HPE2-W09 Q&As

Aruba Data Center Network Specialist Exam

Pass HP HPE2-W09 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.leads4pass.com/hpe2-w09.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by HP Official Exam Center

QUESTION 1

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 broadcasts as unicast packets to each VTEP in the same VNI.

A. Yes

B. No

Correct Answer: A

Headend replication is a method of handling BUM traffic in VXLAN networks without EVPN, where the ingress VTEP replicates every BUM packet and sends them as a separate unicast to the remote egress VTEPs in the same VNI1. This avoids the need for multicast routing in the underlay network, but it can increase the load on the ingress VTEP. Therefore, this correctly describes how the VTEPs handle VXLAN traffic forwarding without EVPN.

QUESTION 2

Can you attach this type of ArubaOS-CX interface to a VRF?

Solution: A loopback interface

A. Yes

B. No

Correct Answer: B

A loopback interface can be attached to a VRF on an ArubaOS-CX switch. A loopback interface is a virtual interface that has an IP address assigned to it and is always up. A loopback interface can be attached to a VRF by using the vrf attach command under its configuration mode1.

QUESTION 3

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

A. Yes

B. No

Correct Answer: A

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

QUESTION 4

Is this a guideline for establishing a Virtual Switching Extension (VSX) Inter-Switch Link (ISL) between two ArubaOS-CX switches?

Solution: Use a link aggregation with multiple 40GbE links or multiple 100GbE links.

A. Yes

B. No

Correct Answer: A

Virtual Switching Extension (VSX) is a high-availability technology that allows two ArubaOS-CX switches to operate as a single logical device. VSX Inter-Switch Link (ISL) is a link between the two VSX switches that is used for both data plane and control plane traffic. It is recommended that the ISL link is a link aggregation with multiple 40GbE links or multiple 100GbE links to provide redundancy and bandwidth1. Therefore, this is a valid guideline for establishing a VSX ISL between two ArubaOS-CX switches.

QUESTION 5

The architect designs a spine and leaf network for a single data center that will use multiple leaf switches as Virtual Tunnel End Points (VTEP). The architect needs to select the type of Integrated Routing and Bridging (IRB) for the solution.

Is this statement about the IRB type true?

Solution: Asymmetric IRB requires a third L3 VNI to route packets between ingress and egress VTEPs.

A. Yes

B. No

Correct Answer: A

Asymmetric IRB requires a third L3 VNI to route packets between ingress and egress VTEPs is a true statement about the IRB type for a spine and leaf network for a single data center that will use multiple leaf switches as Virtual Tunnel End Points (VTEP). Asymmetric IRB is a method of routing traffic between different VXLAN segments using a centralized gateway. In this method, ingress VTEPs route the traffic to the gateway VTEP using a Layer 3 VNI, and egress VTEPs route the traffic to the destination networks using a Layer 2 VNI1. The Layer 3 VNI acts as a transit VNI for inter-VXLAN routing.

QUESTION 6

Is this a use case for disabling split-recovery mode on ArubaOS-CX switches in a Virtual Switching Extension (VSX) fabric?

Solution: In situations in which the primary switch fails and then reboots, you want to make the primary switch wait a period before it takes over as the primary switch.

A. Yes

B. No

Correct Answer: B

Virtual Switching Extension (VSX) is a high-availability technology that allows two ArubaOS-CX switches to operate as a single logical device. Split-recovery mode is a feature that prevents traffic loss when the Inter-Switch Link (ISL) goes outof-sync and keepalive subsequently fails. When split-recovery mode is enabled, the secondary VSX member disables its downstream links until it synchronizes with the primary member. When split-recovery mode is disabled, the secondary VSX member keeps its downstream links up even when it is out-of-sync with the primary member1. Disabling split-recovery mode does not affect how the primary switch waits a period before it takes over as the primary switch after a failure and reboot. The primary switch always takes over as the primary switch immediately when it comes back online, regardless of the split-recovery mode setting. To make the primary switch wait a period before it takes over as the primary switch, you need to configure a preemption delay on both VSX members1. Therefore, this is not a use case for disabling split-recovery mode on ArubaOS-CX switches in a VSX fabric.

QUESTION 7

An ArubaOS-CX is \ssmq DCBX on Interface 1/1/1. You enter this command:

show dcbx interface 1/1/1

Is this where you can see whether the connected converged network adapter (CNA) has accepted the application priorities advertised with DCBX?

Solution: in the Enhanced Transmission Selection (ETS) Local advertisement section

A. Yes

B. No

Correct Answer: B

In the Enhanced Transmission Selection (ETS) Local advertisement section is not where you can see whether the connected converged network adapter (CNA) has accepted the application priorities advertised with DCBX on an ArubaOS-CX switch interface. The ETS Local advertisement section shows the priority groups and bandwidth allocation that the switch advertises to the CNA, but it does not show whether the CNA has accepted them or not. To see whether the CNA has accepted the application priorities advertised with DCBX, you need to look at the Application Priority Local advertisement section, which shows the application protocol identifier (APPID) and priority values that the switch advertises to the CNA, and compare them with the Application Priority Remote advertisement section, which shows the APPID and priority values that the CNA advertises to the switch1.

QUESTION 8

Refer to the exhibit.

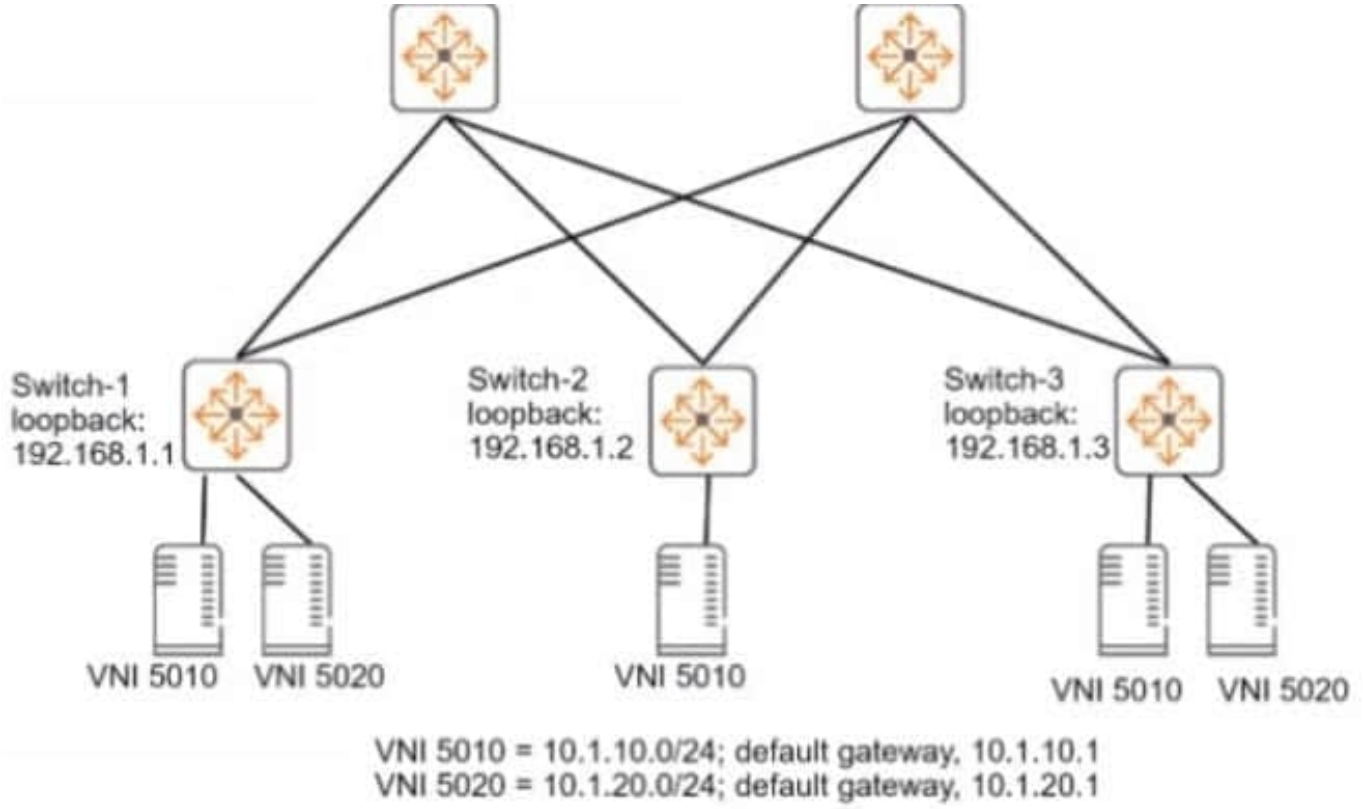

You need to set up an ArubaOS-CX switch to implement Virtual Extensible LAN (VXLAN) WITHOUT Ethernet VPN (EVPN). The exhibit Indicates which servers should be part of the same VXLANs and the desired VNIs for the VXLANs. Assume that the network is already configured to permit each ArubaOS-CX switch to reach each other switch\\'s loopback interface.

Is this part of the process for setting up VXLAN to meet the requirements?

Solution: On Switch-1, create two VXLAN interfaces, one with ID 5010 and one with 1D 5020; both VXLAN interfaces should use 192.168.1.1 as the source IP address.

A. Yes

B. No

Correct Answer: A

VXLAN is a feature of ArubaOS-CX that provides layer 2 connectivity between networks across an IP network1. VXLAN uses a 24-bit identifier called VXLAN Network Identifier (VNI) to segment the layer 2 domain1. VXLAN also uses a tunnel endpoint (VTEP) to encapsulate and decapsulate VXLAN packets1. A VXLAN interface is a logical interface that represents a VNI and is associated with a source IP address and a VRF1. To set up VXLAN without EVPN, you need to create VXLAN interfaces on each switch and configure static VTEP peers1. Based on the exhibit, Switch-1 needs to create two VXLAN interfaces, one with ID 5010 and one with ID 5020, to match the VNIs of the servers connected to it. Both VXLAN interfaces should use 192.168.1.1 as the source IP address, which is the loopback interface of Switch-1. Therefore, this is part of the process for setting up VXLAN to meet the requirements, and the correct answer is yes. For more information on VXLAN and EVPN, refer to the Aruba Data Center Network Specialist (ADCNS) certification datasheet2 and the EVPN VXLAN Guide for your switch model1.

QUESTION 9

Is this a best practice when positioning ArubaOS-CX switches in data center networks? Solution: Deploy Aruba CX 83xx switches as data center spine switches.

A. Yes

B. No

Correct Answer: A

Deploy Aruba CX 83xx switches as data center spine switches is a best practice when positioning ArubaOS-CX switches in data center networks. The Aruba CX 83xx switches are designed for data center spine or core roles, and they provide high performance, scalability, and resiliency. They can support various data center network architectures such as leaf-spine, three-tier, or collapsed core1.

QUESTION 10

Is this a way that Virtual Switching Extension (VSX) differs from Virtual Switching Framework (VSF)?

Solution: VSX permits admins to select which features to synchronize between members while VSF requires manual configuration of Identical features on each member of the VSF fabric.

A. Yes

B. No

Correct Answer: A

VSX permits admins to select which features to synchronize between members while VSF requires manual configuration of identical features on each member of the VSF fabric is a way that Virtual Switching Extension (VSX) differs from Virtual Switching Framework (VSF). VSX is a feature that provides active-active forwarding and redundancy for ArubaOS-CX switches. VSF is a feature that provides active-standby forwarding and redundancy for legacy campus switches. VSX allows admins to select which features to synchronize between members using an opt-in model, while VSF requires manual configuration of identical features on each member using a commander-member model1.

QUESTION 11

Your task is to configure an EVPN solution for a dual-stack IPv4 and IPv6 protocol in the overlay networks. Is this statement about EVPN and IPv6 correct? Solution: The IPv6 and IPv4 overlay networks can be encapsulated with VXLAN and transmitted through the underlay network.

A. Yes

B. No

Correct Answer: A

The IPv6 and IPv4 overlay networks can be encapsulated with VXLAN and transmitted through the underlay network is a true statement about EVPN and IPv6 for configuring an EVPN solution for a dual-stack IPv4 and IPv6 protocol in the overlay networks. VXLAN is a feature that provides Layer 2 extension over Layer 3 networks using UDP encapsulation. VXLAN can support both IPv4 and IPv6 overlay networks over an IPv4 or IPv6 underlay network2.

QUESTION 12

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-I and Switch-2 remains up. The Split-recovery mode is enabled. In this case the secondary switch shutdowns SVIs when keepalive is restored.

A. Yes

B. No

Correct Answer: A

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 loops 1.

QUESTION 13

Refer to the exhibit.

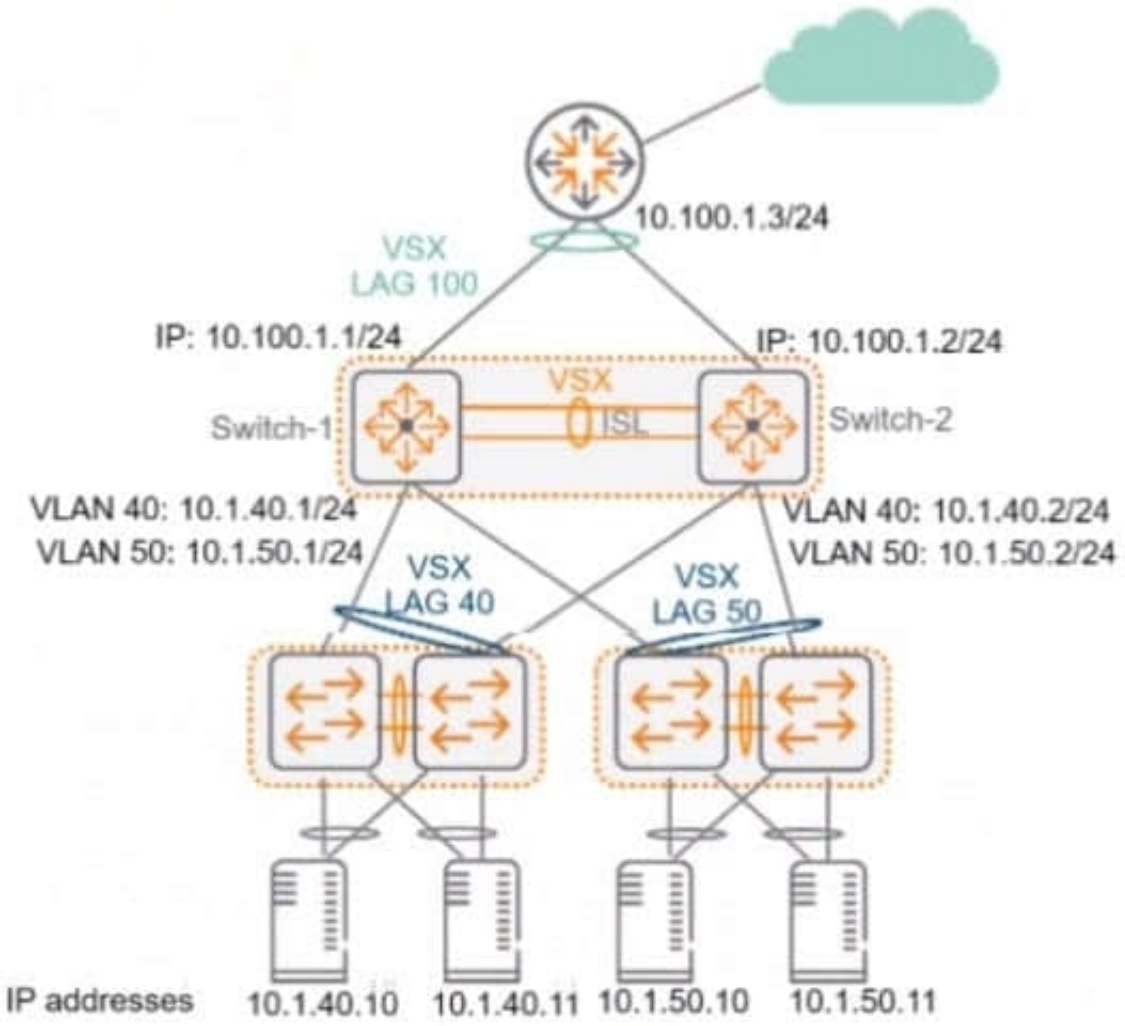

Switch-1, Switch-2, and the router run OSPF on LAG 100, which is a Layer 3 LAG. Does this correctly explain how to control how core-to-access traffic Is forwarded?

Solution: To force the router to use both links, ensure that active gateway is enabled on LAG 100 on both Switch-1 and Switch-2.

A. Yes

B. No

Correct Answer: B

The exhibit shows a network topology where Switch-1 and Switch-2 are part of a Virtual Switching Extension (VSX) fabric, and the router runs OSPF on LAG 100, which is a Layer 3 LAG. The question asks how to control how core-to-access traffic is forwarded, which means how the router chooses between the two links to Switch-1 and Switch-2. To force the router to use both links, ensuring that active gateway is enabled on LAG 100 on both Switch-1 and Switch-2 is not the correct solution. Active gateway is a feature that allows both VSX members to act as the default gateway for downstream devices, using a common virtual MAC address. Active gateway does not affect how upstream devices, such as the router, forward traffic to the VSX members1. To force the router to use both links, the correct solution is to configure equal-cost multi-path (ECMP) in OSPF on the router. ECMP is a feature that allows a router to load balance traffic across multiple paths with the same cost. ECMP can be configured using the maximum-paths command and specifying how many equal-cost paths the router should use2. Therefore, this does not correctly explain how to control how core-to-access traffic is forwarded.

QUESTION 14

Is this a difference between a typical data center network\\'s requirements and a typical campus network\\'s requirements?

Solution: Data center network traffic flows are typically east-west whereas while campus networks experience more north-south traffic.

A. Yes

B. No

Correct Answer: A

A data center network is a network that connects servers, storage devices, and other devices within a data center. A campus network is a network that connects buildings and users within a campus area, such as a university or an enterprise. Data center network traffic flows are typically east-west, which means they are between servers or devices within the data center. This is because data center applications often require high-speed communication and data exchange between servers for processing, analysis, or backup. Campus network traffic flows are typically north-south, which means they are between users or devices and external networks, such as the Internet or a wide area network (WAN). This is because campus users often access online services or resources that are hosted outside the campus network12. Therefore, this is a valid difference between a typical data center network\\'s requirements and a typical campus network\\'s requirements.

QUESTION 15

Refer to the exhibit.

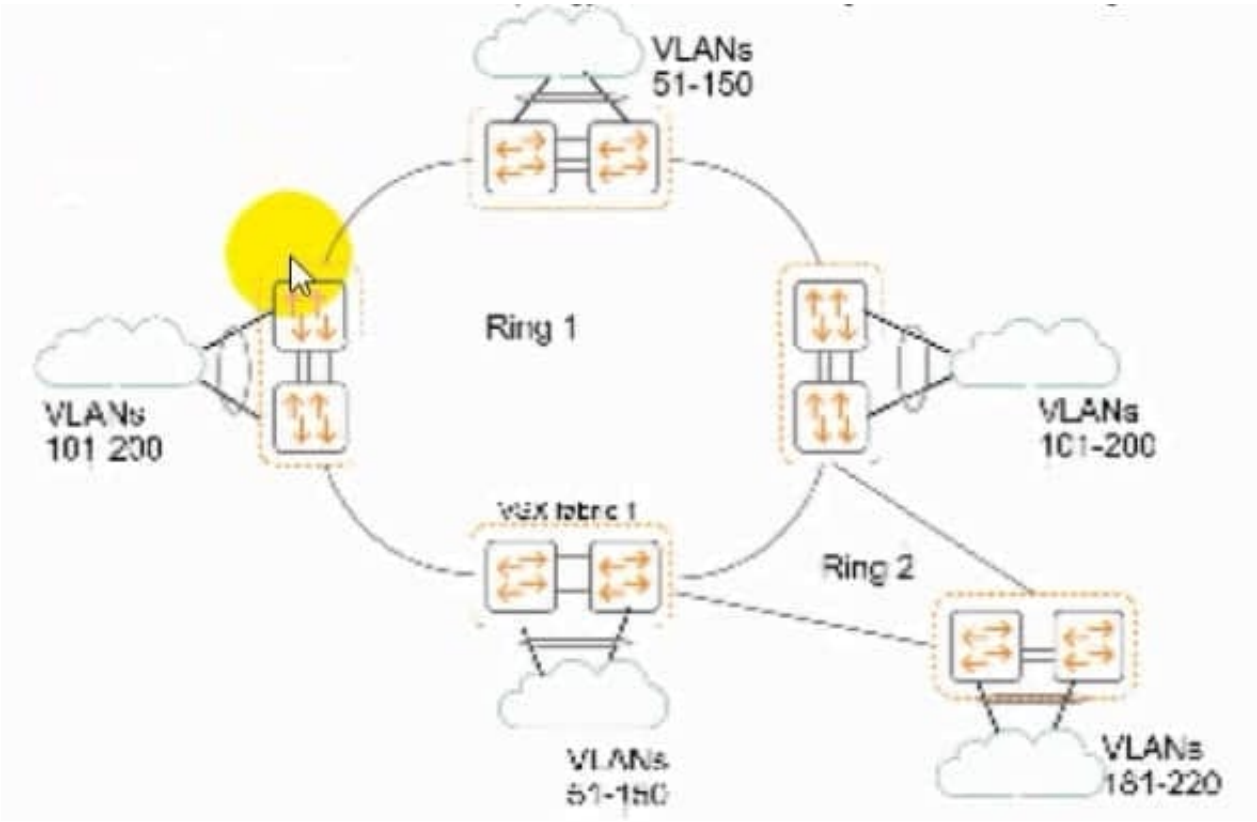

which shows the topology tot an Ethernet Ring Protection Switching (ERPS) solution.

Is this a valid design for the control and protected VLANs on the VSX fabric 1 switches?

Solution: Ring 1, instance 1:

control VLAN: 1000 protected VLANs: 51-135 Ring 1, Instance 2:

control VLAN: 1000 protected VLANs: 136-220 Ring 2, Instance 1: control VLAN:

1001 protected VLANs: 181 -200 Ring 2, Instance 2: control VLAN: 1001 protected VLANs: 201 -220

A. Yes

B. No

Correct Answer: A

Ring1, instance 1: control VLAN: 1000 protected VLANs: 51-135 Ring 1,Instance2: control VLAN: 1001 protected VLANs: 136-220 Ring 2, Instance l: control VLAN: 1002 protected VLANs: l8l -200 Ring 2, Instance2: control VLAN: l003 protected VLANs:201 -220 is a valid design for the control and protected VLANs on the VSX fabric l switches for an Ethernet Ring Protection Switching (ERPS) solution. The control VLANs are unique for each ring instance and do not overlap with any protected VLANs. The protected VLANs are also unique for each ring instance and do not overlap with any control VLANs2.