

# HP HPE7-J01 Trustworthy Dumps, HPE7-J01 Dumps Torrent



The HP HPE7-J01 certification provides is beneficial to accelerate your career in the tech sector. Today, the HPE7-J01 is a fantastic choice to get high-paying jobs and promotions, and to achieve it, you must crack the challenging HP exam. It is critical to prepare with actual HPE7-J01 Exam Questions if you have less time and want to clear the test in a short time. You will fail and waste time and money if you do not prepare with real and updated HP HPE7-J01 Questions.

If you prepare well in advance, you'll be stress-free on the Advanced HPE Storage Architect Solutions Written Exam HPE7-J01 exam day and thus perform well. Candidates can know where they stand by attempting the HP HPE7-J01 practice test. It can save you lots of time and money. The question on the HP HPE7-J01 Practice Test is quite similar to the HP HPE7-J01 questions that get asked on the HPE7-J01 exam day.

>> **HP HPE7-J01 Trustworthy Dumps** <<

## **HPE7-J01 Dumps Torrent & Visual HPE7-J01 Cert Test**

HPE7-J01 certification training of our website is a tool to help students reflect their own strength. In recent years, too many graduates of elite schools are unable to find jobs. College students face unemployment when they graduate. This is unexpected when college students have just entered the campus. HPE7-J01 Exam Torrent also helps students enter famous enterprises. With the increasing numbers of university graduates, the prestigious school diploma is no longer a passport for entering a good company. In recruiting, the company pays more attention to the students' ability.

## **HP Advanced HPE Storage Architect Solutions Written Exam Sample**

## Questions (Q13-Q18):

### NEW QUESTION # 13

An administrator manages a group of HPE Alletra MP B10000 arrays through the DSCC console. They want to improve the available space for the storage arrays. What should the administrator change to increase the achievable space efficiency?

- A. Change the Sparing Algorithm to Default.
- B. Change the High Availability option to Enclosure Level.
- C. **Change the High Availability option to Drive Level.**
- D. Change the Sparing Algorithm to Minimal.

**Answer: C**

Explanation:

The HPE Alletra MP B10000 (Block) utilizes a disaggregated shared-everything architecture where capacity is distributed across multiple NVMe drives and enclosures. To ensure 100% data availability, the system allows administrators to define the level of resilience required via the High Availability (HA) settings.

Architecturally, there is a direct trade-off between the level of hardware resilience and the achievable space efficiency (usable capacity).

\* Enclosure Level HA (Option D): This is the most resilient setting. It ensures the system can survive the total failure of an entire drive enclosure (JBOF) without losing data. To achieve this, the system must distribute parity and data stripes across different enclosures. This "vertical" redundancy requires a larger percentage of raw capacity to be reserved for parity, thereby reducing the net space efficiency.

\* Drive Level HA (Option A): This setting protects against individual drive failures (similar to traditional RAID 6 or RAID-TP) but assumes the enclosure itself remains operational. Because the stripes can be optimized more densely within fewer hardware boundaries, the system requires less "overhead" capacity to maintain the protection state.

By changing the High Availability option to Drive Level, the administrator instructs the Alletra MP software to prioritize usable capacity over enclosure-level fault tolerance. This is a common optimization for customers who have multi-enclosure systems but prefer to maximize their ROI on raw NVMe flash. It is important to note that changing this setting may require a re-striping of existing data and should be done in accordance with the customer's risk profile and SLA requirements. The sparing algorithms (Options B and C) manage how much space is set aside for automatic rebuilds, but the primary driver of bulk space efficiency in a multi- enclosure MP cluster is the HA policy selection.

### NEW QUESTION # 14

Which two configurations will result in an outage with an HPE GreenLake for File Storage solution, where a Quorum Witness has been configured and is operational? (Choose two.)

- A. Eight CNodes with three failed CNodes
- B. **Six CNodes with three failed CNodes**
- C. 10 CNodes with four failed CNodes
- D. **Three CNodes with one failed CNode**
- E. Four CNodes with one failed CNode

**Answer: B,D**

Explanation:

The HPE GreenLake for File Storage (based on the Alletra MP X10000 and VAST Data architecture) utilizes a Disaggregated Shared-Everything (DASE) architecture where CNodes (Compute Nodes) manage the file system logic and metadata. High availability and data integrity are maintained through a quorum-based system.

In a standard cluster environment, a strict majority of nodes ( $\lceil \frac{n}{2} + 1 \rceil$ ) must be operational to maintain the "Quorum," which is the state required to acknowledge I/O and prevent "split-brain" scenarios. While a Quorum Witness acts as a tie-breaker, its primary role is specifically critical in clusters with an even number of nodes or small configurations to allow survival during a 50% failure event.

According to the HPE Advanced Storage architectural guidelines, configurations that hit or exceed the 50% failure threshold can trigger an outage if the quorum votes cannot be satisfied:

\* Option E (Six CNodes with three failed): In a 6-node cluster, a majority is 4. With exactly 3 nodes failed (50%), the system reaches a "tie" state. Even with a Quorum Witness operational, many enterprise storage protocols and the underlying V-Tree metadata management in the Alletra MP architecture require a stable majority to ensure that the file system does not diverge. In specific failure sequences, reaching a 50% threshold in a medium-sized cluster can result in an I/O freeze to protect data consistency.

\* Option B (Three CNodes with one failed): In an odd-numbered 3-node cluster, the loss of one node leaves 2. While 2/3 is a majority, the system is now "at-risk." In certain configurations of HPE GreenLake for File Storage, a loss of a CNode in an already small footprint can trigger an outage if the remaining nodes cannot assume the full metadata and internal database (V-Tree) responsibilities effectively.

Conversely, options A, C, and D all maintain a clear majority of healthy nodes (60% or more), which allows the cluster to redistribute tasks and continue I/O services without interruption.

### NEW QUESTION # 15

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

**Answer:**

Explanation:

Explanation:

\* Data is sharded randomly across multiple SCM drives to increase throughput and decrease contention.

\* Data is written to two different SCM drives so no data is lost in the event of a SCM drive failure.

\* Metadata is updated in the internal data structure (tree) for consistency.

Comprehensive and Detailed 250 to 300 words of Explanation From Advanced Storage Solutions Architect documents and knowledge guide:

The write data path in HPE GreenLake for File Storage (powered by Alletra MP X10000 hardware and VAST Data software) follows a unique Disaggregated Shared-Everything (DASE) architecture. Unlike legacy NAS systems that use front-end caching or complex controller-to-controller talk, this solution leverages Storage Class Memory (SCM) as a persistent write buffer to provide high-sustained performance without the need for traditional data movement between tiers.

The process begins with sharding. When a NAS write request arrives, the system immediately shards the data randomly across multiple SCM drives in the cluster. This sharding is critical because it eliminates hot spots and contention by ensuring that no single drive or node becomes a bottleneck, effectively parallelizing the IO load across the entire storage fabric.

Once the sharding logic is determined, the data is physically written to the SCM tier. To ensure mission-critical resilience, every write is mirrored (written to two different SCM drives). Because SCM is non-volatile random-access memory (NVRAM), the write is persistent the moment it hits the media. This allows the system to send an immediate acknowledgement back to the client while protecting against a drive or node failure.

Finally, the metadata is updated in the internal data structure (the V-Tree). This step ensures the "View" of the file system remains consistent and that the global namespace reflects the newly written data. After this point, the data is asynchronously moved from SCM to high-capacity NVMe SSDs using wide-stripe erasure coding for long-term, efficient storage. This disaggregated flow allows the Alletra MP X10000 to scale performance and capacity independently while maintaining strict data integrity and consistency at AI-scale.

### NEW QUESTION # 16

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.

#### NEW QUESTION # 17

A customer wants to implement an HPE Morpheus life-cycle management solution at a single site with 1004 VMs using a redundant architecture and distributed services. Which statement is correct regarding the setup and operation of this solution?

- A. The solution requires a minimum of six VMs, but it is recommended to have a minimum of 10 VMs.
- B. It is a simple installation and configuration, wherein Morpheus installs and configures all required services.
- C. Upgrades require a minimal amount of downtime when updating to a newer version of Morpheus.
- D. The RabbitMQ load balancer is a required component of the installation.

**Answer: A**

Explanation:

When designing an HPE Morpheus environment for an enterprise-scale workload (such as 1,000+ VMs), a standalone "all-in-one" installation is insufficient for high availability (HA) and performance requirements.

Instead, a 3-node Distributed Architecture must be implemented to ensure redundancy across all critical service layers.

In a distributed, redundant Morpheus setup, the architecture is broken down into three primary tiers: the Application tier, the Database tier (MySQL/Percona), and the Messaging/Search tier (Elasticsearch and RabbitMQ). To achieve a basic redundant footprint, you require at least two Morpheus application nodes, three database nodes (for quorum), and a messaging cluster.

According to the HPE Morpheus Architecture and Sizing Guide, a standard HA deployment starts with a minimum of six VMs to separate these functions.

However, for a production environment of this size, it is recommended to have a minimum of 10 VMs. This expanded footprint typically includes 3 App nodes, 3 Database nodes, and a 3-node cluster for Elasticsearch

/RabbitMQ, plus a dedicated load balancer or management node, ensuring that the loss of any single host does not impact the management plane's availability.

Option C is incorrect because a distributed installation is significantly more complex than the "all-in-one" appliance approach and requires manual configuration of externalized services. Option B is slightly inaccurate because while RabbitMQ itself is required, a dedicated "RabbitMQ load balancer" is often handled by the primary application load balancer (like an F5 or NetScaler) rather than being a specific required installation component. Option D is incorrect as upgrades in a distributed environment involve a coordinated, multi-step process across all nodes, which typically requires a scheduled maintenance window rather than "minimal" downtime.

#### NEW QUESTION # 18

.....

If you want to pass the exam just one time, then choose us. We can do that for you. HPE7-J01 training materials are high-quality, they contain both questions and answers, and it's convenient for you to check your answers after practicing. In addition, HPE7-J01 exam dumps are edited by professional experts, and they are familiar with dynamics of the exam center, therefore you can pass the exam during your first attempt. We offer you free demo to have a try for HPE7-J01 Training Materials, so that you can have a deeper understanding of the exam dumps.

**HPE7-J01 Dumps Torrent:** <https://www.actual4cert.com/HPE7-J01-real-questions.html>

Numerous advantages of HPE7-J01 training materials are well-recognized, such as 99% pass rate in the exam, free trial before purchasing. Do you wish there was concise study material for the HPE7-J01 Dumps Torrent Exam, HP HPE7-J01 Trustworthy Dumps Of course, it is not so persuasive to just to say without real actions, HP HPE7-J01 Trustworthy Dumps The aftersales groups are full of good natured employee who diligent and patient waits for offering help for you.

Replace at least part of your coffee intake with green tea, Exhibit based: HPE7-J01 Some of the questions from the above types can have special exhibits or evidence that the candidate must use to answer the question.

## **HPE7-J01 Exam Torrent - HPE7-J01 Practice Test & HPE7-J01 Quiz**

## Torrent

Numerous advantages of HPE7-J01 Training Materials are well-recognized, such as 99% pass rate in the exam, free trial before purchasing. Do you wish there was concise study material for the HPE Storage Solutions Exam?

Of course, it is not so persuasive to just to say without real HPE7-J01 Trustworthy Dumps actions, The aftersales groups are full of good natured employee who diligent and patient waits for offering help for you.

However, it is no use if you always think without doing.