



HP MASE - EdgCldAr HPE7-V01 Practice Questions

Shared by Peters on 17-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: MultipleChoice

Your customer has a combination of HPE and non-HPE equipment.

What is the correct functionality available in Aruba Fabric Composer for this customer's environment?

Options:

- A- Management and monitoring of HPE Alletra Storage MP B10000 arrays.
- B- Management and monitoring of HPE Synergy Virtual Connect.
- C- Configuration of SAN zoning on ArubaOS-CX 8400 switches.
- D- Acceleration of spine-leaf fabric provisioning on ArubaOS-CX switches.

Answer:

D

Explanation:

The correct answer is D because Aruba Fabric Composer is designed for data-center fabric automation, especially ArubaOS-CX based spine-leaf and EVPN-VXLAN environments. It discovers data-center infrastructure and automates the provisioning and ongoing configuration of both spine-and-leaf fabrics and Layer 2 two-tier topologies. That makes it directly relevant to customers with mixed HPE and non-HPE environments where repeatable fabric deployment, validation, and operational consistency are required. Option A is incorrect because HPE Alletra Storage MP B10000 management belongs to HPE storage management tools such as Data Services Cloud Console and array-specific integrations, not Aruba Fabric Composer. Option B is incorrect because HPE Synergy Virtual Connect is managed through HPE OneView / Synergy Composer, not Aruba Fabric Composer. Option C is also not the best answer because SAN zoning is normally performed through Fibre Channel SAN fabric tools or storage-networking workflows, not through Aruba Fabric Composer as a generic function of ArubaOS-CX 8400. The exam logic is clear: Fabric Composer's value is accelerating, standardizing, and validating data-center network-fabric provisioning. Reference/topics: Aruba Fabric Composer, data-center fabric automation, ArubaOS-CX, spine-leaf provisioning.

Question 2

Question Type: MultipleChoice

Your customer is in the early stages of an AI project. They are requesting a POC demonstrating how they can utilize on-prem data with a limited amount of rack space, in an air-cooled environment.

Performance with the AI use case of Retrieval-Augmented Generation (RAG) will be the primary factor in their decision making.

Which HPE Compute offering would be suitable in this scenario?

Options:

- A- HPE ProLiant Compute DL380a Gen12 with NVIDIA H200 NVL GPUs.
- B- HPE ProLiant Compute DL384 Gen12 with NVIDIA GH200 NVL2.
- C- HPE ProLiant Compute DL320 Gen11 with NVIDIA L4 GPUs.
- D- HPE Cray XD670 with NVIDIA H200 Tensor Core GPUs.

Answer:

B

Explanation:

The correct answer is B. The customer is not asking for the lowest-cost AI inference server; the scenario explicitly states that RAG performance is the primary decision factor, while also requiring limited rack space and air cooling. HPE ProLiant Compute DL384 Gen12 with NVIDIA GH200 NVL2 is the strongest fit because it is designed for high-performance enterprise AI and provides coherent CPU/GPU memory architecture with very high memory bandwidth, which is especially valuable for RAG and larger language-model workflows. HPE product information describes the DL384 Gen12 as a 2U air-cooled 19-inch rack server with NVIDIA GH200 NVL2, and HPE states that the platform provides outstanding coherent memory and memory bandwidth to run larger models faster and more cost-effectively. Option A is powerful, but the DL380a Gen12 is a larger GPU platform and is less optimized for the "limited rack space" RAG-specific requirement than the DL384. Option C is compact but lower-performance, making it weaker when performance is the primary factor. Option D is excessive for an early-stage on-prem POC. Reference/topics: AI workload sizing, RAG architecture, DL384 Gen12, GH200 NVL2, air-cooled AI infrastructure.

Question 3

Question Type: MultipleChoice

Your customer is currently using an HPE SimpliVity 380 Gen10 solution and is planning for a generation refresh to an HPE SimpliVity 380 Gen11.

The existing cluster is planned to be maintained as a test and dev platform for another 24--36 months while production workloads are moved to the new environment.

The new solution will be purchased as HPE GreenLake for Private Cloud Business Edition, and the customer would like to include the existing Gen10 cluster within the same management platform.

Which statement is true about combining the management of both the old and new generations?

Options:

- A- The new PCBE platform can be downgraded for on-prem management for interoperability with the Gen10.
- B- The existing SimpliVity environment needs to be redeployed as PCBE to co-exist with the new solution.
- C- The new PCBE platform can manage the old platform by utilizing the Data Services Connector OVA.
- D- The existing SimpliVity environment can be imported into the HPE GreenLake platform using Data Services Connector.

Answer:

D

Explanation:

The correct answer is D because existing HPE SimpliVity systems can be brought under HPE GreenLake / Data Services Cloud Console management by using the Data Services Connector. The connector is deployed as a virtual machine, registered with HPE GreenLake, and used to enable or import supported HPE SimpliVity systems into the cloud-managed PCBE operating model. This fits the customer's transition scenario exactly: production is moving to the new Gen11 PCBE environment, but the older Gen10 SimpliVity cluster remains useful for test and development. Option A is incorrect because the new PCBE platform should not be downgraded to preserve legacy-style management. Option B is also wrong because the existing cluster does not need to be redeployed as a brand-new PCBE installation. Option C is directionally close but imprecise: the management plane is HPE GreenLake / DSCC using Data Services Connector, not the new PCBE system "managing" the old platform directly. Reference/topics: HPE GreenLake, Private Cloud Business Edition, Data Services Connector, SimpliVity onboarding and lifecycle management.

Question 4

Question Type: MultipleChoice

One of your customers is planning to refresh their current campus wired and wireless networking equipment. They are currently moving to a pay-per-use model for their whole IT solution.

They are wondering if there is a possibility to get new networking equipment in this model without capital expenses. They are also interested in moving LAN management to an external service provider. You are presenting HPE GreenLake for Aruba Networking.

Which statement is true for this customer situation?

Options:

- A- Aruba networks can be managed by Aruba experts with 24/7 hardware monitoring.
- B- HPE GreenLake for Aruba Networking only offers Aruba software licenses in a subscription-based format.
- C- HPE GreenLake for Aruba Networking is the central monitoring cloud software for Aruba equipment only.
- D- HPE GreenLake for Aruba Networking offers HPE experts support for third-party wireless devices.

Answer:

A

Explanation:

The correct answer is A because the customer is asking for a network refresh without traditional capital purchase and also wants externalized operational management. HPE GreenLake for Networking is positioned as a Network-as-a-Service model that allows customers to consume HPE Aruba Networking capabilities in a cloud-like operating model rather than buying every component upfront. HPE's GreenLake for Networking QuickSpecs describe it as a comprehensive NaaS offering for consuming HPE Aruba Networking Edge Services Platform capabilities in a cloud-like way. HPE Managed Services can also provide remote monitoring, management, and optimization of supported infrastructure, which aligns with the customer's desire to move LAN management to an external provider. Option B is too narrow because the offer is not merely a software-license subscription. Option C confuses HPE GreenLake for Networking with Aruba Central-style cloud management. Option D is incorrect because the service is built around HPE Aruba Networking solutions, not generic expert support for arbitrary third-party wireless platforms. Reference/topics: HPE GreenLake for Networking, NaaS positioning, HPE Managed Services, business-outcome mapping.

Question 5

Question Type: MultipleChoice

You are designing an HPE solution for SAP HANA. You are considering an as-a-service approach as one of the variants to be proposed to the customer.

What is true regarding HPE GreenLake for SAP HANA?

Options:

- A- HPE GreenLake for SAP HANA requires an onsite HPE Alletra Storage MP X10000.
- B- HPE GreenLake for SAP HANA offers billing based on GB of RAM used by SAP HANA.
- C- HPE GreenLake for SAP HANA can be used for SAP HANA instances running in the public cloud.
- D- HPE GreenLake for SAP HANA offers billing based on the amount of users using SAP HANA.

Answer:

B

Explanation:

The correct answer is B. HPE GreenLake for SAP HANA is positioned as an on-premises, consumption-based infrastructure model for SAP HANA rather than a public-cloud runtime service or a user-count licensing model. The key sizing and financial driver is the SAP HANA memory footprint because SAP HANA is an in-memory database platform. HPE describes the GreenLake SAP HANA model as a complete pay-per-use infrastructure solution with an on-premises appliance, operating system, and services, and states that the simplified monthly invoice is based on the total gigabytes of RAM used by the SAP HANA system. Option A is incorrect because the service is not defined by a mandatory HPE Alletra Storage MP X10000 array requirement. Option C is incorrect because this offer is specifically built to deliver cloud-like consumption while keeping SAP HANA infrastructure under customer control, typically on-premises. Option D is incorrect because SAP HANA infrastructure consumption is not metered by end-user count. Reference/topics: HPE GreenLake consumption model, SAP HANA workload sizing, as-a-service positioning, financial-impact presentation.

Question 6

Question Type: MultipleChoice

Your customer plans to replace a legacy HPE BladeSystem solution with the HPE Synergy platform.

The customer requirements include:

30 compute nodes to support their virtualization workloads

Redundancy within the management subsystem

An in-frame storage solution to support vSAN

A guidance/automation technology with HPE Aruba to virtualize WWNs

Simplified management with a single logical enclosure

You are reviewing the following proposal:

3 x HPE Synergy 12000 frames

2 x HPE Virtual Connect SE 100Gb F32 Modules for Synergy

4 x HPE Synergy 50Gb Interconnect Link Modules

3 x HPE D3940 Storage Modules with maximum number of drives

6 x HPE Composer2 + 30 x HPE OneView Advanced licenses

30 x SY 480 Gen11 compute modules, each with 2 CPUs and appropriate mezzanine cards

Which statement about the proposal is true to help meet the customer requirements?

Options:

A- Additional HPE OneView Advanced licenses should be added for the HPE Virtual Connect modules and HPE D3940 Storage Modules.

B- HPE 32Gb Fibre Channel switches should be added to the proposal to enable WWN virtualization and profile mobility.

C- Each frame should be equipped with two HPE 12Gb SAS switches to provide connectivity between compute nodes and D3940 Storage Module.

D- 3 x HPE Composer 2 modules can be removed from the proposal, as you need one HPE Composer2 module per HPE Synergy frame for enhanced availability.

Answer:

C

Explanation:

The correct answer is C because HPE Synergy D3940 is an in-frame direct-attached storage

module, and compute modules access it through HPE Synergy 12Gb SAS Connection Modules. For a redundant in-frame storage design, each Synergy frame that contains a D3940 storage module should include two SAS connection modules to provide redundant paths between the compute modules and the storage module. This is especially important for virtualization and software-defined storage designs such as vSAN, where storage path availability is a core design requirement. Option A is incorrect because HPE OneView Advanced licensing is associated with managed server infrastructure licensing, not separate additional licenses for every Virtual Connect or D3940 module. Option B is not required for this in-frame DAS design; Fibre Channel switches are used for SAN connectivity, not for D3940 SAS connectivity. Option D misstates Composer design: Composer redundancy is planned at the management-appliance level, not as a simple one-per-frame removal rule. This aligns with the course's Synergy, composable infrastructure, virtualization, and storage-networking design objectives.

Question 7

Question Type: MultipleChoice

Your customer has HPE Superdome Flex running VMware ESXi 8.x. Current setup consists of:

One Rack Management Controller

One HPE Superdome Flex Base chassis

One HPE Superdome Flex Expansion chassis

The existing complex has a single nPartition. The customer plans to expand the existing nPartition with two additional expansion chassis.

Which statement about the planned expansion is true?

Options:

- A- If the existing nPartition is managed using HPE OneView, it must be removed before an expansion.
- B- Customer has to replace one of the expansion chassis with base chassis, to provide I/O redundancy.
- C- In order to add two additional chassis, customer has to power off the existing nPartition.
- D- Existing nPartition must be placed in the maintenance mode using Rack Management Controller CLI.

Answer:

C

Explanation:

The correct answer is C because adding additional expansion chassis changes the physical hardware boundary of the HPE Superdome Flex complex and the resources assigned to the existing nPartition. Such a change cannot be treated like a live virtual-machine operation or a simple online software adjustment. The existing nPartition must be powered off before the physical expansion and partition-resource changes can be completed safely and consistently. Option A is incorrect because HPE OneView management does not require removing the nPartition from management merely to perform a hardware expansion; HPE OneView supports management of Superdome Flex-class infrastructure. Option B is incorrect because adding expansion chassis does not require replacing an expansion chassis with another base chassis for I/O redundancy in the way stated. Option D is also insufficient because placing the partition into a maintenance mode from the Rack Management Controller CLI is not equivalent to powering off the partition for a chassis-level expansion. The course includes Superdome Flex and scale-up compute designs for database and mission-critical workloads, so this question validates correct lifecycle planning for large HPE scale-up systems.

Question 8

Question Type: MultipleChoice

Your customer has a networking environment based on Aruba CX 8325 switches configured with Quality of Service settings that give higher priority to iSCSI traffic between HPE ProLiant servers and HPE Alletra 6000 storage.

They are planning to add HPE Synergy frames with HPE Virtual Connect modules that will connect to the same Aruba infrastructure. The servers will be deployed in different VLANs. They need to update their environment to provide a lossless fabric for RoCE traffic that will traverse between the HPE Synergy frames and through multiple switches.

Your customer has two sets of Synergy frames in different logical enclosures connected with Aruba CX 8325 switches. They want to update their environment to provide a lossless fabric for RoCE traffic.

What is the correct environmental or design change that will meet the customer's new requirement?

Options:

- A- Use RoCEv1 and configure DCBX on the switch facing the CNA in the servers.
- B- Use RoCEv2 and properly design networking map for the Layer 3 interfaces on the path.
- C- The Aruba CX 8325 switches must be changed to Aruba CX 8400 to support Data Center Bridging eXchange (DCBX).
- D- The HPE Synergy frames must be connected to different switches for RoCE traffic and

Ethernet traffic.

Answer:

B

Explanation:

The correct answer is B. RoCEv2 is the appropriate choice when RDMA traffic must traverse routed Layer 3 networks, multiple VLANs, and multiple switches. RoCEv1 is Layer 2 only and is therefore unsuitable when the design requires traffic to move between different routed segments or logical enclosures through the data-center network. A lossless RoCEv2 design must be planned end to end: priority flow control, ECN, QoS classification, MTU consistency, and routing paths must be aligned across the server adapters, HPE Synergy Virtual Connect domain, and Aruba CX fabric. Option A is incorrect because RoCEv1 cannot satisfy the stated Layer 3 traversal requirement. Option C is incorrect because Aruba CX 8325 is a data-center switch family capable of participating in advanced data-center designs; replacing it with CX 8400 is not inherently required just to support DCBX-style fabric behavior. Option D creates artificial separation and does not solve the multi-switch lossless-fabric requirement. Reference/topics: RoCEv2, lossless Ethernet, Aruba CX data-center switching, HPE Synergy Virtual Connect networking, workload-driven network design.

Question 9

Question Type: MultipleChoice

Your customer is planning to deploy several HPE ProLiant DL360 Gen11 servers with Windows Server 2025, for a bare-metal use case with no hypervisors.

The customer is concerned about the time it takes to scan and deploy the servers for new firmware and driver updates.

What HPE management tool can be used in this scenario to apply firmware updates to servers, with direct integration to HPE repositories?

Options:

- A- HPE OneView
- B- HPE Compute Ops Management
- C- HPE iLO Amplifier
- D- Microsoft Endpoint Configuration Manager

Answer:

B

Explanation:

The correct answer is B. HPE Compute Ops Management is the correct cloud-based management service for fleet-level HPE ProLiant lifecycle operations, including firmware baseline management and update execution. HPE describes Compute Ops Management as supporting HPE ProLiant ML/DL/MicroServer Gen10, Gen10 Plus, and Gen11 servers with iLO 5 and later, and includes automatic firmware-baseline update capability. HPE also provides firmware-update workflows for servers managed by Compute Ops Management. This fits the scenario precisely: bare-metal DL360 Gen11 servers, Windows Server installed directly, no hypervisor layer, and a requirement to reduce the operational time required to scan and deploy firmware/driver updates. HPE OneView is better suited for infrastructure managed through OneView domains and profiles, especially composable or converged environments. HPE iLO Amplifier is not the current strategic answer for GreenLake-era ProLiant fleet lifecycle management. Microsoft Endpoint Configuration Manager can manage operating system software, but it is not the HPE-native server firmware lifecycle tool with direct HPE repository integration. Reference/topics: HPE Compute Ops Management, ProLiant Gen11 lifecycle management, firmware baselines, HPE GreenLake cloud operations.

Question 10

Question Type: MultipleChoice

Your customer is planning a Scality RING solution. Approximately 750 TB of data must be stored on the SDS. The customer prefers SSD or NVMe drives.

What is true regarding this solution?

Options:

- A- HPE servers, for example, such as Alletra Storage Server 4120, can be used for Scality RING deployments.
- B- The minimum capacity for ordering Scality RING perpetual licenses is 1000TB.
- C- Scality RING Single Site Hardware Lifetime licenses are provided at a lower cost compared to perpetual licenses.
- D- Scality licenses cannot be purchased from HPE.

Answer:

A

Explanation:

The correct answer is A because Scality RING is positioned by HPE as a scale-out object and file storage solution built on qualified HPE server platforms, including HPE Alletra Storage Server models. For a customer planning approximately 750 TB of usable software-defined storage, the design focus is the validated HPE server configuration, drive type, usable-capacity target, and Scality license sizing. HPE documentation describes the Scality RING solution as consisting of HPE server systems, Scality license and maintenance, and installation services, with licensing based on unique usable data capacity. The minimum RING configuration is not 1000 TB; HPE references a lower starting usable-capacity threshold, so option B is incorrect. Option D is also incorrect because Scality solutions and licenses are explicitly available through HPE solution packaging. Option C is not the best design statement for this customer use case because the key technical validation is platform support and capacity alignment, not a generic claim about lower-cost licensing. This maps directly to the course topic covering scale-out unstructured data platforms, including Scality, Qumulo, Cohesity, and related HPE storage-server designs.

Question 11

Question Type: MultipleChoice

Your customer has requested a solution with HPE Alletra Storage MP B10000 that needs four SFP28 ports per node, for iSCSI connectivity.



Refer to DOM 505004:

Qty	Product #	Product Description
1	2UJ1SA	HPE Virtual Rack Service
1	5005AA	HPE Alletra Storage MP B19100 Base Configuration
1	5002AA	CO1 Factory Integrated
1	5000AA	HPE Alletra Storage ArrayO/S per TB 5-year LTU
1	5000DA	CO1 Factory Integrated
1	9S1813	CO1 HPE Configurator Defined Build Instruction Option
1	R1CT1A	HPE Alletra Storage MP 10000 2U Chassis
1	R1CT1A	CO1 Factory Integrated
2	5005AA	HPE Alletra Storage MP B19120 Controller Node
2	5002AA	CO1 Factory Integrated
2	R1CB1A	HPE Alletra Storage MP 10210QE 4-port Host Bus Adapter
2	R1CD0A	CO1 Factory Integrated
1	R0Q33A	HPE C13 - C14 250V 15Amp Blue 1.8m WW Power Cord (North America, Europe, Japan, Korea, Australia, New Zealand, Taiwan, China and India)
1	R0R33A	CO1 Factory Integrated
1	R2Z30A	HPE C13 - C14 250V 15Amp Gray 1.8m WW Power Cord
1	R0S30A	CO1 Factory Integrated
1	S0A46A	HPE Storage Data Encryption LTU
1	S0A49A	CO1 Factory Integrated
20	5002AA	HPE Alletra Storage MP B10000 per TB 5-year Software and Support SaaS
1	H6MA6A	HPE 3Y Tech Care Essential Service
1	H6MAA0000W	HPE Alletra 5kg MP Rate Config Setup
1	H6MAA0000Q	HPE Alletra Storage MP 2U Chassis Setup
2	H6MAA0000M	HPE Alletra 5kgMP BC Switchless Node Setup
2	H6MAA5000Q	HPE Alletra 5kg MP 10210QE 4-port HBA Setup
12	H6MAA2000Q	HPE Alletra 5kg MP B1910 NVMe SSD Setup
2	H6MAA6000Q	HPE Alletra 5T02 MP 60 TC COE HBA Support
1	H6B6AA	HPE Alletra Storage MP B10000 SVC
1	H4124A1	HPE Technical Interactive Startup SVC
1	H4132A1	V2V HPE Alletra Storage Array OS Startup SVC
1	H4134A1	V2V HPE Alletra 5T0 MPB10000 2NFS Setup SVC

A junior architect has prepared a BOM for the solution.

What needs to be changed in the BOM to meet the requirement?

Options:

- A- Remove the 52564A HBAs and add Q2P64B XCVRs to the BOM.
- B- Add two 52564A HBAs to the BOM.
- C- Add four Q2P64B XCVRs to the BOM.
- D- Replace the 52564A HBAs with R7C82A HBAs in the BOM.

Answer:

C

Explanation:

The correct answer is C. The customer requirement is explicit: four SFP28 iSCSI ports per controller node. An HPE Alletra Storage MP B10000 controller pair has two nodes, so the design requires eight usable SFP28 optical connections in total when all four 25GbE ports on each node are to be populated. The BOM already includes the Ethernet host bus adapters required for the iSCSI personality; the gap is the number of optical transceivers. HPE Alletra Storage MP B10000 supports host-port configuration by node, and HPE QuickSpecs list host port limits by controller model. HPE also identifies the 10/25GbE 4-port Host Bus Adapter as the correct Ethernet HBA class for Alletra Storage MP connectivity, while the Q2P64B item is the 25Gb SFP28 short-wave extended-temperature optical transceiver. Option A incorrectly removes the HBAs. Option B adds unnecessary HBAs instead of completing the optics population. Option D replaces adapters unnecessarily when the correct Ethernet HBA capacity is already present. The required correction is to add the missing four SFP28 transceivers so that all eight 25GbE ports can be used for iSCSI. Reference/topics: B10000 host-port sizing, iSCSI connectivity, SFP28 transceiver validation, BOM review.

Question 12

Question Type: MultipleChoice

Your customer has a networking environment based on Aruba CX 8325 switches configured with Quality of Service settings that give higher priority to iSCSI traffic between HPE ProLiant servers and HPE Alletra 6000 storage. They are planning to add HPE Synergy frames with HPE Virtual Connect modules that will connect to the same Aruba infrastructure.

The frames will be deployed in different VLANs. They need to update their environment to provide a lossless fabric for RoCE traffic that will traverse between the HPE Synergy frames and through multiple switches.

Your customer has two sets of Synergy frames in different logical enclosures connected with Aruba CX 8325 switches. They want to update their environment to provide lossless traffic for RoCE traffic.

What is the correct environmental or design change that will meet the customer's new requirement?

Options:

- A- Use RoCEv1 and configure DCBX on the switch facing the CNA in the servers.
- B- Server Profiles should include redundant connections to a RoCE network.
- C- Create DCB queues using server's operating system utility.
- D- Assign two different priority values to RoCE traffic.

Answer:

B

Explanation:

The correct answer is B. In HPE Synergy, network connectivity for compute modules is defined through HPE OneView server profiles. When RoCE is required, the server profile must explicitly include RoCE connections mapped to the appropriate RoCE network. HPE OneView guidance for adding a RoCE network states that server profiles should be created or edited with at least two RoCE connections mapped to the network. That directly supports the customer's requirement for resilient, lossless RDMA connectivity between Synergy frames and the Aruba data-center fabric. Option A is incorrect because RoCEv1 is Layer 2 only and is not appropriate for a design traversing multiple switches and VLAN/routed boundaries. Option C is incomplete because operating-system queue creation alone does not establish Synergy Virtual Connect fabric policy or redundant RoCE connectivity. Option D is technically wrong because RoCE traffic should be consistently classified and lossless behavior must be enforced end to end, not split across two different priority values. Reference/topics: HPE Synergy networking, HPE OneView server profiles, RoCE networks, DCB/lossless Ethernet, Aruba CX data-center fabric.



To Get Premium Files for HPE7-V01 Visit

<https://www.p2pexams.com/products/hpe7-v01>

For More Free Questions Visit

<https://www.p2pexams.com/hp/pdf/hpe7-v01>

20%
DISCOUNT

P2P
exams