



HP HPE7-J01

Advanced HPE Storage Architect Written Exam

Thank You for Downloading HPE7-J01 Updated Exam Questions

<https://certs4sale.com/hp/hpe7-j01-pdf-exam-dumps>

Version: 4.0

Question: 1

A storage administrator wants to set up NAS replication between two HPE StoreOnce appliances. A corresponding NAS library was created between the two systems, primary and secondary, where the secondary will receive the replicated information. When the administrator tries to create the NAS mapping on the primary StoreOnce solution, the administrator is presented with an error stating they do not have permission. What should the administrator do to solve this issue?

- A. On the primary StoreOnce appliance, generate an access token and configure this token on the secondary StoreOnce appliance.
- B. On the secondary StoreOnce appliance, set the Replication Permissions to allow Enabled Public Access.
- C. On the secondary StoreOnce appliance, define the primary StoreOnce appliance, along with the respective username and password.
- D. On the primary StoreOnce appliance, define the secondary StoreOnce appliance, along with the respective username and password.

Answer: C

Explanation:

In the context of HPE StoreOnce Catalyst and NAS replication, security is governed by a bi-directional trust or permission-based handshake. When configuring replication between two StoreOnce appliances, the "Target" (Secondary) system acts as the gatekeeper. The error described occurs because the primary system is attempting to push data or create a mapping to a destination that has not authorized it.

According to the HPE StoreOnce documentation regarding Replication Permissions, the secondary appliance must explicitly grant permission to the source appliance before any mapping or data transfer can occur. This is a security measure designed to prevent unauthorized data ingestion or "rogue" replication tasks from consuming storage resources on the target system. To resolve the permission error, the administrator must log into the Secondary StoreOnce appliance (the target) and navigate to the replication settings. There, they must add the Primary StoreOnce appliance as an authorized "source" by specifying its network address (FQDN or IP) and providing the necessary

credentials (username and password) that the primary system will use to authenticate.

Unlike simpler protocols where a "Public Access" setting (Option B) might exist, HPE StoreOnce requires a defined relationship for NAS and Catalyst replication to ensure data integrity and multi-tenancy security. Option A refers to token-based authentication, which is more common in modern cloud-native Alletra environments via DSCC, but not the standard for legacy StoreOnce NAS replication. Option D is incorrect because the permission must be granted at the receiving end, not the sending end. Once the secondary system has the primary's details stored in its Replication Permissions list, the primary appliance will be able to successfully "discover" the target libraries and establish the mapping without further permission errors.

Question: 2

A customer wants to implement HPE Cloud Bank Storage with the detach option LTU feature. Which statement is correct regarding the implementation of this feature?

- A. HPE Cloud Bank Store must be connected using the connect (read-write) option.
- B. A detached Cloud Bank datastore can only be reconnected to the HPE StoreOnce system from which it was detached.
- C. HPE Services is always required when reconnecting a detached datastore.
- D. Support for reading and writing to the detached datastore is supported with the detach license installed.

Answer: A

Explanation:

HPE Cloud Bank Storage is an extension of HPE StoreOnce Catalyst that enables the movement of deduplicated data to public, private, or hybrid cloud object storage. The Cloud Bank Detach feature is a critical lifecycle management capability designed for long-term retention and disaster recovery scenarios.

According to the HPE StoreOnce User Guide, the "Detach" operation is a specific administrative action that essentially "unplugs" the Catalyst store from the local StoreOnce appliance while leaving the data intact in the cloud bucket (e.g., AWS S3 or Azure Blob). For an administrator to initiate the detach process, the Cloud Bank store must currently be in a Read-Write (RW) state on the StoreOnce system. If a store is currently connected as Read-Only (often the case after a disaster recovery sync), it cannot be detached until it is promoted or was originally connected in a Read-Write capacity.

Once the detach operation is executed using the required Detach Capacity LTU (License to Use), the store enters a "Detached" state. In this state, the data in the cloud becomes immutable and the store is removed from the local StoreOnce system's active management. It is important to note that once detached, the store can only be reconnected to a StoreOnce system (either the original or a new one for DR) in a Read-Only state for recovery purposes. Statement D is incorrect because you cannot write to a detached store. Statement B is incorrect because one of the primary value propositions of Cloud Bank is portability—allowing you to connect a detached store to a completely different StoreOnce appliance in a different region for recovery. Finally, while HPE Services are available for complex DR planning, they are not a technical requirement for the software-defined reconnection of a detached store.

Question: 3

A customer is interested in a backup repository solution with long-term data retention. The customer has the following requirements:

Needs to leverage secondary storage for development operations and development testing

Fast granular restore and instant recovery features

Cost-effective, yet scalable solution that provides built-in replication features

What is the best solution for this customer?

- A. HPE Alletra 5000s and Veeam
- B. HPE dHCI and Cohesity
- C. HPE Alletra 5000s and Scalify RING
- D. HPE Alletra 4000s and Commvault

Answer: A

Explanation:

The requirements provided point toward a "Secondary Storage" use case where the data must be more than just a "cold" backup; it needs to be "active" for DevOps and testing. The HPE Alletra 5000 (the successor to the HPE Nimble Storage Adaptive Flash arrays) is specifically engineered for this hybrid role.

Architecturally, the Alletra 5000 utilizes the CASL (Content Aware Storage Architecture) file system. This allows it to perform high-speed inline deduplication and compression, making it a cost-effective

repository for long-term retention. Crucially for the customer's DevOps requirement, Alletra 5000 supports Zero-Copy Clones. This means the storage administrator can instantly create multiple copies of production datasets for development and testing without consuming additional storage space or impacting the performance of the primary backup repository.

When paired with Veeam Backup & Replication, the solution meets the "fast granular restore" and "instant recovery" requirements perfectly. Veeam's vPower technology enables Instant VM Recovery, which allows a virtual machine to be started directly from the compressed and deduplicated backup file on the Alletra 5000. Because the Alletra 5000 includes a flash tier for metadata and frequently accessed data, it provides the necessary IOPS to run these recovered VMs or DevTest workloads with near-production performance.

In contrast, while Cohesity (Option B) is a strong secondary platform, HPE dHCI is a primary infrastructure solution and not just a backup repository. Scality RING (Option C) is an object storage solution geared toward massive scale and petabyte-level archives, but it lacks the performance characteristics for "instant recovery" and seamless DevOps cloning found in the Alletra 5000. HPE Alletra 4000 (Option D) is a high-density data server (formerly Apollo) which provides the raw hardware but lacks the integrated CASL-based intelligence and "Better Together" orchestration that the Alletra 5000/Veeam partnership offers for this specific customer profile.

Question: 4

Which HPE system can be integrated into a factory-built HPE Qumulo solution for a customer?

- A. HPE Apollo 4200
- B. HPE Alletra 5000
- C. HPE ProLiant 360
- D. HPE Apollo 4500

Answer: A

Explanation:

The HPE Solutions for Qumulo are a result of a strategic partnership designed to provide a high-performance, scale-out NAS (Network Attached Storage) platform for unstructured data. According to the HPE Solutions with Qumulo Reference Architecture, the primary hardware platform utilized for these factory-built, integrated solutions is the HPE Apollo 4000 series, specifically the HPE Apollo 4200.


The Apollo 4200 is chosen for this role because it is a density-optimized, storage-centric server that provides an ideal balance of compute and massive internal storage capacity within a standard 2U rack footprint. Architecturally, the Apollo 4200 supports an "SSD-first" hybrid configuration or an all-flash configuration, which aligns perfectly with Qumulo's file system requirements. Qumulo's software uses the SSDs for a high-speed metadata layer and write-cache, while utilizing high-capacity HDDs for the data plane, ensuring that even with billions of files, the system maintains near-flash performance.

While the HPE ProLiant DL325 is also used for specific all-NVMe nodes in the Qumulo portfolio, the Apollo 4200 remains the foundational building block for the hybrid and archive nodes that comprise the bulk of enterprise deployments. The HPE Apollo 4500 (Option D) is a 4U system that, while part of the Apollo family, is not the standard integrated platform for the mainstream Qumulo joint offering. The HPE Alletra 5000 (Option B) is a block-storage-focused platform derived from the Nimble lineage, and the ProLiant DL360 (Option C) is a general-purpose 1U compute server that lacks the internal drive density required for a high-capacity scale-out file storage solution. By selecting the Apollo 4200, customers benefit from a pre-validated, factory-integrated solution that simplifies the deployment of massive file lakes for workloads like video surveillance, medical imaging, and big data analytics.

Question: 5

HOTSPOT

An administrator wants to build a scale-out solution using two HPE Alletra 6000 storage arrays. Only the first array has been initialized. Use your cursor to place a + where the administrator can add the second uninitialized array to the first storage array.



Answer:

Explanation:

Hardware

To expand a single-array configuration into a multi-array scale-out cluster (known as an HPE Nimble/Alletra 6000 Group), the administrator must navigate to the physical resource management section of the interface. In the HPE Alletra 6000 and legacy Nimble Storage UI, this is located under the Hardware tab in the primary navigation menu.

Once the administrator clicks on Hardware, the system displays the "Arrays" view. If a second, uninitialized Alletra 6000 array is physically connected to the same network and has been discovered by the management software (via the Discovery protocol), it will appear in a list of "Discovered Arrays" or within the main hardware pane with an option to be added to the current group. By selecting the uninitialized array and clicking the add/plus icon within that specific Hardware sub-menu, the administrator triggers the group expansion process. This process non-disruptively integrates the new array's controllers and capacity into the existing management plane, allowing for seamless data migration and performance scaling across the member arrays.

The exhibit provided (image_64b359.png) currently shows the Data Protection sub-tab under the Manage menu. This area is dedicated to managing volume collections, snapshots, and replication partners rather than hardware membership. To perform the hardware-level task of adding a new physical node to the cluster, the administrator must first move away from the "Manage" view and select Hardware. This architectural separation ensures that logical data management tasks are distinct from physical infrastructure scaling. Once the second array is added, the "Manage" view will then reflect the combined resources of both arrays, allowing for group-wide volumes that can span the performance and capacity of the entire scale-out solution.

THANK YOU FOR DOWNLOADING HPE7-J01 UPDATED EXAM QUESTIONS

Note: Thanks For Trying The Demo Of Our HPE7-J01 Exam Product

Visit Our Site to Purchase the Full Set of Actual HPE7-J01 Exam Questions With Answers.

Money Back Guarantee



Click The Link Below

<https://certs4sale.com/hp/hpe7-j01-pdf-exam-dumps>