Accelerate Your Certification with HPE7-V01 Practice Questions

Last chance to prepare smart! Get your hands on free HP Advanced HPE Edge-to-Cloud Solutions exam PDF questions. Study real HPE7-V01 dumps with verified answers and fast-track your certification success with PrepBolt HPE7-V01 exam pdf questions and answers.

Thank you for Downloading HPE7-V01 exam PDF Demo

https://prepbolt.com/HPE7-V01.html

Question 1

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 2

Question Type: MultipleChoice

You are designing a solution for a customer that will support a virtualization workload based on HPE Synergy. Your customer has the following requirements:

Dual-socket HPE Synergy compute nodes running ESXi 8.x with maximum memory supported

Each node should have access to shared storage from HPE Alletra Storage MP B10000

All virtual machines should be protected against host failure

All virtual machines with vGPU must be protected with Fault Tolerance

All compute nodes should boot from disks located in the D3940 Storage Modules protected with RAID 1

All frames should be managed as one logical enclosure to simplify management

All compute nodes should support two network ports with the maximum Ethernet connectivity speed

You are reviewing a proposal that you must verify. The proposal consists of:

3 x HPE Synergy 12000 frames

2 x HPE Synergy Composer2 modules

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

4 x HPE Synergy 50Gb Interconnect Link Modules

3 x HPE Synergy D3940 Storage Modules, each with 20 SSD drives

1152 core VMware vSphere Enterprise Plus subscriptions

36 x HPE Synergy 480 Gen11 compute nodes

Each SY480 Gen11 is configured with:

A single Platinum CPU with 32 cores

HPE Synergy 6820C 25/50Gb Converged Network Adapter

HPE Synergy 5830C 32Gb Fibre Channel Host Bus Adapter

HPE MR416i-o Gen11 12G SAS Mezzanine Controller

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

Options:

A- 2 HPE Synergy 50Gb Interconnect Link Modules should be removed from the proposal, as the total

number of interconnects required is six.
B- HPE MR416i-o Smart RAID controller can be removed because there are no external SAS disk enclosures included in the proposal.
C- Two HPE Virtual Connect SE 100Gb F32 Modules for Synergy can be removed from the proposal because all requirements are met without them.
D- To access HPE Alletra MP B10000 arrays, FC upgrade licenses for HPE Virtual Connect SE 100Gb F32 Modules for Synergy must be purchased.
E- vSphere Enterprise Plus subscriptions can be replaced with vSphere Standard subscriptions.

Answer:

C, D

Explanation:

The correct answers are C and D. In a three-frame HPE Synergy design using the primary/satellite fabric model, a redundant pair of HPE Virtual Connect SE 100Gb F32 modules can provide the intelligent fabric function while Interconnect Link Modules extend connectivity to satellite frames. Therefore, the proposal can be reduced by removing two unnecessary F32 modules if the remaining pair and link modules still satisfy the required redundant 25/50Gb connectivity. HPE describes the F32 architecture as consolidating data center network connections and scaling bandwidth across multiple Synergy frames using a primary/satellite design. Option D is also correct because Fibre Channel uplink functionality on the HPE Virtual Connect SE 100Gb F32 requires a separate Synergy 32Gb FC upgrade license. HPE OneView documentation states that the Synergy 32Gb FC Upgrade license is required for this interconnect to use Fibre Channel uplinks. Option B is false because the D3940 boot design still needs the SAS/RAID path. Option E is false because Fault Tolerance and advanced availability requirements make Standard an unsafe downgrade. Reference/topics: Synergy fabric validation, F32 FC licensing, D3940 boot architecture, VMware licensing review.

Question 3

Question Type: MultipleChoice

customer's requirements]

Your customer wants to implement a disaster recovery solution for their virtualization workload. They also want to implement a protection method that will protect the workload in case of a site failure.

Their existing environment includes the following configurations:

All ESXi systems are located in the same datacenter

Virtual machines are protected by an HA cluster

Some of the VMs are protected with Fault Tolerance

HPE Synergy 480 Gen11 is the compute platform

HPE Alletra Storage MP B10000 array for storage

Orchestrated failover for all VMs using VMware tools

The customer plans to replicate the same setup at the second site. The total number of virtual machines that must be protected is 200. Your customer also expects an RPO of 0.

Which component or configuration is a necessary part of a solution that will meet all customer requirements?

Options:

A- HPE Synergy frames from the secondary site added to the logical enclosure configured in the primary site.
B- Storage policy-based management (SPBM) must be configured for a group of the virtual machines.
C- Synchronous replication configured between sites to replicate the volumes used to build the datastores.
D- Asynchronous replication configured between sites used to replicate the volumes used to build the datastores.

Answer:

C

Explanation:

The correct answer is C. RPO 0 is the decisive requirement. A zero recovery point objective means the customer cannot lose acknowledged writes if a site fails. That requirement cannot be satisfied by asynchronous replication because asynchronous copy always introduces replication lag. The storage volumes backing the VMware datastores must be synchronously replicated between the two sites so that committed data exists at both locations before the write is considered protected. HPE Alletra Storage MP B10000 is designed for mission-critical block workloads and supports high-availability replication models, including synchronous Remote Copy / Peer Persistence designs for VMware environments. HPE documentation for Active Peer Persistence describes synchronous Remote Copy as the mechanism used for high-availability protection across arrays. Option A is incorrect because stretching or joining HPE Synergy frames across sites is not the mechanism that provides RPO 0. Option B can help with VMware policy-driven placement, but SPBM alone does not replicate data. Option D fails the RPO requirement. Reference/topics: DR architecture, RPO 0, synchronous replication, B10000 Remote Copy, VMware datastore protection.

Question 4

Question Type: MultipleChoice

You presented a solution to your customer that will replace their legacy HPE BladeSystem.

The new solution is based on the HPE Synergy platform and consists of the following components:

5 HPE Synergy 12000 frames

2 HPE Virtual Connect SE 100Gb F32 Modules for Synergy

8 HPE Synergy 25/50Gb Interconnect Link Modules

60 HPE Synergy 480 Gen11

Which statement about the business and financial benefit of the new solution for the customer is true?

Options:

A- HPE Synergy 480 Gen11 compute modules have an extension that will allow the customer to install 4 GPUs in half-height format compute nodes.
B- HPE Virtual Connect SE 100Gb F32 Module for Synergy can support FC traffic once an upgrade license for each module is added in HPE OneView.
C- Compared to the legacy HPE BladeSystem, the HPE Synergy platform supports more expansion slots in the form of mezzanine cards.
D- Compared to the legacy HPE BladeSystem, the HPE Synergy platform supports more interconnect modules, which provides better flexibility and expansion possibilities.

Answer:

D

Explanation:

The correct answer is D. The business value of the proposed HPE Synergy design is the more flexible, scalable composable-fabric architecture compared with traditional blade designs. HPE Synergy uses a primary/satellite networking architecture with Virtual Connect and Interconnect Link Modules, allowing network connectivity to scale across multiple frames while reducing upstream cabling, top-of-rack dependency, and aggregation-port consumption. HPE describes the Virtual Connect SE 100Gb F32 design as using a disaggregated, rack-scale primary/satellite architecture to consolidate data center network connections, reduce hardware, and scale network bandwidth across multiple HPE Synergy frames. That is the financial point: the customer gains higher fabric flexibility and expansion capability without replicating a full set of expensive intelligent fabric modules in every frame. Option A is not the relevant benefit and overstates the GPU expansion concept. Option B is technically true about FC capability, but it does not describe the primary business and financial value of the proposed multi-frame architecture. Option C is a narrower compute-module feature and is not the strongest financial argument. Reference/topics: HPE Synergy business value, Virtual Connect primary/satellite fabric,

network consolidation, financial-impact presentation.

Question 5

Question Type: MultipleChoice

A customer requested a quote for a 64Gb Fibre Channel system between two data centers with a distance of 95 KM / 60 MI between the locations, with a requirement for 14 F-ports on each switch.

Refer to exhibit

Which statement is true about the given configuration?

Options:

A- A Short Wave transceiver must be removed to install the Long Wave transceiver.
B- A Fibre Channel license needs to be modified for Extended Fabric to support the required distance.
C- A SFP56 or SFP56-DD port can be used for the Long Wave transceiver when all SFPs are inserted.
D- A port range of 39 to 47 can be used for the LW transceiver if ports 0 to 23 are used for SW.

Answer:

B

Explanation:

The correct answer is B. The controlling requirement is the 95 km / 60 mile inter-data-center Fibre Channel distance. Standard Fibre Channel fabrics have native distance limits, and long-distance SAN designs require the correct licensing and buffer-credit behavior to maintain stable inter-switch link operation. HPE SAN design guidance states that B-series switches provide extended fabric support up to 10 km by default, but for distances greater than 10 km an Extended Fabrics license is required. Therefore, the configuration must include or modify the Fibre Channel software licensing to support the required long-distance fabric behavior. Option A is not the architectural issue; optics can be changed, but the long-distance fabric still needs the correct licensing and configuration. Option C is misleading because port form factor compatibility alone does not satisfy 95 km operation. Option D is too specific to port numbering and does not solve the distance requirement. The quote must be validated for both physical optics and the SAN feature license that enables extended-distance behavior. Reference/topics: HPE B-series Fibre Channel switching, Extended Fabric, inter-data-center SAN design, proposal validation.

Question 6

Question Type: MultipleChoice

Your customer is expecting to see an increase in the amount of data stored on their HPE Alletra Storage MP B10000.

They are asking to create a plan to support a raw capacity of 450 TB all flash.

Refer to exhibit

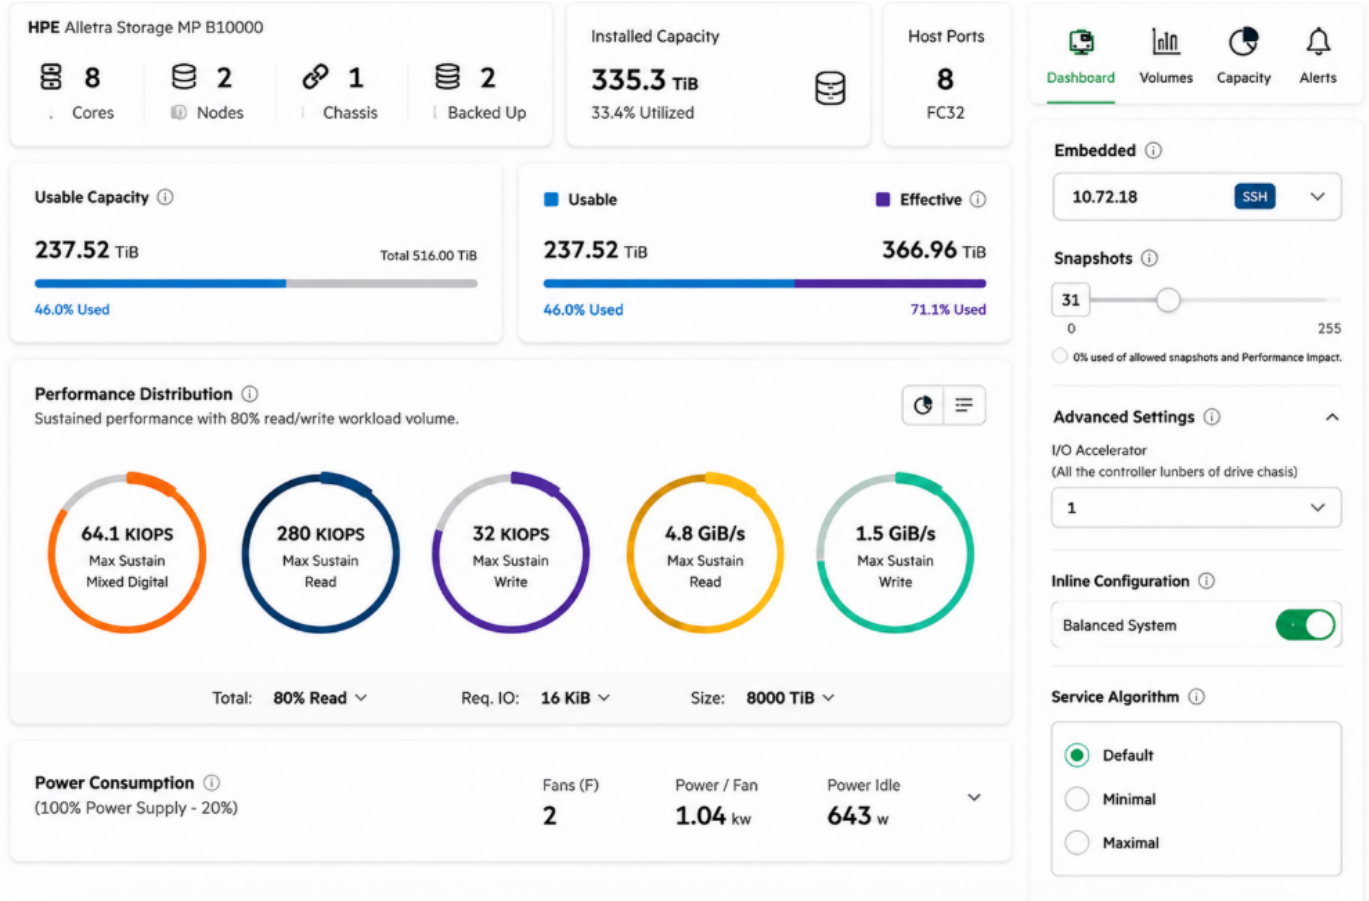

What change would you recommend to the configuration to achieve this?

Options:

A- Add 100Gb HBAs to B10120 nodes and additional drive shelf.
B- Replace the QLC drives with higher capacity TLC drives.
C- Add more QLC drives to the system and change the sparing to minimal.
D- Replace the nodes with B10120 model and add additional drive shelf.

Answer:

C

Explanation:

The correct answer is C. The exhibit shows an HPE Alletra Storage MP B10000 configuration based on QLC flash, with raw capacity below the customer's 450 TB target. The correct capacity adjustment is to increase the number of QLC drives in the system while selecting the sparing policy that maximizes available usable capacity for the intended design. Replacing QLC with TLC is not the right direction for a high-capacity all-flash requirement because QLC media is positioned for dense, cost-efficient capacity expansion. Adding host bus adapters does not increase raw storage capacity, so option A is unrelated to the requirement. Replacing controller nodes is also unnecessary because the problem is media capacity, not controller performance. HPE Alletra Storage MP B10000 QuickSpecs identify scale-out capacity characteristics, including maximum raw capacity per shelf, confirming that the platform is intended to scale capacity through shelves and drive population rather than by changing host-connectivity cards. The correct plan is to add QLC capacity to reach the raw-capacity target and adjust sparing appropriately. Reference/topics: B10000 capacity planning, QLC all-flash sizing, sparing policy, GreenLake proposal validation.

Thank You for trying HPE7-V01 PDF Demo

To try our HPE7-V01 practice exam software visit link below

https://prepbolt.com/HPE7-V01.html

Start Your HPE7-V01 Preparation

Use Coupon “SAVE50” for extra 50% discount on the purchase of Practice Test Software. Test your HPE7-V01 preparation with actual exam questions.