

DUMPSBOSS.

Advanced HPE Storage Integrator Solutions Written Exam

HP HPE7-J02

Version Demo

Total Demo Questions: 7

Total Premium Questions: 73

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Describe, differentiate, and apply industry standard storage technologies and architectures	15
Topic 2, Gather and analyze customer information	10
Topic 3, Recommend and position HPE storage solutions	24
Topic 4, Design HPE storage solutions	21
Topic 5, Validate HPE storage solution design	3
Total	73

QUESTION NO: 1

Your customer wants to use their HPE Alletra Storage MP B10000 array to store persistent data for Kubernetes-based applications. After deploying the CSI driver using Helm and creating the secret with the command `kubectl create -f hpe-backed.yaml`, what is the next required step to enable the containerized applications to consume persistent volumes on the Alletra MP array?

- A. Update the Helm repository by using `helm repo update` to recognize the CSI driver
- B. Create a PersistentVolumeClaim by using `kubectl create -f my-pvc.yaml`
- C. Create a PersistentVolume by using `kubectl create -f pv.yaml`
- D. Define a StorageClass by running `kubectl create -f storageclass.yaml`

ANSWER: D

Explanation:

Define a StorageClass by running `kubectl create -f storageclass.yaml`. In Kubernetes, a StorageClass supplies the provisioner and the storage-specific parameters that Kubernetes uses for dynamic volume provisioning. For an HPE Alletra Storage MP B10000 deployment, the StorageClass associates a PersistentVolumeClaim request with the installed HPE CSI driver and identifies the desired array-side provisioning behavior, such as the backend configuration, service level, or other driver-supported volume attributes. The previously created secret provides the CSI driver with the credentials and connection information it needs to communicate securely with the storage system; the StorageClass then makes that configured driver available as a consumable storage offering in the cluster. Once the StorageClass exists, application teams can request capacity through PersistentVolumeClaims. Kubernetes uses the selected StorageClass to invoke the CSI driver, which creates and presents the backing persistent volume on the Alletra MP array and binds it to the claim. This dynamic-provisioning workflow allows applications to request storage declaratively without requiring an administrator to pre-create a separate volume for every workload. Kubernetes documents that StorageClasses define the classes of storage offered to users and are used by dynamic provisioners when claims request them. See the [Kubernetes StorageClass documentation](#) and HPE's [HPE container-platform resources](#).

QUESTION NO: 2

You are troubleshooting a storage environment using HPE Alletra Storage MP B10000 in a Peer Persistence configuration. A customer reports high latency when accessing data from applications.

When reviewing the MPIO path status on the host, what should be the path status for connections to the storage array located in the remote data center?

- A. Active Optimized
- B. Passive
- C. Active Non-optimized
- D. Standby

ANSWER: C

Explanation:

Active Non-optimized is the expected ALUA/MPIO state for host paths leading to the array in the remote data center in a normal HPE Peer Persistence configuration. Peer Persistence presents a common volume to hosts through both participating storage systems, while maintaining a preferred, or optimized, access path to the array that currently owns and services the volume. The paths through the remote array remain available and usable, but I/O on them requires the additional inter-array transport to reach the volume's owning system. Consequently, those paths are advertised as active but non-optimized rather than unavailable. Host multipathing software can therefore recognize the locality preference and preferentially direct I/O through the optimized local paths, while retaining remote connectivity for transparent failover and continuity if the preferred site or paths become unavailable. During latency troubleshooting, seeing Active Non-optimized for remote-site paths is therefore a normal topology indication; it confirms that ALUA path-state information is being provided to

the host for path selection. Refer to HPE's [Alletra Storage MP B10000 documentation portal](#) for the applicable Peer Persistence and host-configuration documentation, and the [SCSI ALUA specification overview](#) for the active/optimized and active/non-optimized access-state model.

QUESTION NO: 3

Refer to the exhibit.

The array is experiencing frequent cache misses for read operations.

Which action plan would you suggest to correct the issue?

- A. Upgrade the setup with additional expansion shelf.
- B. Upgrade the controllers to Alletra 5050 models.
- C. Increase the pinned cache size.
- D. Increase the flash to disk ratio.

ANSWER: D

Explanation:

Increase the flash to disk ratio. On an HPE Nimble Storage hybrid array, SSD capacity is used as an adaptive read cache for the active data set, while HDDs provide the capacity tier. A read-cache miss means requested data is not currently resident in that flash cache and must be retrieved from the disk tier, increasing response time for read-intensive workloads. Adding sufficient flash capacity relative to the disk capacity and active working set expands the available adaptive cache and raises the likelihood that frequently accessed blocks remain on SSDs. This directly addresses repeated read-cache misses and reduces the performance impact of backend disk reads. The sizing decision should be based on workload behavior observed through HPE InfoSight, particularly cache utilization, miss rates, and the active data-set characteristics, rather than capacity alone. HPE describes its Nimble Storage architecture as combining flash with capacity media and predictive analytics to optimize workload performance; review the HPE Nimble Storage information at [HPE Nimble Storage](#) and use the documentation and guidance available through the [HPE Support Center](#) when planning the flash expansion.

QUESTION NO: 4

Your customer is a software development company using Kubernetes, Canonical, and Tanzu. They are expanding their Dev/Ops container environment and are considering **HPE Private Cloud Business Edition (PCBE)**.

What is a unique benefit that HPE PCBE would bring to this customer's environment?

- A. Full stack intelligence
- B. Multi-tenant VASA
- C. Infinite scale to the customer at no additional cost
- D. Container level protection

ANSWER: A

Explanation:

HPE Private Cloud Business Edition provides **Full stack intelligence** by combining a cloud-operating experience with integrated management of the private-cloud infrastructure stack. For a DevOps organization running Kubernetes-based workloads and platforms such as Canonical and Tanzu, this gives IT teams centralized visibility and operational control over the compute, storage, virtualization, and cloud-management components that support those workloads. Rather than requiring separate infrastructure-management workflows for each layer, the platform is designed to simplify lifecycle operations, including deployment, monitoring, capacity management, and upgrades, through an HPE GreenLake cloud experience.

This is especially valuable as the customer expands its container environment: development teams can consume resources with greater agility while infrastructure teams retain governance, performance insight, and a consistent operational model. HPE positions Private Cloud Business Edition as a turnkey, self-managed private-cloud solution that delivers a cloud experience while keeping workloads and data under customer control. Its value is therefore the integrated, full-stack operational intelligence and automation supporting modern application environments. See the [HPE private cloud solutions overview](#) and the [HPE GreenLake for Private Cloud Business Edition page](#).

QUESTION NO: 5

The storage solution based on the exhibit is deployed at a customer site.



How can the sequential read performance values be enhanced for this configuration?

- A. By adding another expansion shelf
- B. By increasing the amount of 10/25 Gb NICs
- C. By replacing 10/25 Gb with a 32 Gb FC HBA
- D. By adding SCM to the solution

ANSWER: B

Explanation:

By increasing the amount of 10/25 Gb NICs is correct because sustained sequential-read throughput is strongly influenced by the aggregate bandwidth available across the storage system's front-end host connections. Large sequential reads transfer substantial contiguous blocks of data, so adding appropriately configured 10/25 Gb Ethernet NIC ports creates additional paths through which data can be delivered to the connected hosts. With suitable host-side adapters, switch capacity, multipathing, and workload distribution, those additional links can raise the available aggregate read bandwidth and improve the sequential-read value shown for the configuration.

This approach aligns the improvement with the relevant potential bottleneck: host-facing connectivity. The ports should be configured using supported redundancy and multipath practices, and the Ethernet switching infrastructure must provide sufficient nonblocking bandwidth so that it does not constrain the new links. HPE's storage guidance emphasizes designing host connectivity with adequate ports, paths, and fabric bandwidth for the expected workload, while HPE's Ethernet guidance describes 10/25GbE as a high-bandwidth infrastructure option. See the [HPE Primera 600 Storage QuickSpecs](#) and [HPE Ethernet switch portfolio information](#).

QUESTION NO: 6

A customer is replacing a legacy deduplication appliance with HPE StoreOnce. The backup architect provides protected-data volumes, daily change rates, retention periods, expected deduplication assumptions, replication requirements, and backup and restore throughput targets.

Which HPE tool should be used to translate these requirements into an HPE StoreOnce recommendation?

- A. HPE CloudPhysics
- B. HPE NinjaSTARS
- C. HPE Ninja Online
- D. One Config Advanced

ANSWER: B

Explanation:

HPE NinjaSTARS is the appropriate tool because it is designed to size and configure StoreOnce solutions from backup-workload requirements. It accounts for the retention design, capacity assumptions, expected data reduction, performance targets, and replication considerations that affect an appropriate StoreOnce configuration.

CloudPhysics assesses infrastructure and virtualized workload behavior for modernization planning. Ninja Online is a broader HPE configuration and sizing environment but is not the StoreOnce-focused right-sizing tool. One Config Advanced supports quoting and configuration processes after the required solution sizing has been established.

QUESTION NO: 7

A customer plans to replace a VMware cluster and its third-party storage with HPE infrastructure. CPU, memory, capacity, and I/O utilization vary substantially between month-end processing and normal business days. The customer wants a sizing recommendation that reflects both normal and peak workload behavior.

Which assessment approach should you recommend?

- A. Deploy the HPE CloudPhysics Observer and collect data over a representative period
- B. Export a one-time vCenter inventory and size from configured VM resources
- C. Use HPE InfoSight to collect telemetry from the third-party storage array
- D. Use a product configurator without collecting workload metrics

ANSWER: A

Explanation:

Deploy the HPE CloudPhysics Observer and collect utilization data over a representative period. The Observer obtains inventory, configuration, capacity, and performance information from the vCenter-managed environment. Historical data allows the assessment to identify demand patterns, peak periods, and rightsizing opportunities rather than basing a design on a one-time inventory.

A static vCenter export can describe the current configuration, but it cannot reliably characterize workload peaks or utilization trends. InfoSight data is valuable for supported HPE infrastructure but is not the appropriate collection method for a competitive assessment of an existing third-party environment. A general configuration tool does not collect the required workload metrics.