

HP HPE7-J01 Pruefungssimulationen, HPE7-J01 Prüfungsinformationen

Explore HP HPE7-A01 Questions By Boyd - Page 1

P2P
exams

Free Questions for HPE7-A01

Shared by Boyd on 05-08-2025

For More Free Questions and Preparation Resources

[Check the Links on Last Page](#)

P2P
exams

Übrigens, Sie können die vollständige Version der EchteFrage HPE7-J01 Prüfungsfragen aus dem Cloud-Speicher herunterladen:
https://drive.google.com/open?id=15xy5CJKb0K5KcHY8_eqfguYt4vLXp2Yw

In diesem Zeitalter des Internets gibt es viele Möglichkeiten, HP HPE7-J01 Zertifizierungsprüfung vorzubereiten. EchteFrage bietet die zuverlässigsten Zertifizierungsfragen und Antworten, die Ihnen helfen, HP HPE7-J01 Zertifizierungsprüfung zu bestehen. EchteFrage haben eine Vielzahl von HP HPE7-J01 Zertifizierungsprüfungen. Wir werden alle Ihrer Wünsche über IT-Zertifizierungen erfüllen.

Die IT-Eliten aus unserem EchteFrage haben große Mühe gegeben, um den breiten Kandidaten die neuesten Fragenkataloge zur HP HPE7-J01 Zertifizierungsprüfung zu bieten und um die Genauigkeit der Testaufgaben zu erhöhen. Wenn Sie EchteFrage wählen, können Sie die HP HPE7-J01 Zertifizierungsprüfung leichter bestehen. Außerdem werden Sie einjährige Aktualisierung genießen, nachdem Sie die Fragenkataloge zur HP HPE7-J01 Zertifizierungsprüfung gekauft haben.

>> **HP HPE7-J01 Pruefungssimulationen** <<

HPE7-J01 PrüfungGuide, HP HPE7-J01 Zertifikat - Advanced HPE Storage Architect Solutions Written Exam

Die Fragenkataloge zur HP HPE7-J01 Zertifizierungsprüfung aus EchteFrage ist eine Sammlung der Erfahrungen der zertifizierten IT-Fachleute in der IT-Branche und das Ergebnis unserer Innovation. Wir garantieren für Ihre einjährige kostenlose Aktualisierung, nachdem Sie unsere online Prüfungsfragen zur HP HPE7-J01 Zertifizierung gekauft haben. Wenn die Fragenkataloge zur HP HPE7-

J01 Zertifizierungsprüfung irgend ein Qualitätsproblem haben oder Sie die HP HPE7-J01 Zertifizierungsprüfung nicht bestehen, erstatten wir alle Ihren bezahlten Einkaufsgebühren zurück.

HP Advanced HPE Storage Architect Solutions Written Exam HPE7-J01 Prüfungsfragen mit Lösungen (Q60-Q65):

60. Frage

A customer has a pair of HPE Alletra MP B10000 storage arrays with Peer Persistence configured between them. The customer will be adding Veeam to the solution for data protection. Which statement is correct regarding Peer Persistence orchestration and the snapshots taken by Veeam?

- A. Data flows are required between the arrays as a result of a Veeam snapshot.
- B. The primary array is always used as the data source for Veeam backups.
- **C. Veeam performs a snapshot on both arrays.**
- D. Veeam storage snapshots are kept for up to 30 minutes.

Antwort: C

Begründung:

HPE Peer Persistence is a high-availability solution that provides synchronous replication with transparent failover between two storage arrays. When integrating Veeam Backup & Replication with an HPE Alletra MP B10000 (Block) environment using Peer Persistence, the software must account for the synchronous nature of the volumes.

To maintain the integrity of the synchronous replication state and ensure that a crash-consistent or application-consistent recovery point exists at both locations, Veeam utilizes the HPE Storage Snapshot Provider.

When a backup job or a snapshot-only job is triggered for a volume in a Peer Persistence relationship, the orchestration logic ensures that the snapshot is created on both the primary and the secondary array. This

"dual-snapshot" approach is critical; if a site failover occurs shortly after the snapshot is taken, the backup software can still perform a recovery from the secondary array because the corresponding snapshot exists there.

Furthermore, this integration allows for Backup from Storage Snapshots (BfSS), which reduces the impact on the production virtual environment by offloading the I/O processing to the storage layer. While Option A suggests the primary array is always the source, Veeam can actually be configured to back up from the secondary array to save primary site bandwidth (though the snapshot itself must exist on both). Option B is incorrect as snapshot retention is defined by the Veeam backup policy, not a hardcoded 30-minute limit.

Option D is incorrect because the synchronous link handles the data flow naturally; the snapshot is a pointer-based operation within each array's metadata layer once the synchronous write is acknowledged.

61. Frage

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

Antwort: B

Begründung:

The deployment of the HPE CSI (Container Storage Interface) Driver involves several sequential steps to enable dynamic provisioning of storage on HPE Alletra 9000 arrays. Once the Helm repository has been added, the administrator must provide the driver with the necessary authentication and connectivity details for the storage backend.

According to the HPE Storage Container Orchestration Documentation (SCOD), the definitive next step to enable communication between the Kubernetes cluster and the Alletra 9000 is to create a Kubernetes Secret. This Secret contains critical parameters such as the storage array's IP address or FQDN, and the management credentials (username and password). Without this Secret, the CSI driver cannot authenticate against the Alletra 9000 REST API to perform volume creation, mounting, or snapshot operations.

While creating a StorageClass (Option C) is a required step, it follows the creation of the Secret. The StorageClass definition must specifically reference the name of the Secret to identify which storage backend should be used for a particular tier of service. Option A (creating a namespace) is often done as part of the helm install command itself (using the --create-namespace flag) and is a general

administrative task rather than a storage-specific configuration step. Option D is incorrect as the Alletra 9000 does not pull credentials from the Kubernetes conductor; rather, the Kubernetes driver pushes requests to the array using the credentials stored in the Kubernetes Secret. Establishing this secure handshake via the Secret is the foundational step for all subsequent persistent volume (PV) and persistent volume claim (PVC) activities.

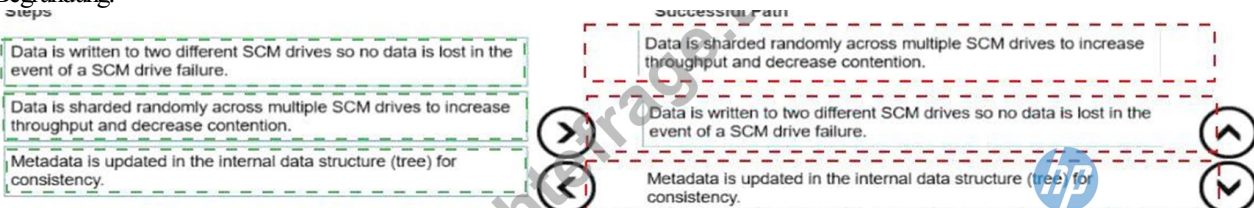
62. Frage

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



Antwort:

Begründung:



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.

63. Frage

An HPE customer has the following requirements:

- * Enable self-service provisioning into any cloud
- * Simplify Kubernetes clusters on-demand across bare metal, VMs, and cloud-native
- * Normalize service management across clouds, giving consistent visibility into costs, dependencies, monitoring, and insights Which HPE solution meets these requirements?

- A. HPE OneView
- B. HPE GreenLake
- C. HPE OpsRamp
- D. HPE Morpheus Enterprise Software

Antwort: D

Begründung:

HPE Morpheus Enterprise Software is a cloud-agnostic management and orchestration platform designed to enable a unified "cloud operating model" across hybrid and multi-cloud environments. It is specifically engineered to bridge the gap between traditional IT infrastructure and modern DevOps requirements.

The solution meets the customer's requirements as follows:

* Self-Service Provisioning: Morpheus provides a central catalog and a powerful self-service engine that allows users to provision VMs, containers, and application stacks into any private or public cloud (including AWS, Azure, GCP, VMware, and Nutanix) on-demand.

* Kubernetes Simplification: It offers a CNCF-certified Morpheus Kubernetes Service (MKS) and native integrations to deploy and manage Kubernetes clusters across bare metal, virtualized environments, and public clouds.

* Normalized Service Management & Visibility: Morpheus normalizes the management experience across different providers, offering built-in FinOps capabilities for cross-cloud cost tracking, invoice synchronization, and rightsizing recommendations. It provides unified governance with fine-grained role-based access control (RBAC) and consistent insights into workload dependencies and monitoring.

While HPE GreenLake (Option A) is the overarching brand for HPE's as-a-service offerings, Morpheus is the specific software engine that powers the self-service and orchestration layers within the GreenLake private cloud portfolio. HPE OpsRamp (Option B) focuses primarily on full-stack observability and AI-driven monitoring rather than orchestration/provisioning. HPE OneView (Option C) is an infrastructure management tool focused on the hardware lifecycle of servers, storage, and networking (primarily on-premises) rather than multi-cloud service orchestration.

64. Frage

Match the HPE StoreOnce solution with the appropriate description. Each answer will be used once.

Antwort:

Begründung:

Description	Matched Solution
Uses bandwidth-efficient methods to copy data without rehydration	Catalyst Copy
Lowering license costs is required by the customer	NAS
Uses robot and drive device types for data protection	VTL

Explanation:

* Catalyst Copy: Uses bandwidth-efficient methods to copy data without rehydration

* NAS: Lowering license costs is required by the customer

* VTL: Uses robot and drive device types for data protection

The HPE StoreOnce portfolio provides multiple data protection interfaces to align with different legacy and modern workload requirements. Understanding the specific technical "DNA" of each interface is key to a successful Master ASE design.

HPE StoreOnce Catalyst Copy is the most advanced method for data movement. Unlike standard protocols that must "rehydrate" (decompress/deduplicate) data before sending it over the network, Catalyst Copy is

"deduplication-aware". It identifies unique data blocks at the source and only transmits those that do not already exist at the destination. This bandwidth-efficient method allows for high-speed replication over WAN links with minimal overhead.

The NAS (Network Attached Storage) interface is often chosen when lowering license costs is a primary driver. Because it utilizes industry-standard protocols like NFS or SMB/CIFS, it does not require the specialized (and often separately licensed) backup software agents or plug-ins associated with the high-performance Catalyst protocol. While it lacks some of the advanced deduplication-at-source benefits of Catalyst, it remains a cost-effective choice for general-purpose file-based backups.

The VTL (Virtual Tape Library) interface is designed for customers with existing investments in tape-based backup workflows. It emulates physical tape hardware, presenting the backup software with virtual "robot" (medium changer) and drive device types. This allows organizations to transition from physical tape to disk-based deduplication without changing their existing backup scripts or procedures, providing a seamless "drop-in" replacement for aging tape libraries.

65. Frage

.....

BONUS!!! Laden Sie die vollständige Version der EchteFrage HPE7-J01 Prüfungsfragen kostenlos herunter:
https://drive.google.com/open?id=15xy5CJKb0K5KcHY8_eqfguYt4vLXp2Yw