Vendor:HP

Exam Code:HPE2-W09

Exam Name: Aruba Data Center Network Specialist Exam

Version: Demo

QUESTION 1

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

Solution: The keepalive goes down, ISL link remains up. Switch-1 and Switch-2 remains up. The Split-recovery mode is disabled. In this case the secondary switch shutdowns Svls.

A. Yes

B. No

Correct Answer: B

The keepalive goes down, ISL link remains up. Switch-1 and Switch-2 remains up. The Split-recovery mode is disabled. In this case the secondary switch shutdowns SVIs is not 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 keepalive goes down, but the ISL link remains up, both VSX nodes continue to forward traffic normally and do not shut down their SVIs because they can still exchange synchronization messages over the ISL link1.

QUESTION 2

Does this correctly describe routing information advertised by a VXLAN Tunnel Endpoint (VTEP) that uses EVPN?

Solution: IMET routes advertise the MAC addresses that the VTEP has learned locally in a VXLAN.

A. Yes

B. No

Correct Answer: B

IMET routes advertise the MAC addresses that the VTEP has learned locally in a VXLAN is not a correct description of routing information advertised by a VXLAN Tunnel Endpoint (VTEP) that uses EVPN. IMET routes are one of the types of routes that EVPN uses to advertise multicast information for VXLAN networks. IMET routes advertise the IP addresses of VTEPs that can join multicast groups for VXLAN segments2.

QUESTION 3

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

Solution: You can run NAE with VSX, but only the primary VSX member will actually run agents during normal operation.

A. Yes

B. No

Correct Answer: A

Network Analytics Engine (NAE) is a built-in analytics framework for network assurance and remediation on ArubaOS-CX switches. NAE allows monitoring, troubleshooting, and proactive network management using scripts and agents. Virtual Switching Extension (VSX) is a high-availability technology that allows two ArubaOS-CX switches to operate as a single logical device. You can run NAE with VSX, but only the primary VSX member will actually run agents during normal operation. The secondary VSX member will only run agents if the primary member fails or is rebooted1. Therefore, this correctly describes NAE limitations on ArubaOS-CX switches.

QUESTION 4

AtubaOS-CX switches are acting as Virtual Extensible LAN (VXLAN) Tunnel Endpoints (VTEPs) WITHOUT Ethernet VPN (EVPN).

Does this correctly describe how the VTEPs handle VXLAN traffic forwarding? Solution: VTEPs that use headend replication forward unicasts with unknown destination MAC addresses as unicast packets to each VTEP in the same VNl.

A. Yes

B. No

Correct Answer: A

VTEPs that use headend replication forward unicasts with unknown destination MAC addresses as unicast packets to each VTEP in the same VNI is a correct description of how the VTEPs handle VXLAN traffic forwarding. Headend replication is a method of replicating VXLAN packets at the ingress VTEP instead of using multicast routing. The ingress VTEP sends a copy of the VXLAN packet to each egress VTEP that belongs to the same VNI using unicast tunnels1.

QUESTION 5

Is this a rule for configuring schedule profiles on an ArubaOS-CX switch? Solution: With the exception of a single strict priority queue, all queues must use the same scheduling algorithm.

A. Yes

B. No

Correct Answer: A

With the exception of a single strict priority queue, all queues must use the same scheduling algorithm is a rule for configuring schedule profiles on an ArubaOS-CX switch. A schedule profile defines how traffic is scheduled across different queues on an interface. Each queue can have one of three scheduling algorithms: strict priority (SP), weighted round robin (WRR), or weighted fair queuing (WFQ). However, only one queue can have SP, and all other queues must have either WRR or WFQ1.

QUESTION 6

A customer\\'s servers use ISCSI, and they send data and storage traffic on the same pair of I OGbE links. Is this a best practice for supporting the ISCSI requirements?

Solution: Use Virtual Routing and Forwarding (VRF) to tunnel iSCSI traffic through the network spine on the same links that data traffic uses.

A. Yes

B. No

Correct Answer: B

ISCSI is a protocol that allows storage devices to communicate over IP networks. ISCSI traffic has different requirements than data traffic, such as low latency, high throughput, and reliability. Therefore, it is not a best practice to send data and storage traffic on the same pair of 10GbE links, as this can cause congestion and performance degradation. It is also not a best practice to use Virtual Routing and Forwarding (VRF) to tunnel ISCSI traffic through the network spine on the same links that data traffic uses. VRF is a technology that creates multiple isolated Layer 3 domains on a physical network, each with its own routing table. VRF does not provide any benefits for ISCSI traffic, as it does not guarantee bandwidth, priority, or quality of service. VRF also adds overhead and complexity to the network configuration1. Therefore, this is not a valid way to support the ISCSI requirements.

QUESTION 7

Refer to the exhibit.

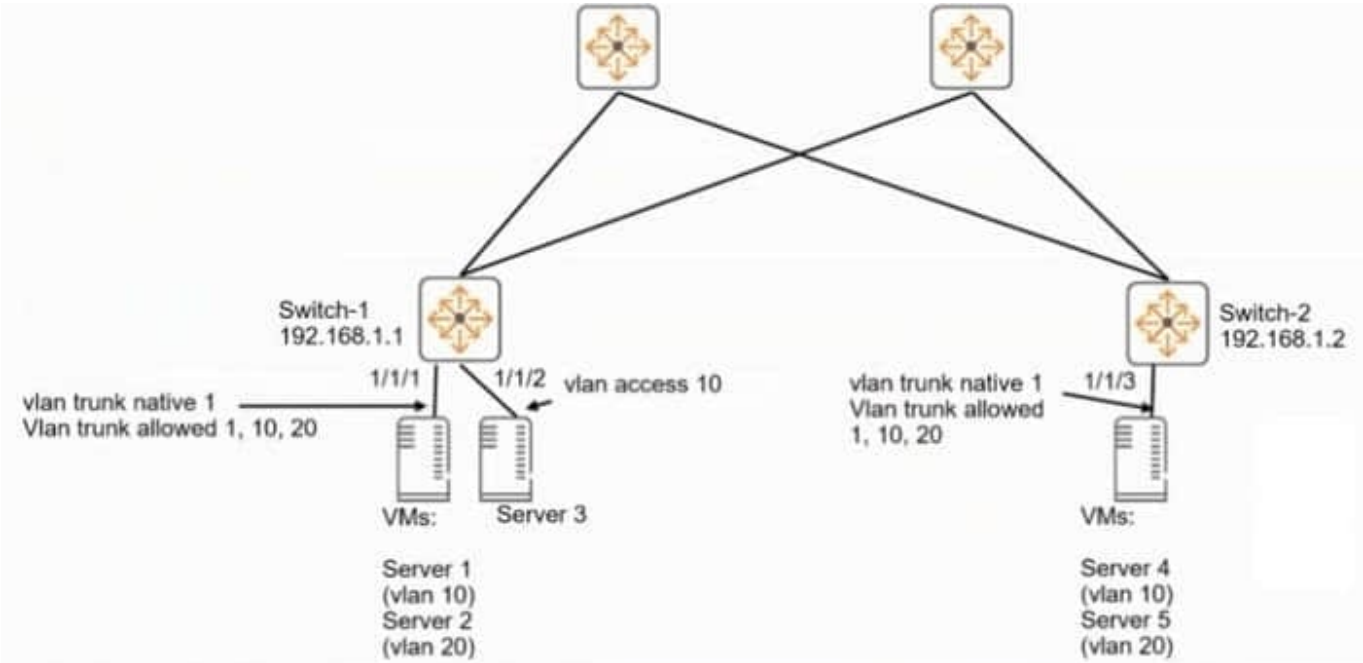


```
switch# show interface vxlan1 vteps
Source          Destination     Origin  Status       VNI   VLAN
192.168.1.1     192.168.1.2     static  Operational  1000  10
192.168.1.1     192.168.1.2     static   Operational 2000  20
```

Switch-1 and Switch-2 ate ArubaOS-CX switches that implement VXLAN WITHOUT Ethernet VPN (EVPN). Switch-2 uses the same VNI-to-VLAN mappings as Switch-1. Is this how the specified servers communicate?

Solution: Server 1 and Server 4 require routing services within the VXLANs to communicate with each other.

A. Yes

B. No

Correct Answer: B

The exhibit shows a network topology where Switch-1 and Switch-2 are ArubaOS-CX switches that implement VXLAN without Ethernet VPN (EVPN). Switch-2 uses the same VNI-to-VLAN mappings as Switch-1. The question asks how the specified servers communicate, which means Server 1 and Server 4. Server 1 and Server 4 are in different VLANs and different VNIs, which means they are in different layer 2 segments. To communicate with each other, they require routing services between the VXLANs. However, using Virtual Routing and Forwarding (VRF) to tunnel iSCSI traffic through the network spine on the same links that data traffic uses is not the correct way to provide routing services. VRF is a technology that creates multiple isolated Layer 3 domains on a physical network, each with its own routing table. VRF does not provide any benefits for iSCSI traffic, as it does not guarantee bandwidth, priority, or quality of service. VRF also adds overhead and complexity to the network configuration1. To provide routing services between the VXLANs, the correct way is to use VXLAN routing with EVPN or distributed anycast gateway (DAG). VXLAN routing with EVPN allows the switches to exchange MAC and IP information using BGP EVPN control plane, and to perform routing between different VNIs using a centralized or distributed model2. DAG allows the switches to act as anycast gateways for their local hosts, and to route traffic between different VNIs using a symmetric or asymmetric model3. Therefore, this does not correctly describe how the specified servers communicate.

QUESTION 8

Refer to the exhibit.

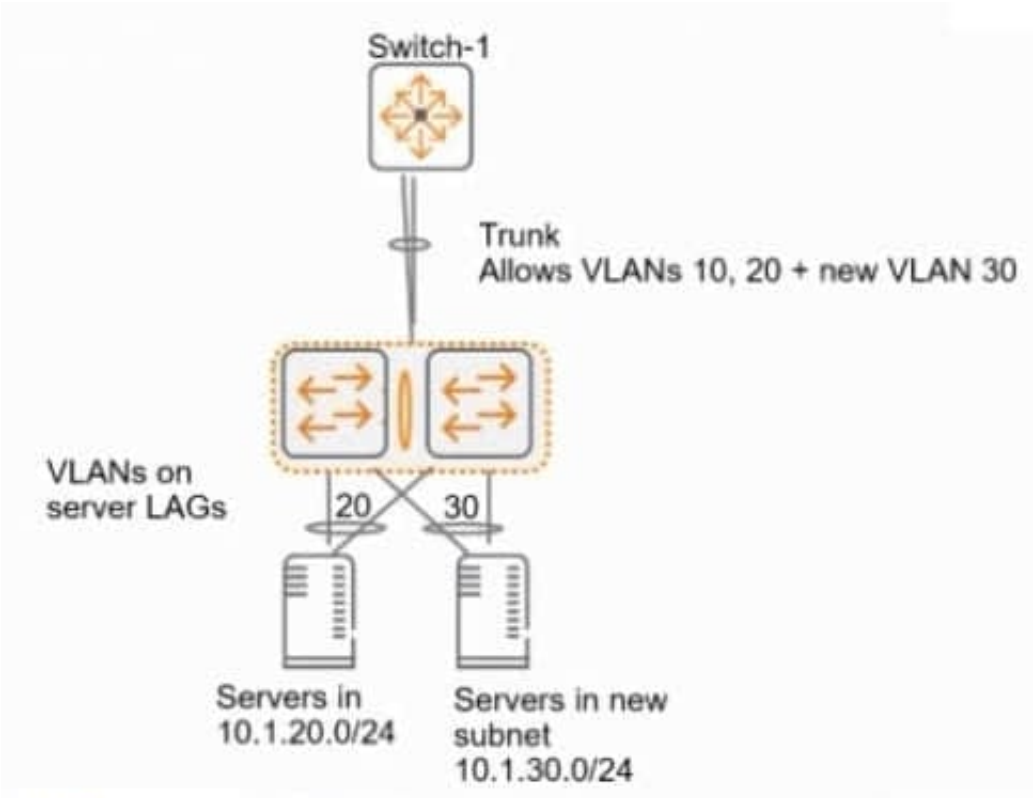


```
Switch-1 show ip route all-vrf

Displaying ipv4 routes selected for forwarding

'[x/y]' denotes [distance/metric]

10.1.10.0/24, vrf A
   via vlan10, [0/0], connected
10.1.10.1/32, vrf A
   via vlan10, [0/0], local
10.1.20.0/24, vrf B
   via vlan20, [0/0], connected
10.1.20.1/32, vrf B
   via vlan20, [0/0], local
```

You are adding a VLAN 30, subnet 10.0.30.0/24 to the network shown in the exhibit. (This network is simplified to just the relevant switches for this item.) This subnet belongs in VRF A, and you have added a Layer 3 VLAN 30 interface

attached to this VRF on Switch-1. You want to make the services in this VLAN available to devices in 10.1.20.0/24 in VRF B.

Is this part of a valid setup for meeting these requirements?

Solution: Add a route with this command: ip route 10.1.20.0/24 vlan20 vrf A

A. Yes

B. No

Correct Answer: B

Adding a route with this command: ip route 10.1.20.0/24 vlan20 vrf A is not part of a valid setup for meeting these requirements. This command would add a static route for 10.1.20.0/24 in VRF A, but it would not be able to reach VLAN 20 on Switch-2 because Switch-2 does not have a VLAN interface for VLAN 20 in VRF A. To make the services in VLAN 30 available to devices in 10.1.20.0/24 in VRF B, you need to use inter- VRF routing or route leaking between VRF A and VRF B on Switch-11.

QUESTION 9

Refer to the exhibit.

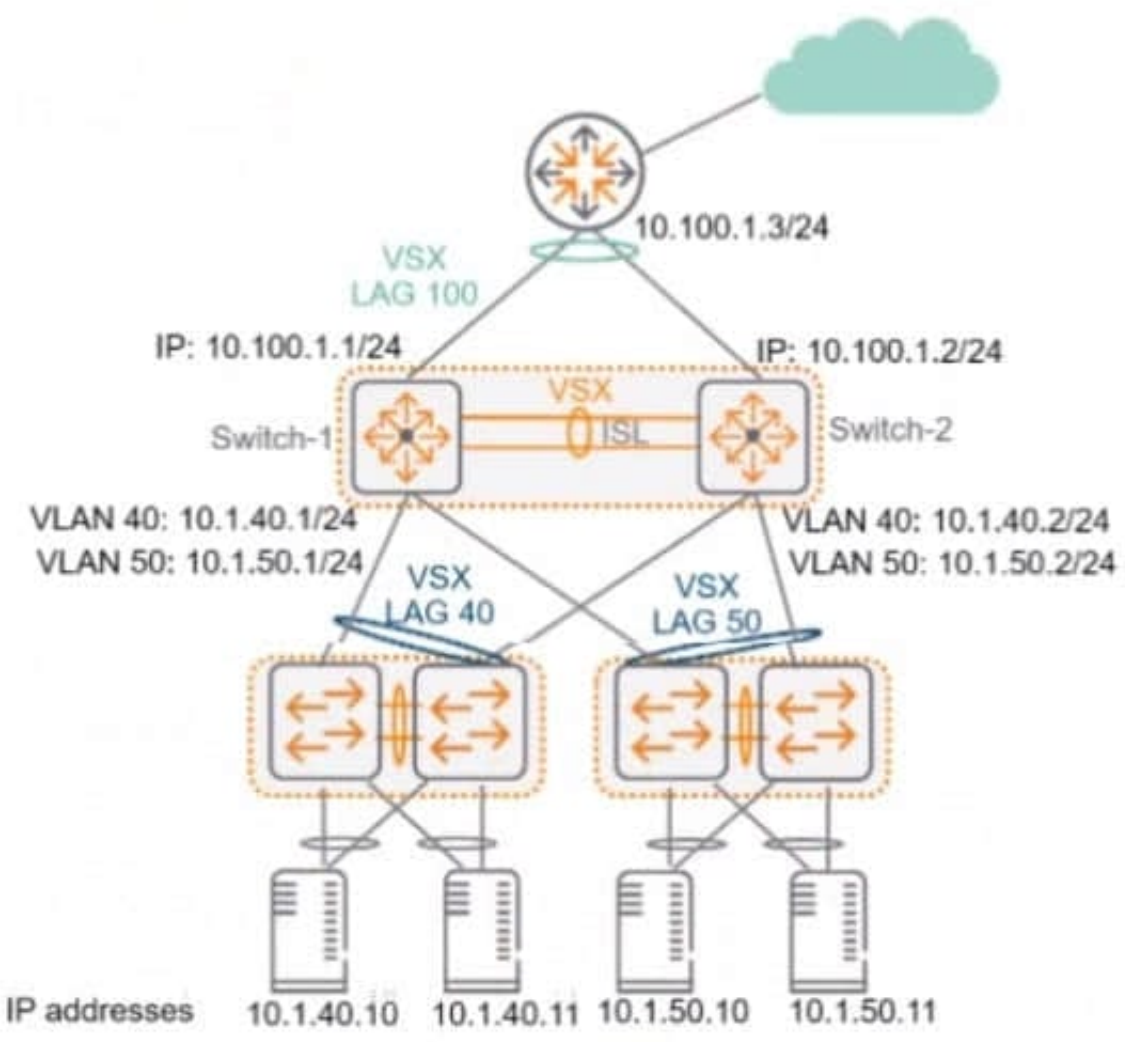

Switch-1, Switch-2, and the router run OSPF on LAG 100, which is a Layer 3 LAG. Does this correctly explain how to control how core-to-access traffic Is forwarded? Solution: To reduce the amount of traffic sent over the ISL between Switch-1 and Switch-2. enable Equal Cost Multi Path (ECMP) on both Switch-1 and Switch-2.

A. Yes

B. No

Correct Answer: B

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 traffic 1.

QUESTION 10

Is this how you should position switches in the ArubaOS-CX portfolio tor data center networks?

Solution: Deploy Aruba CX 8400 switches as core switches for very large three-tier data center networks.

A. Yes

B. No

Correct Answer: A

Deploying Aruba CX 8400 switches as core switches for very large three-tier data center networks is how you should position switches in the ArubaOS-CX portfolio for data center networks. ArubaOS-CX is an operating system that provides advanced features and automation capabilities for data center networks1. It runs on various switch models that are designed for different roles and scenarios in the data center1. Aruba CX 8400 switches are modular switches that offer high performance, scalability, and reliability for the core layer of very large three-tier data center networks1. The statement is true because it correctly describes how to position Aruba CX 8400 switches in the ArubaOS-CX portfolio for data center networks.

QUESTION 11

Is this a requirement for implementing Priority Flow Control (PFC) on an ArubaOS-CX switch interface? Solution: configuring trust of Cos on the interface

A. Yes

B. No

Correct Answer: A

Configuring trust of CoS on the interface is a requirement for implementing Priority Flow Control (PFC) on an ArubaOS-CX switch interface. PFC is a feature that allows a switch to pause traffic on a per-class basis using IEEE 802.1Qbb frames. To use PFC, the switch must trust the CoS values in the incoming frames and map them to priority groups and queues1.

QUESTION 12

Is this a use case for implementing Enhanced Transmission Selection (ETS) on an ArubaOS-CX switch? Solution: to enable the switch to assign the correct priority and bandwidth to traffic that it transmits to servers

A. Yes

B. No

Correct Answer: A

Enhanced Transmission Selection (ETS) is a network scheduling algorithm that allows the switch to assign different priority and bandwidth values to different traffic classes1. This can be useful for transmitting traffic to servers that have different requirements for latency, jitter, or throughput. For example, ETS can prioritize voice or video traffic over data traffic, or allocate more bandwidth to backup or replication traffic. ETS is configured using the Data Center Bridging Exchange (DCBx) protocol, which advertises the configuration to peer devices2. Therefore, implementing ETS on an ArubaOS-CX switch is a valid use case for enabling the switch to assign the correct priority and bandwidth to traffic that it transmits to servers.