

DUMPSBOSS.

Advanced HPE Storage Architect Written Exam

HP HPE7-J01

Version Demo

Total Demo Questions: 8

Total Premium Questions: 86

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Gather and analyze customer storage requirements	5
Topic 2, Design HPE storage solutions	23
Topic 3, Recommend HPE storage solutions	26
Topic 4, Implement HPE storage solutions	26
Topic 5, Troubleshoot HPE storage solutions	5
Topic 6, Mix Questions	1
Total	86

QUESTION NO: 1

An administrator is creating Virtual Protection Groups (VPGs) in Zerto to replicate information locally and to a remote disaster site. What is the maximum number of VPGs with which a VM can be associated?

- A. One
- B. Two
- C. Three
- D. Four

ANSWER: C

Explanation:

Three is correct. Zerto supports one-to-many protection, enabling the same protected virtual machine to participate in as many as three Virtual Protection Groups (VPGs). Each VPG is an independently configured replication and recovery unit: it has its own target site, journal, recovery settings, and service-level objectives. This permits a protection design in which the same VM is replicated for distinct operational purposes, such as local recovery, remote disaster recovery, and an additional off-site or long-retention copy.

This capability is useful when a business needs separate recovery domains or locations without treating the VM as a single, inflexible replication relationship. The architect must still size source and target resources appropriately, because every configured protection relationship produces replication traffic and requires journal and target-storage capacity. VPG membership also continues to preserve the write-order consistency required for the set of workloads grouped within each individual VPG. Zerto documentation describes VPGs as the core entity used to define protected workloads and recovery configuration, and its product documentation covers multi-site and multiple-protection design considerations. See the [Zerto Documentation portal](#) and HPE's [storage solutions documentation and resources](#) for current product guidance.

QUESTION NO: 2

A company has a pair of Alletra 9000s, managed via the HPE GreenLake Data Services Cloud Console (DSCC). An administrator installed Kubernetes locally but requires persistent storage using the Alletra 9000s. After installing the helm repo for the HPE CSI Driver for Kubernetes, what is the next step the administrator should perform to use the Alletra 9000s for persistent storage?

- A. Create a Kubernetes namespace for the HPE CSI Driver.
- B. Create a secret to allow the HPE CSI Driver to communicate with the Alletra 9000s.
- C. Create a storage class that references the Alletra 9000s on the Kubernetes conductor.
- D. Add the Kubernetes conductor credentials to the Alletra 9000s in the HPE GreenLake DSCC.

ANSWER: B

Explanation:

Create a secret to allow the HPE CSI Driver to communicate with the Alletra 9000s. The HPE CSI Driver needs backend connection and authentication information before it can issue storage operations to an Alletra 9000 array. This information is supplied to Kubernetes as a Secret, typically through a backend configuration file that includes the array management address and authorized array credentials. Keeping these values in a Kubernetes Secret avoids placing sensitive credentials directly in driver or workload manifests.

Once the Secret is available in the driver's namespace, the installed CSI components can use it when communicating with the array management interface. The driver can then create, publish, expand, and delete persistent volumes on behalf of Kubernetes. A subsequent StorageClass can select the appropriate provisioning behavior and reference the backend configuration needed for the intended service level. This sequence establishes the required trusted connection to the Alletra storage system before application persistent-volume claims are dynamically provisioned. HPE's Storage Container Orchestrator documentation provides the CSI deployment and backend-configuration workflow, while Kubernetes documents

QUESTION NO: 3

An architect is designing a two-site VMware disaster recovery solution with Zerto. Which two components are required to establish the fundamental replication architecture between the production and disaster recovery sites? (Choose two.)

- A. A Zerto Virtual Manager is installed at each site.
- B. A Virtual Replication Appliance is installed on each participating hypervisor host.
- C. A single Zerto Virtual Manager at the production site manages both sites.
- D. Replica disks must be retained only at the production site.
- E. Each protected VM requires an HPE StoreOnce Catalyst Store.

ANSWER: A B

Explanation:

A Zerto Virtual Manager is required at each site to manage the local Zerto environment and coordinate the paired-site relationship. Virtual Replication Appliances are deployed on the participating hypervisor hosts to intercept writes and move replication data. Replica disks and journals are maintained at the recovery site for a production-to-DR protection design; they are not dependent on a single ZVM at the production site.

QUESTION NO: 4

A customer has three HPE Alletra 9000 arrays at separate production sites and plans to use an N-to-1 Remote Copy topology to one disaster recovery site. Which two design considerations are required for the disaster recovery array? (Choose two.)

- A. Size the target for aggregate replicated capacity and failover performance
- B. Configure separate Remote Copy relationships and target mappings for each source
- C. Deploy one dedicated disaster recovery array for each production array
- D. Use synchronous replication because fan-in requires zero data loss
- E. Eliminate replication bandwidth planning because the target is shared

ANSWER: A B

Explanation:

In an N-to-1, or fan-in, topology, one target array receives replicated volumes from multiple source arrays. The target must therefore be sized for the aggregate replicated capacity and for the performance required if protected applications are started at the disaster recovery site. Each source array also requires its own correctly configured Remote Copy relationship and target-volume mapping on the common target array.

The topology does not require one physical target array per source array. The use of asynchronous or synchronous replication is determined by the workload recovery requirements and available latency; fan-in itself does not require synchronous replication. Source-to-target replication bandwidth remains necessary because all protected changes must reach the disaster recovery array.

QUESTION NO: 5 - (SIMULATION)

Order the steps for a write data path and a successful write IO in HPE GreenLake for File Storage using NAS.

ANSWER: See the explanation for the answer

Explanation:

Correct order for a successful NAS write I/O:

1. **Data is sharded randomly across multiple SCM drives** to increase throughput and reduce contention.
2. **Data is written to two different SCM drives**, providing mirrored persistence and protection if an SCM drive fails.
3. **Metadata is updated in the internal data structure (tree)** to maintain a consistent filesystem view.

After the write is persistent on the protected SCM tier and the metadata is updated, the write can be acknowledged to the NAS client. Data is subsequently destaged asynchronously to the capacity tier.

QUESTION NO: 6

A customer needs to replace its HPE Alletra 6000 storage array with either an HPE Alletra 9000 managed with GreenLake or an IBM FlashSystem solution. Which two advantages should an HPE Partner share with the customer as to why the customer should opt for the HPE storage solution? (Choose two.)

- A. Fractional scaling
- B. Full stack automation with hybrid cloud
- C. Ransomware detection
- D. AIOps for infrastructure
- E. Granular disaster recovery

ANSWER: B D

Explanation:

“Full stack automation with hybrid cloud” is a key HPE advantage because HPE Alletra 9000 is designed for a cloud operational experience through the HPE GreenLake platform. GreenLake centralizes management and data-services operations, allowing teams to use a consistent, cloud-style control plane for on-premises infrastructure while integrating storage operations into broader hybrid-cloud workflows. This model helps reduce operational overhead through policy-based management, automation, and API-driven workflows rather than treating the array as an isolated infrastructure component.

“AIOps for infrastructure” is also correct because HPE InfoSight provides predictive analytics and cross-stack operational intelligence for HPE storage. InfoSight analyzes telemetry from the storage environment and correlates it with infrastructure and workload behavior to identify performance risks, capacity trends, and operational issues. For an organization moving from Alletra 6000 to Alletra 9000, this delivers continuity of an intelligence-led operations model while adding the enterprise availability and performance capabilities of the Alletra 9000 platform. Together, GreenLake cloud operations and InfoSight AIOps position HPE around simplified lifecycle management and proactive infrastructure operations. See [HPE Alletra Storage 9000](#) and [HPE InfoSight](#).

QUESTION NO: 7

A customer has a mix of applications for VMware VMs, HPE containers, and bare metal solutions. The customer is an early adopter of containers and is already using HPE Ezmeral Runtime Enterprise. They have the following criteria:

The customer runs applications on VMware VMs.

The customer wants to run multiple workloads on bare metal, VMs, and containers.

The customer wants a fully managed hybrid multi-cloud environment.

The customer wants to integrate into DevOps toolchains for immediate productivity.

Which solution will best fit the customer's needs?

- A. HPE GreenLake for VCF
- B. HPE GreenLake for Microsoft Azure Stack HCI
- C. HPE Private Cloud Business Edition with HPE VME
- D. HPE Private Cloud Enterprise

ANSWER: D

Explanation:

HPE Private Cloud Enterprise is the best fit because it is designed to provide a managed, cloud-operated private-cloud environment for heterogeneous enterprise workloads. The customer needs one operational model across VMware virtual machines, containerized applications, and bare-metal resources rather than an offering focused solely on a particular virtualization stack. Private Cloud Enterprise supports this multigenerational workload approach while giving IT and application teams self-service consumption, policy-based governance, lifecycle operations, and a public-cloud-like experience from the GreenLake platform.

Its managed-service model directly addresses the requirement for a fully managed hybrid multi-cloud environment: HPE operates and maintains the underlying platform, allowing the customer to focus on application delivery rather than infrastructure maintenance. It also aligns well with an organization already adopting HPE Ezmeral Runtime Enterprise, since Kubernetes-based container operations can be incorporated alongside virtualized and physical workloads. DevOps productivity is supported through service catalogs, automation, infrastructure-as-code capabilities, and integrations with established development and operations toolchains. These capabilities let teams provision and manage resources consistently while retaining appropriate enterprise controls. HPE describes Private Cloud Enterprise as a fully managed private-cloud service, and its cloud-management capabilities are intended to automate and govern workloads across hybrid environments. See [HPE Private Cloud Enterprise](#) and [HPE Morpheus Enterprise Software](#).

QUESTION NO: 8 - (SIMULATION)

Match the fibre channel topology to its use case. Each answer will be used once.

ANSWER: See the explanation for the answer

Explanation:

Correct topology-to-use-case matches:

Meshed fabric — many-to-many connectivity and high-performance requirements.

Cascaded fabric — localized data access, with servers and storage commonly attached to the same switch.

Core-edge fabric — a mix of local and distributed data access.

Ring fabric — diverse geographic conditions and locations.

A meshed design provides multiple paths and high performance; cascaded designs suit localized access; core-edge designs scale for mixed access patterns; and ring designs provide resilient connectivity across geographically separated sites.