

DUMPSBOSS.

Advanced HPE Edge-to-Cloud Solutions

HP HPE7-V01

Version Demo

Total Demo Questions: 6

Total Premium Questions: 67

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Explain and describe the HPE Edge-to-Cloud strategy	6
Topic 2, Gather and analyze customer business and technical requirements	1
Topic 3, Design HPE Edge-to-Cloud solutions	58
Topic 4, Present HPE Edge-to-Cloud solutions	2
Total	67

QUESTION NO: 1

Your customer has an ArubaOS-CX leaf-spine fabric and a mixed server environment that includes HPE and non-HPE servers running VMware workloads.

They want to simplify fabric deployment and improve visibility of virtual-machine VLAN connectivity.

Which two capabilities can Aruba Fabric Composer provide? (Select two.)

- A. Automate provisioning of an ArubaOS-CX spine-leaf fabric.
- B. Provide VMware VLAN visibility for workloads on non-HPE servers.
- C. Manage and monitor HPE Alletra Storage MP B10000 arrays.
- D. Apply firmware baselines to HPE ProLiant servers.
- E. Configure Fibre Channel zoning on external SAN switches.

ANSWER: A B

Explanation:

Aruba Fabric Composer provides centralized automation and orchestration for ArubaOS-CX data-center fabrics. It can accelerate provisioning of spine-leaf topologies by applying consistent fabric, interface, VLAN, routing, and overlay configuration through centralized workflows.

Fabric Composer can also integrate with VMware vCenter to provide visibility into virtual-machine networking constructs, including VLAN and port-group relationships, regardless of the server vendor hosting the VMware workloads. It does not manage HPE Alletra storage arrays or perform server firmware lifecycle operations.

QUESTION NO: 2

Which statement about HPE SimpliVity is true?

- A. HPE SimpliVity 380 Gen11 cluster with compute nodes requires the same CPU in all nodes.
- B. Mixing AMD and Intel chassis in the same federation is fully supported.
- C. A single AMD-based HPE SimpliVity node can support up to 4 non-HPE SimpliVity servers.
- D. A single Intel-based HPE SimpliVity node can support up to 8 non-HPE SimpliVity servers.

ANSWER: D

Explanation:

A single Intel-based HPE SimpliVity node can support up to 8 non-HPE SimpliVity servers. HPE SimpliVity supports compute-node designs in which standard, non-SimpliVity servers run workloads while consuming the hyperconverged storage, data services, backup, and disaster-recovery capabilities delivered by SimpliVity nodes. This architecture is useful when an organization needs to scale compute capacity independently without adding data-optimization hardware with every server. For the applicable Intel-based SimpliVity platform design, HPE specifies a maximum attachment ratio of eight non-HPE SimpliVity compute servers for each SimpliVity node. The ratio is a supportability and sizing limit: it helps ensure that compute demand remains aligned with the storage and data-service resources contributed by the SimpliVity nodes. A complete implementation must still follow the relevant HPE support matrix, validated hypervisor configuration, software version requirements, and cluster sizing guidance. HPE describes SimpliVity as an HCI platform that combines compute, storage, and VM-centric data services, while its documentation library provides the platform-specific deployment and configuration guidance used to validate these designs. See [HPE SimpliVity](#) and the [HPE SimpliVity documentation library](#).

QUESTION NO: 3

Your customer is designing an HPE SimpliVity stretched cluster across two data centers. The customer wants the environment to continue operating safely if one site fails or if communication between the two sites is interrupted.

Which two statements are true regarding the arbiter design? (Select two.)

- A. The arbiter should be deployed in the same site as the primary SimpliVity nodes.
- B. The arbiter should be deployed in a third failure domain separate from both cluster sites.
- C. The arbiter stores a replicated copy of all protected virtual-machine data.
- D. The arbiter provides an independent quorum function to help prevent split-brain conditions.
- E. The arbiter replaces the requirement for redundant connectivity between the two cluster sites.

ANSWER: B D

Explanation:

An arbiter should be deployed in a third failure domain, separate from both sites hosting the stretched-cluster nodes. This independent placement allows it to provide quorum information if the two data centers can no longer communicate with each other.

The arbiter does not host application virtual machines or SimpliVity data. Its purpose is to provide an independent voting mechanism that helps prevent split-brain behavior and supports a safe decision about which site can continue operations. Deploying it within either stretched-cluster site would reduce the required fault-domain independence.

QUESTION NO: 4

Refer to exhibit from a customer's Aruba CX switch.

show interface vxlan vni vteps

VNI : 150 VLAN : 150

Routing : Disabled VRF : default

Vxlan-status : operational

VTEPs

Origin Source Destination VTEP-STATUS

static 172.17.255.2 172.17.255.1 static operational

Which statement is true in this scenario?

- A. A static VXLAN setup needs to be created between the peer switches.
- B. An EVPN VXLAN setup will dynamically establish the VTEPs between peer switches.
- C. A Foundation license needs to be applied to the switches to scale the VXLAN configuration.
- D. A VTEP setup requires the same model of the CX switch on each location.

ANSWER: A

Explanation:

A static VXLAN setup needs to be created between the peer switches. The command output explicitly identifies the VTEP origin as *static* and reports the VTEP status as *static operational*. This confirms that the VXLAN tunnel endpoint relationship for VNI 150 is configured manually between the local VTEP source address, 172.17.255.2, and the remote peer VTEP destination address, 172.17.255.1. The operational state also confirms that the required underlay IP connectivity and VXLAN configuration are functioning for that configured peer.

The VNI-to-VLAN association shown in the output maps VNI 150 to VLAN 150, while routing is disabled in the default VRF. This is consistent with a Layer 2 VXLAN overlay, where the VLAN's traffic is encapsulated across the IP underlay to the explicitly configured remote VTEP. Static VXLAN is appropriate when VTEP peers are defined directly and the deployment does not rely on a dynamic control plane to discover or advertise VTEP membership. ArubaOS-CX VXLAN documentation describes configuring VXLAN interfaces, VNIs, source VTEP addresses, and static VTEP peers as part of a static VXLAN design. See the ArubaOS-CX [VXLAN Guide](#) and the [static VXLAN configuration guidance](#).

QUESTION NO: 5

A customer is about to do a server refresh with the intent of using HPE ProLiant DL360 Gen11 servers, with CPUs exceeding 270 watt power usage. The customer has existing third-party, 42U, air-cooled racks they are planning to reuse.

Which cooling option will allow your customer to install the servers with their existing racks, with the least amount of changes to the setup?

- A. Use of liquid-cooled racks.
- B. Use of closed-loop liquid cooling.
- C. Use of direct liquid cooling.
- D. Use of high performance fans.

ANSWER: B

Explanation:

Use of closed-loop liquid cooling is the appropriate choice because it supports the thermal requirements of HPE ProLiant DL360 Gen11 configurations using processors above 270 W while preserving the customer's existing air-cooled rack environment. A closed-loop liquid-cooling solution is contained within the server cooling assembly: heat is transferred from the processor through a sealed liquid loop and rejected through the server's radiator and fans. Consequently, it does not require the rack to be converted into a facility-water or direct-liquid-cooling deployment. This is especially valuable for a refresh in which the customer intends to retain third-party 42U racks and minimize changes to rack plumbing, coolant distribution units, and data-center water infrastructure. HPE's DL360 Gen11 configuration guidance identifies processor thermal limits and applicable cooling requirements; the required cooling option must be selected when configuring high-power CPU SKUs. The solution should also be validated through HPE's supported configuration tools to ensure the selected processors, closed-loop liquid-cooling kit, power supplies, and rack installation meet the platform's requirements. See the [HPE ProLiant DL360 Gen11 Server QuickSpecs](#) and the [HPE ProLiant server portfolio](#) for platform and configuration information.

QUESTION NO: 6

Your customer is refreshing its campus wired and wireless infrastructure. The customer wants to avoid a conventional capital purchase and also wants an external provider to perform ongoing LAN operations.

Which two statements describe how HPE GreenLake for Aruba Networking and associated managed services can address these requirements? (Select two.)

- A. The customer can consume networking infrastructure through a subscription-based, cloud-like model.
- B. Aruba experts can provide managed operations with continuous hardware monitoring.
- C. The offering is limited to Aruba software licenses and requires customer-owned hardware.
- D. The offering can manage third-party wireless infrastructure as an Aruba managed service.
- E. The offering only provides a monitoring portal and does not include infrastructure consumption options.

ANSWER: A B

Explanation:

HPE GreenLake for Aruba Networking provides a cloud-like consumption model for networking infrastructure, allowing the customer to consume hardware, software, and associated services without a traditional upfront capital acquisition model.

Managed services can complement the consumption model by providing Aruba experts for ongoing operations, including continuous hardware monitoring and operational support. GreenLake for Aruba Networking is not limited to software licensing or to use as a monitoring portal.