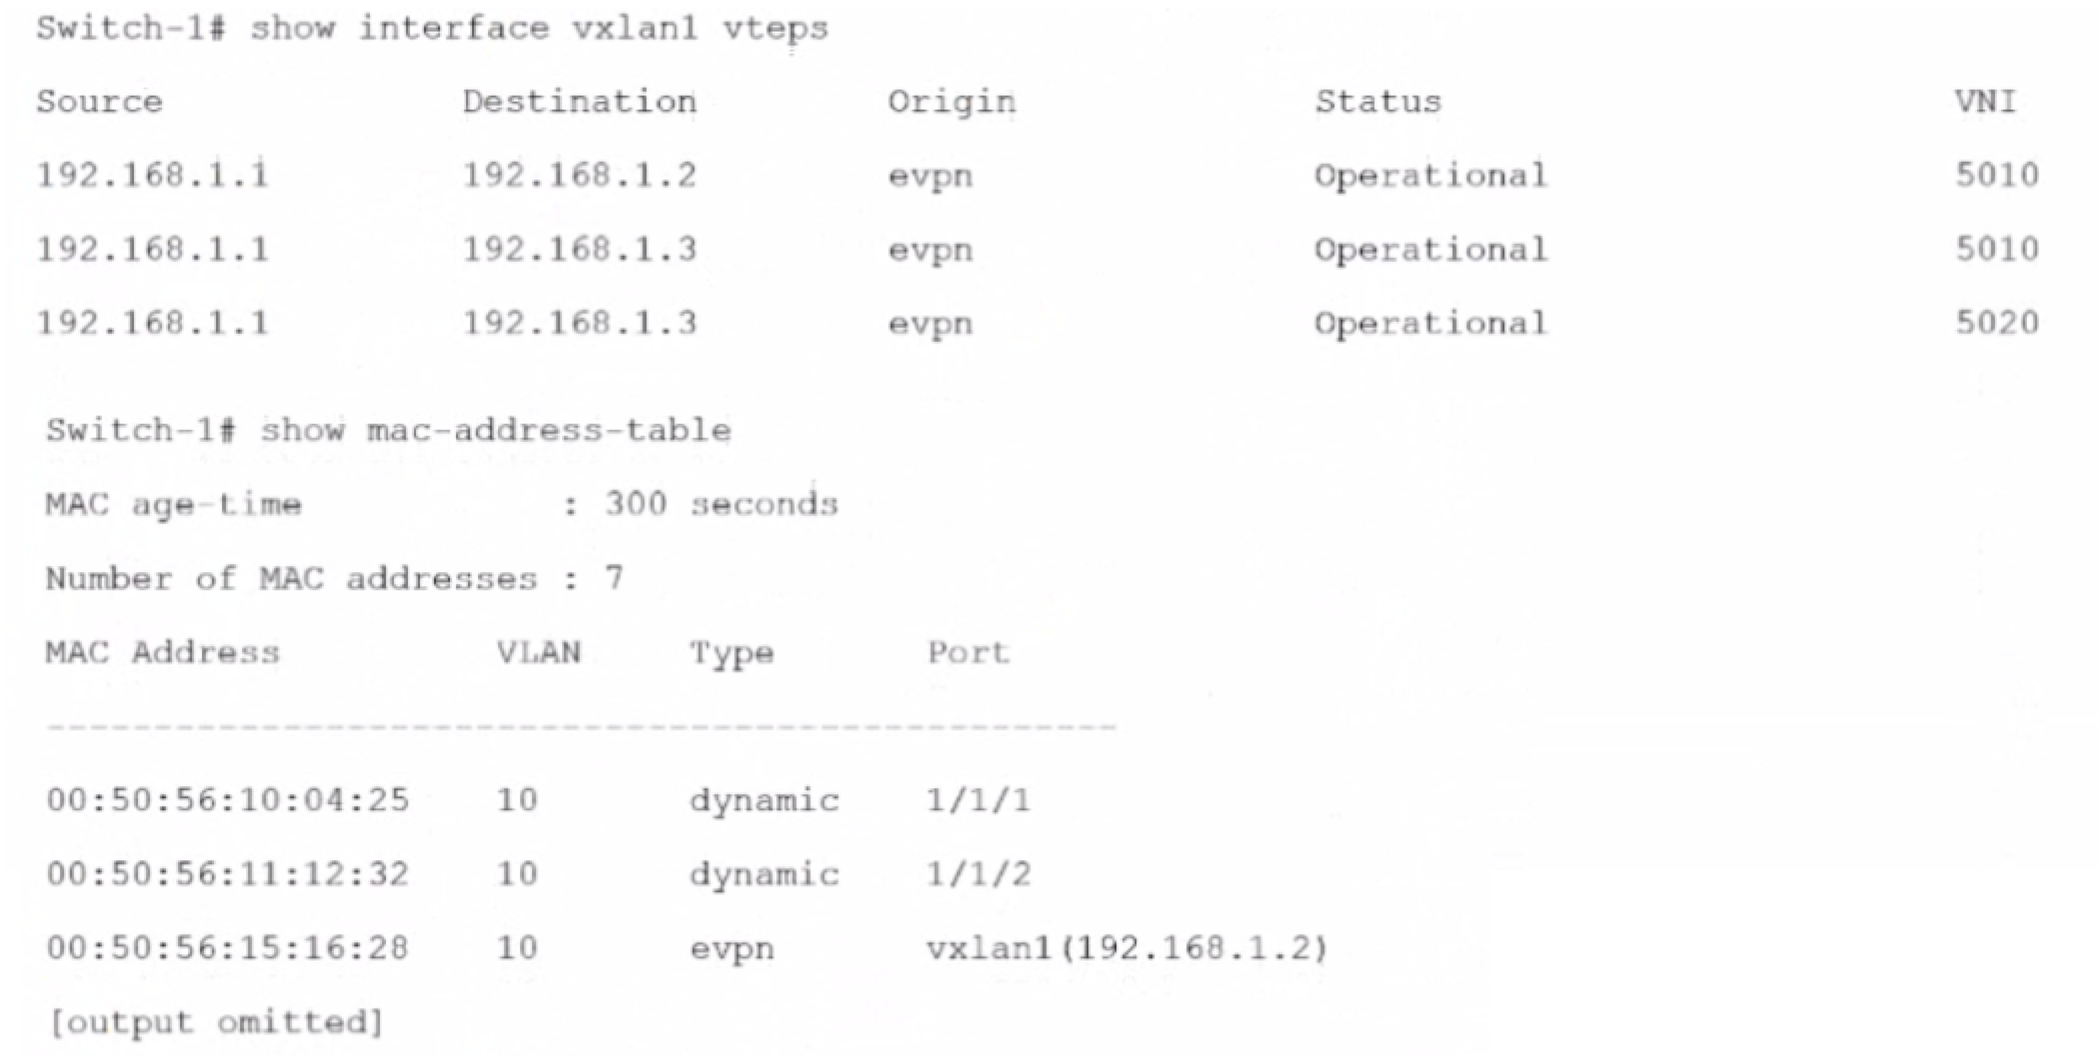

```
Switch-1# show interface vxlan1 vteps

Source          Destination     Origin          Status          VNI

192.168.1.1     192.168.1.2     evpn            Operational     5010

192.168.1.1     192.168.1.3     evpn            Operational     5010

192.168.1.1     192.168.1.3     evpn            Operational     5020
```

```
Switch-1# show mac-address-table

MAC age-time            : 300 seconds

Number of MAC addresses : 7

MAC Address          VLAN     Type       Port

--------------------------------------------------

00:50:56:10:04:25    10       dynamic    1/1/1

00:50:56:11:12:32    10       dynamic    1/1/2

00:50:56:15:16:28    10       evpn       vxlan1(192.168.1.2)

[output omitted]
```

Is this how the switch-1 handles the traffic?

Solution: A broadcast arrives in VLAN 10 on Switch-1. Switch 1 forwards the frame on all interfaces assigned to VLAN10. except the incoming interface. It replicates the broadcast, encapsulates each broadcast with VXLAN. and sends the VXLAN traffic to 192.168.1.2 and 192.168.1.3.

Options:

A- Yes

B- No

Answer:

A

Explanation:

A broadcast arrives in VLAN 10 on Switch-1. Switch 1 forwards the frame on all interfaces assigned to VLAN10, except the incoming interface. It replicates the broadcast, encapsulates each broadcast with VXLAN, and sends the VXLAN traffic to 192.168.1.2 and 192.168.1.3 is a correct explanation of how the switch handles the traffic. Switch-1, Switch-2, and Switch-3 are ArubaOS-CX switches that use VXLAN and EVPN to provide Layer 2 extension over Layer 3 networks. VXLAN is a feature that uses UDP encapsulation to tunnel Layer 2 frames over Layer 3 networks using VNIs. EVPN is a feature that uses BGP to advertise multicast information for VXLAN networks using IMET routes. Switch-1 receives a broadcast in VLAN 10, which belongs to VNI 5010. Switch-1 forwards the frame on all interfaces assigned to VLAN 10, except the incoming interface, as per normal Layer 2 switching behavior. Switch-1 replicates the broadcast, encapsulates each broadcast with VXLAN, and sends the VXLAN traffic to both 192.168.1.2 and 192.168.1.3, which are

Switch-3's and Switch-2's loopback interfaces respectively.

Question 6

Question Type: **MultipleChoice**

Refer to the exhibits.

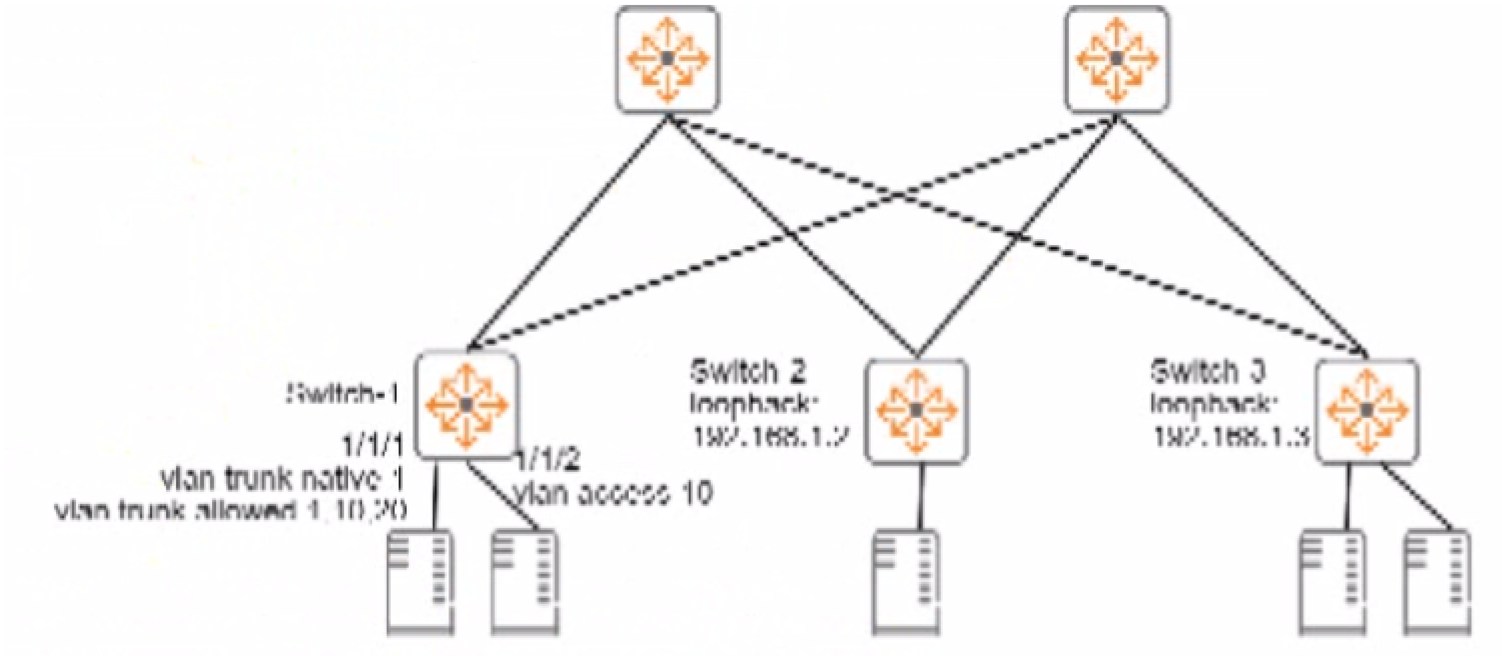

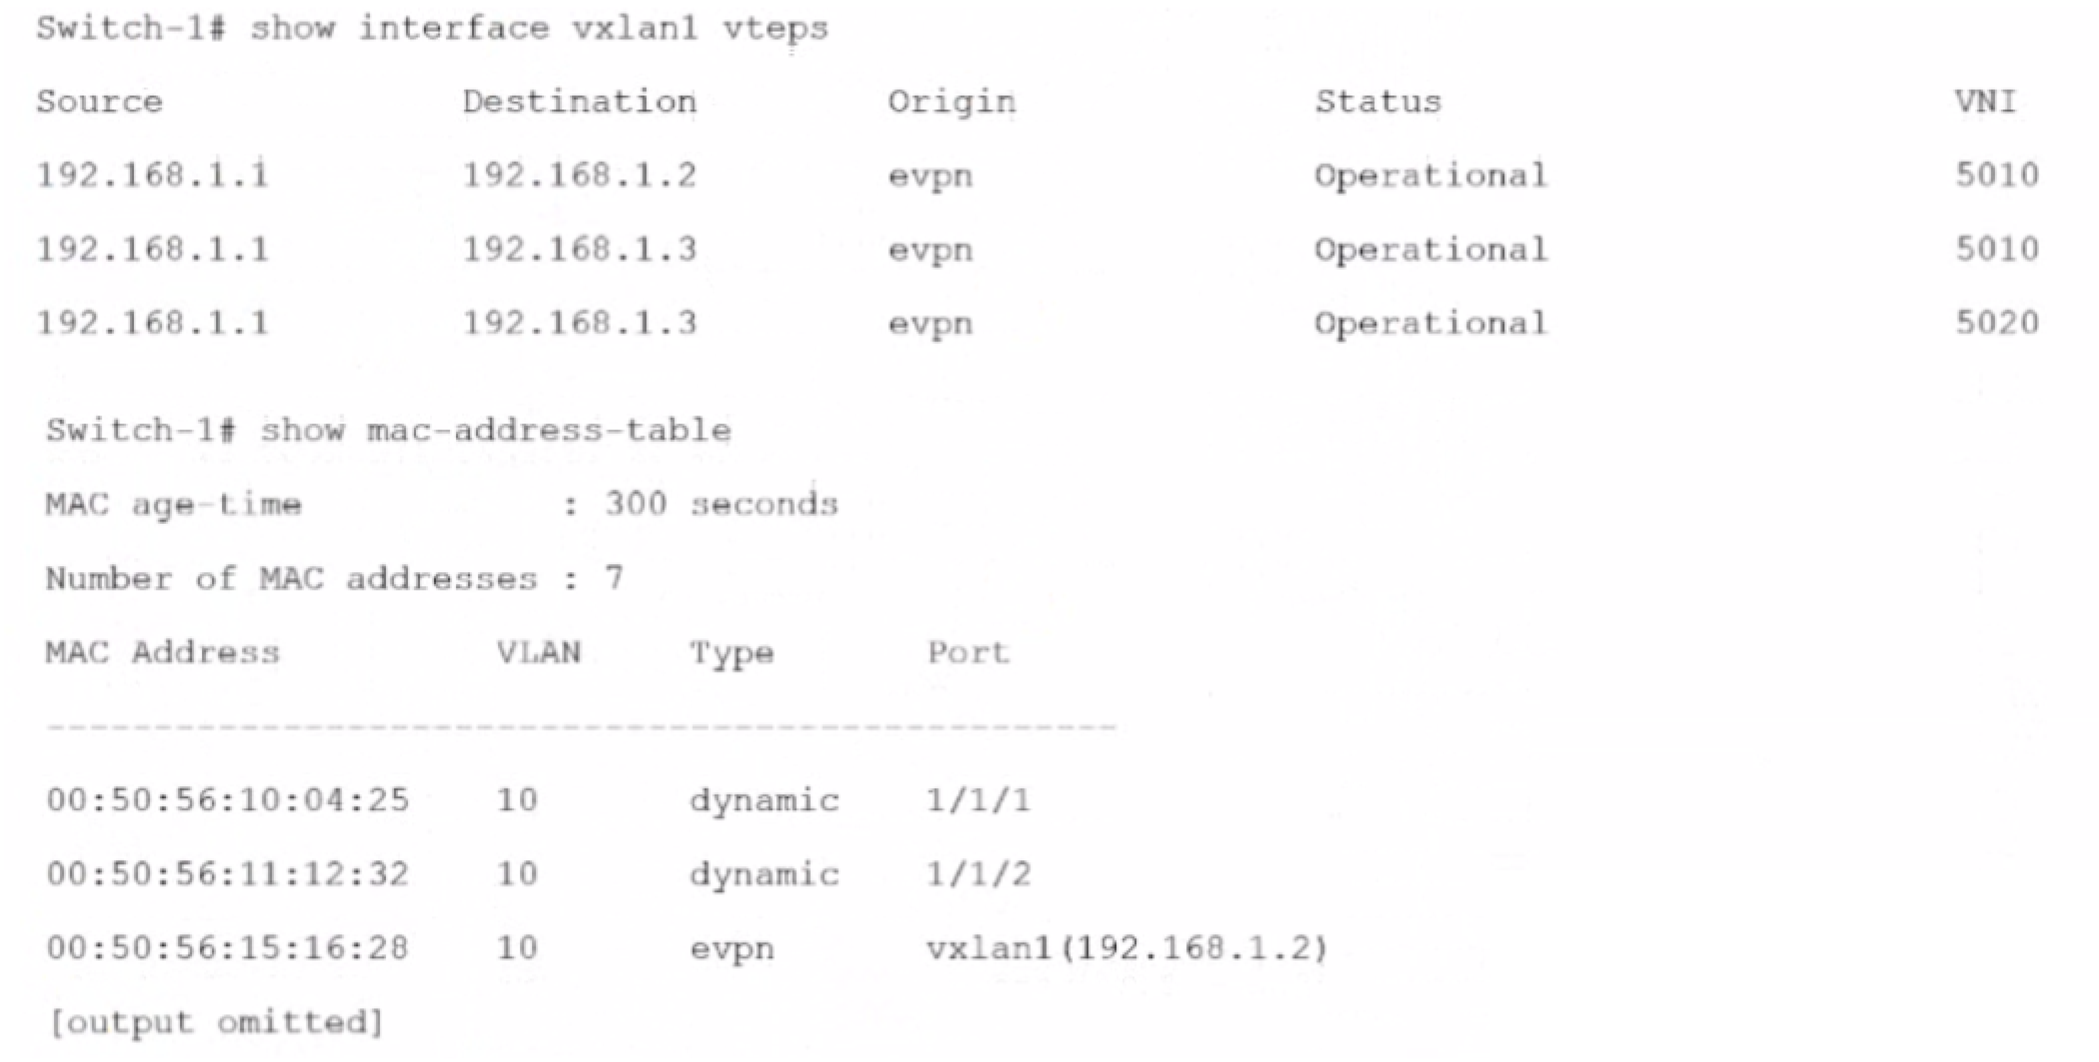

```
Switch-1# show interface vxlan1 vteps

Source          Destination     Origin          Status              VNI

192.168.1.1     192.168.1.2     evpn            Operational         5010

192.168.1.1     192.168.1.3     evpn            Operational         5010

192.168.1.1     192.168.1.3     evpn            Operational         5020

Switch-1# show mac-address-table

MAC age-time            : 300 seconds

Number of MAC addresses : 7

MAC Address          VLAN      Type        Port

--------------------------------------------------

00:50:56:10:04:25    10        dynamic     1/1/1

00:50:56:11:12:32    10        dynamic     1/1/2

00:50:56:15:16:28    10        evpn        vxlan1(192.168.1.2)

[output omitted]
```

Is this how the switch handles the traffic?

Solution: A broadcast arrives in VLAN 10 on Switch-1. Switch 1 forwards the frame on all interfaces assigned to VLAN10, except the incoming interface. It encapsulates the broadcast with VXIAN and sends it to 192.168.1.2. but not 192.168.1.3.

Options:

A- Yes

B- No

Answer:

B

Explanation:

A broadcast arrives in VLAN 10 on Switch-1. Switch 1 forwards the frame on all interfaces assigned to VLAN10, except the incoming interface. It encapsulates the broadcast with VXLAN and sends it to 192.168.1.2, but not 192.168.1.3 is not a correct explanation of how the switch handles the traffic. Switch-1, Switch-2, and Switch-3 are ArubaOS-CX switches that use VXLAN and EVPN to provide Layer 2 extension over Layer 3 networks. VXLAN is a feature that uses UDP encapsulation to tunnel Layer 2 frames over Layer 3 networks using VNIs. EVPN is a feature that uses BGP to advertise multicast information for VXLAN networks using IMET routes. Switch-1 receives a broadcast in VLAN 10, which belongs to VNI 5010. Switch-1 forwards the frame on all interfaces assigned to VLAN 10, except the incoming interface, as per normal Layer 2 switching behavior.However, Switch-1 does not encapsulate the broadcast with VXLAN and send it only to 192.168.1.2, which is Switch-3's loopback interface, but rather replicates the broadcast, encapsulates each broadcast with VXLAN, and sends the VXLAN traffic to both 192.168.1.2 and 192.168.1.3, which are Switch-3's and Switch-2's loopback interfaces respectively1.

Question 7

Question Type: **MultipleChoice**

Refer to the exhibits.

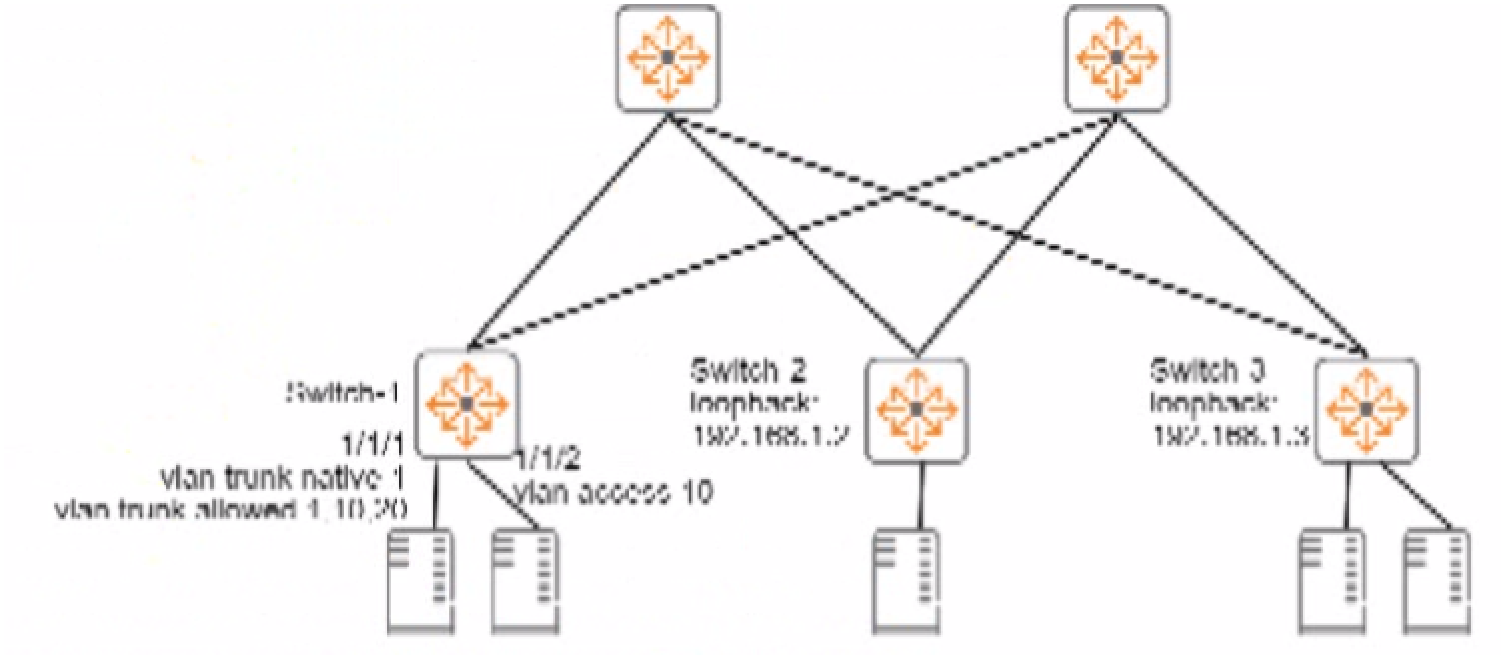

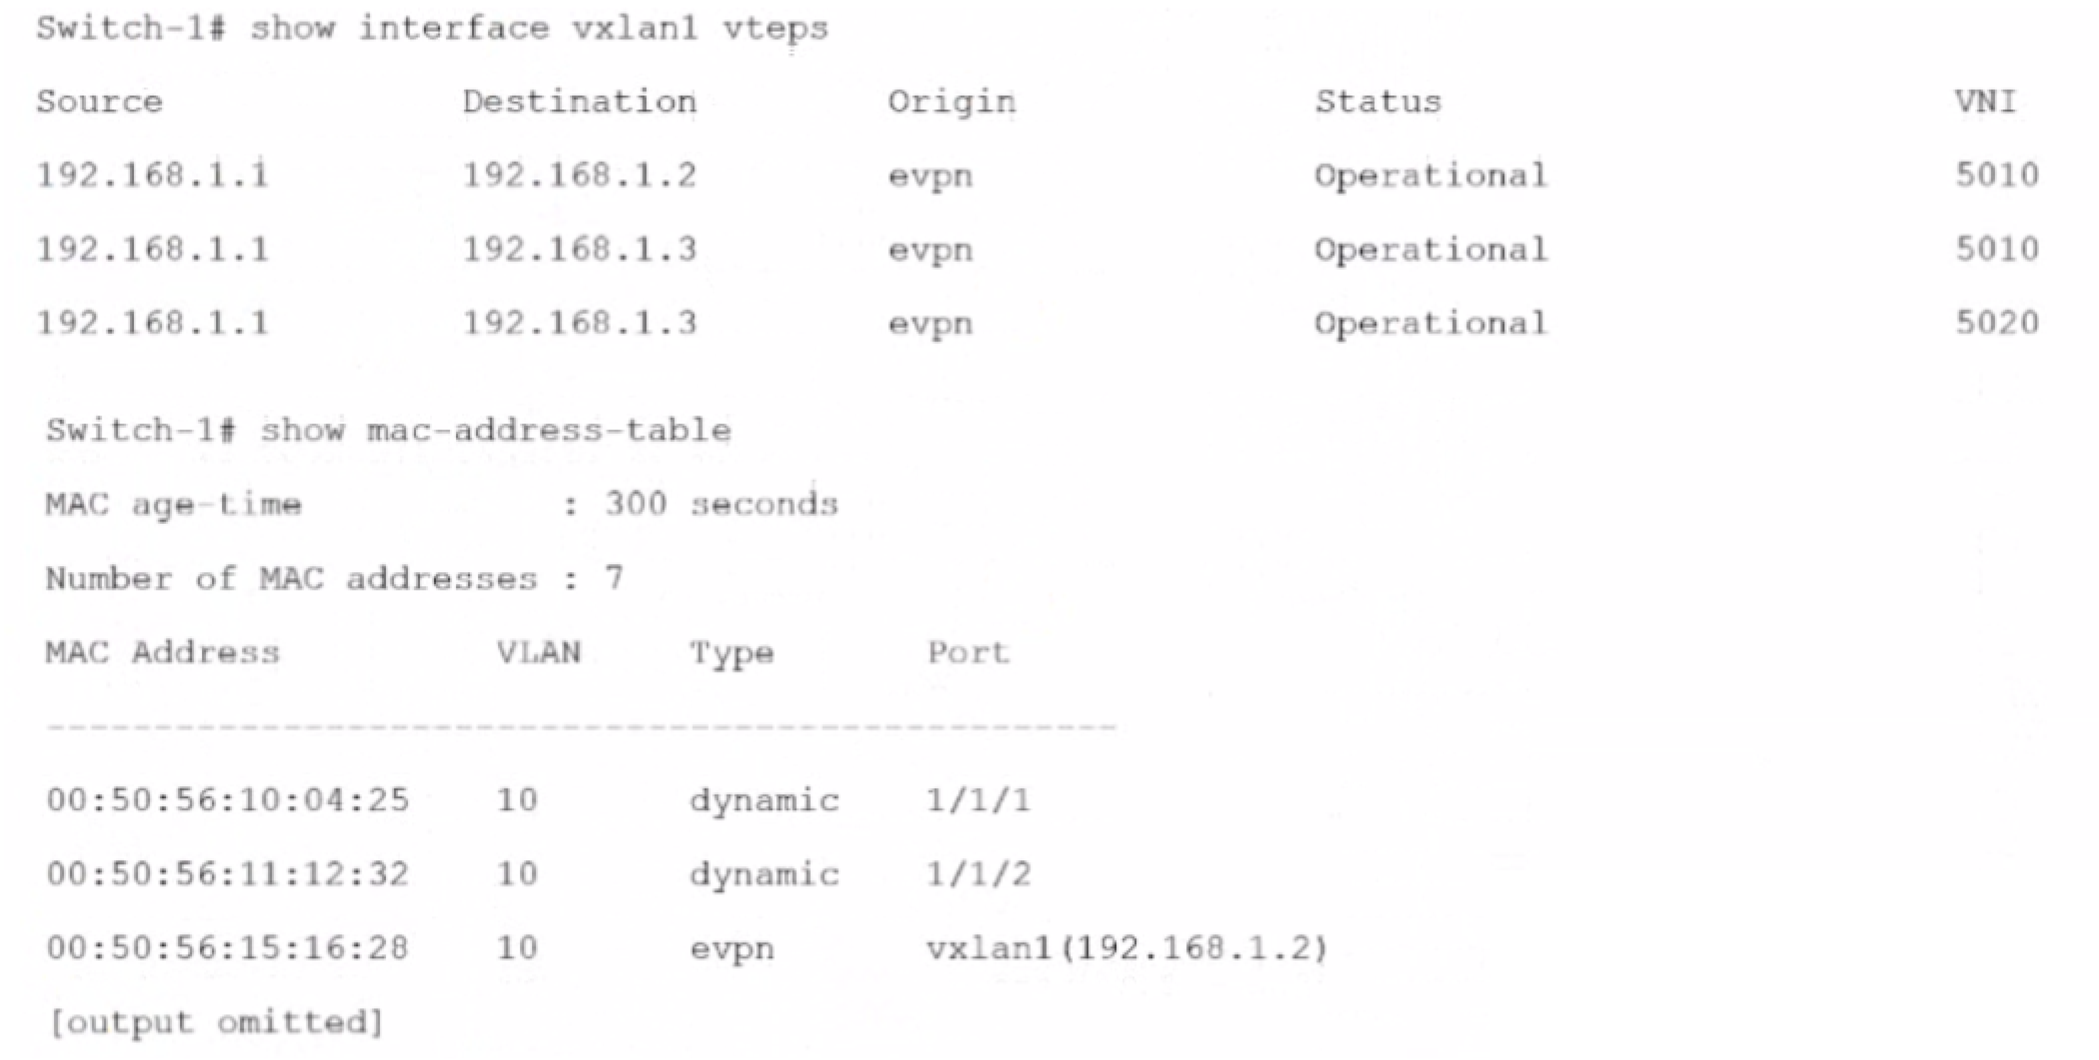

```
Switch-1# show interface vxlan1 vteps

Source          Destination      Origin          Status              VNI

192.168.1.1     192.168.1.2      evpn            Operational         5010

192.168.1.1     192.168.1.3      evpn            Operational         5010

192.168.1.1     192.168.1.3      evpn            Operational         5020
```

```
Switch-1# show mac-address-table

MAC age-time            : 300 seconds

Number of MAC addresses : 7

MAC Address          VLAN     Type       Port

--------------------------------------------------

00:50:56:10:04:25    10       dynamic    1/1/1

00:50:56:11:12:32    10       dynamic    1/1/2

00:50:56:15:16:28    10       evpn       vxlan1(192.168.1.2)

[output omitted]
```

Is this how the switch handles the traffic?

Solution: A frame with destination MAC address, 00:50:56:15:16:28, arrives with a VLAN 10 tag on 1/1/1 on Switch-1. Switch-1 encapsulates the frame with VXLAN and an IP header destined to 192.168.1.2.

Options:

A- Yes

B- No

Answer:

A

Explanation:

A frame with destination MAC address, 00:50:56:15:16:28, arrives with a VLAN 10 tag on 1/1/1 on Switch-1. Switch-1 encapsulates the frame with VXLAN and an IP header destined to 192.168.1.2 is a correct explanation of how the switch handles the traffic. Switch-1, Switch-2, and Switch-3 are ArubaOS-CX switches that use VXLAN and EVPN to provide Layer 2 extension over Layer 3 networks. VXLAN is a feature that uses UDP encapsulation to tunnel Layer 2 frames over Layer 3 networks using VNIs. EVPN is a feature that uses BGP to advertise MAC and IP addresses of hosts connected to VTEPs. Switch-1 receives a frame with destination MAC address, 00:50:56:15:16:28, which belongs to VM-2 on Switch-3. Switch-1 learns from EVPN that VM-2 is reachable through VTEP 192.168.1.2, which is Switch-3's loopback interface.Switch-1 encapsulates the frame with VXLAN and an IP header destined to 192.168.1.2 and sends it over the underlay network1.

Question 8

Question Type: **MultipleChoice**

You need to integrate Aruba Fabric Composer (AFC) with customer datacenter software. Is this integration possible?

Solution: Aruba Fabric Composer (AFC) with Nutanix Hypervisor (AHV)

Options:

A- Yes

B- No

Answer:

A

Explanation:

Aruba Fabric Composer (AFC) with Nutanix Hypervisor (AHV) integration is possible. AFC is a tool that provides automation and orchestration for managing data center networks composed of ArubaOS-CX switches. AFC can integrate with various data center software such as VMware vSphere, Nutanix AHV, Microsoft Hyper-V, etc.AFC can discover, monitor, and configure Nutanix AHV clusters and hosts using REST APIs1.

Question 9

Question Type: **MultipleChoice**

Does this correctly describe Network Analytics Engine (NAE) limitations on ArubaOS-CX switches?

Solution: Different switches have different limitations for the number of NAE scripts, monitors, and agents supported.

Options:

A- Yes

B- No

Answer:

A

Explanation:

Different switches have different limitations for the number of NAE scripts, monitors, and agents supported is a correct description of Network Analytics Engine (NAE) limitations on ArubaOS-CX switches. NAE is a feature that provides automation and analytics for managing ArubaOS-CX switches. NAE scripts are scripts that run on switches and collect data from various sources. NAE monitors are rules that define conditions and actions for NAE agents. NAE agents are instances of NAE scripts and monitors that run on switches.Different switches have different limitations for the number of NAE scripts, monitors, and agents supported depending on their hardware resources1.

Question 10

Question Type: **MultipleChoice**

Refer to the exhibit.

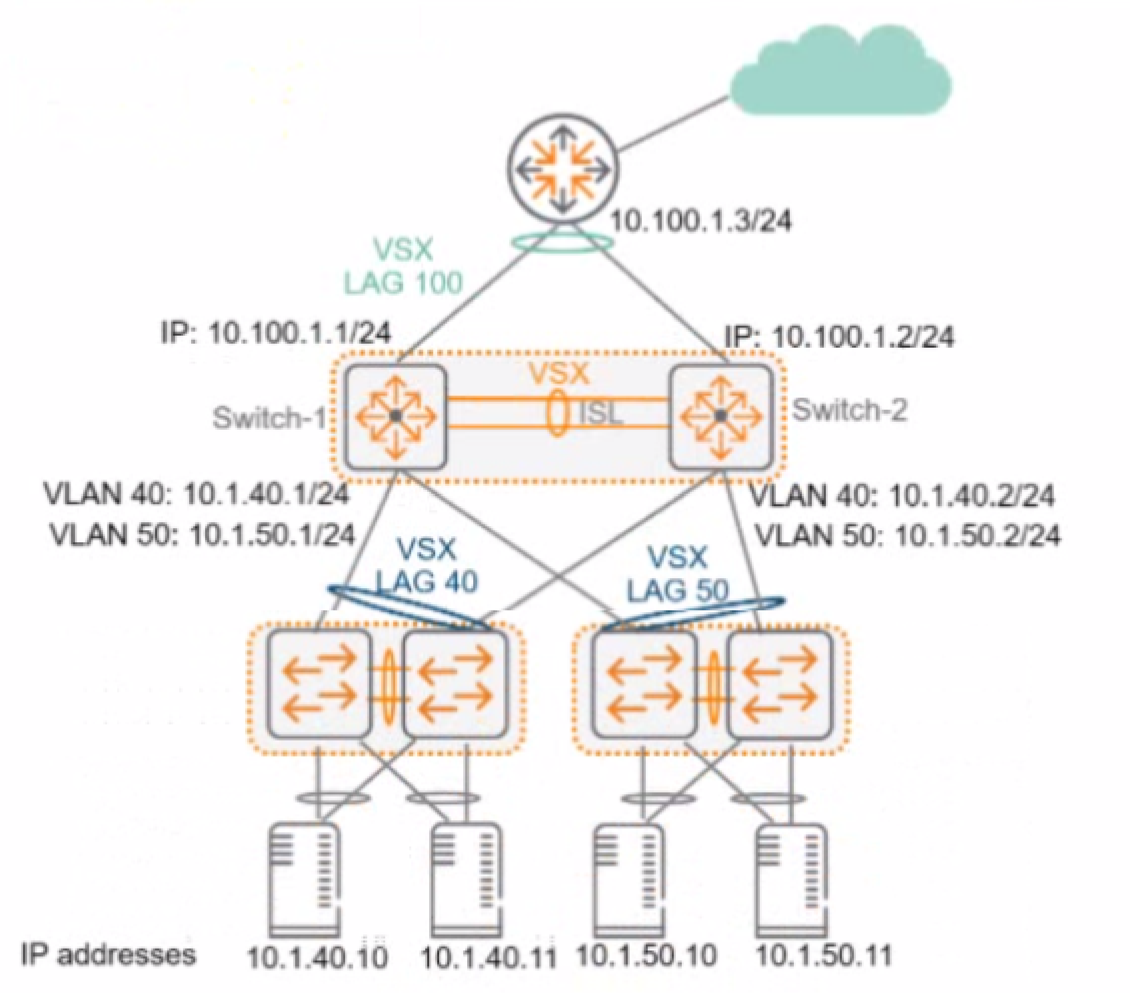

Switch-1, Switch-2, and the router run OSPF on LAG 100, which is a Layer 3 LAG. Does this correctly explain how to control how core-to-access traffic Is forwarded?

Solution: To reduce the amount of traffic sent over the ISI between Switch-1 and Swltch-2. enable active forwarding on LAG 100 on both Switch-1 and Switch-2.

Options:

A- Yes

B- No

Answer:

A

Explanation:

To reduce the amount of traffic sent over the ISL between Switch-1 and Switch-2, enable active forwarding on LAG 100 on both Switch-1 and Switch-2 is a correct explanation of how to control how core-to-access traffic is forwarded. Switch-1, Switch-2, and the router run OSPF on LAG 100, which is a Layer 3 LAG. Active forwarding is a feature that allows a switch to select one link as active and one link as standby for each direction of traffic in a LAG.Enabling active forwarding on LAG 100 on both Switch-1 and Switch-2 would reduce the amount of traffic sent over the ISL by sending traffic over only one link instead of both1.

Question 11

Question Type: MultipleChoice

Refer to the exhibit.

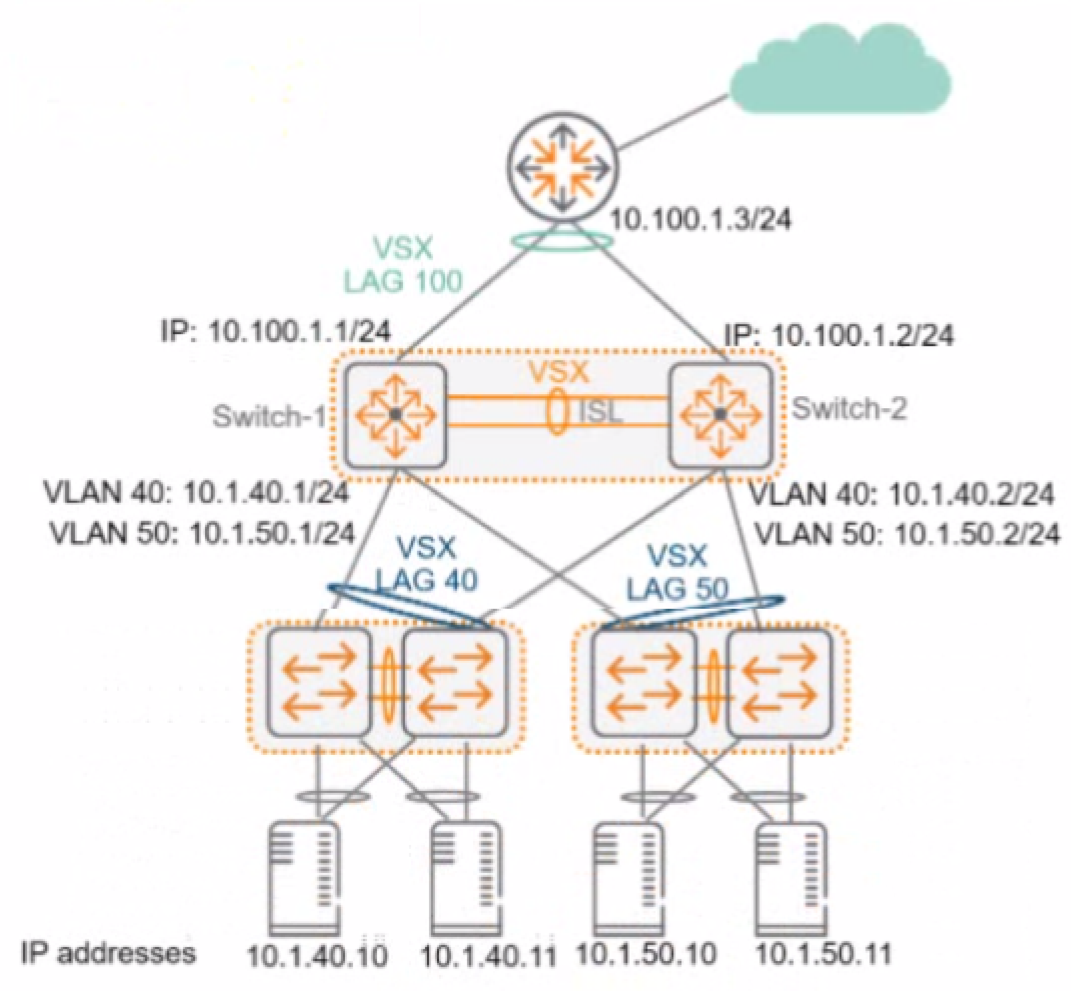

Switch-1, Switch-2, and the router run OSPF on LAG 100, which is a Layer 3 LAG. Does this correctly explain how to control how core-to-access traffic Is forwarded?

Solution: To reduce the amount of traffic sent over the ISL between Switch-1 and Switch-2. enable Equal Cost Multi Path (ECMP) on both Switch-1 and Switch-2.

Options:

A- Yes

B- No

Answer:

B

Explanation:

To reduce the amount of traffic sent over the ISL between Switch-1 and Switch-2, enable Equal Cost Multi Path (ECMP) on both Switch-1 and Switch-2 is not a correct explanation of how to control how core-to-access traffic is forwarded. Switch-1, Switch-2, and the router run OSPF on LAG 100, which is a Layer 3 LAG. ECMP is a feature that allows a router to load balance traffic destined to some network that is reachable through multiple equal cost route nexthops. Enabling ECMP on Switch-1 and Switch-2 would not reduce the amount of traffic sent over the ISL, but rather increase it by sending traffic over both links instead of one.A better way to reduce the amount of traffic sent over the ISL would be to enable active forwarding on LAG 100 on both Switch-1 and Switch-2, which would make one link active and one link standby for each direction of traffic1.

Question 12

Question Type: MultipleChoice

Is this a way that a data center technology can help meet requirements for multi-tenancy?

Solution: Virtual Extensible LAN (VXLAN) provides millions of IDs to scale for the needs of a multi-tenant environment

Options:

A- Yes

B- No

Answer:

A

Explanation:

Virtual Extensible LAN (VXLAN) provides millions of IDs to scale for the needs of a multi-tenant environment is a way that a data center technology can help meet requirements for multi-tenancy. Multi-tenancy is the ability to provide logical separation and isolation of network resources for different tenants or customers on a shared physical infrastructure. VXLAN is a feature that provides Layer 2 extension over Layer