



Vendor: HP

Exam Code: HPE7-J01

Exam Name: Advanced HPE Storage Architect Solutions
Written Exam

Version: DEMO

QUESTION 1

A storage administrator is creating a disaster recovery solution for HPE Alletra 9000 storage arrays. Currently, the company has three storage arrays at three different primary sites. When implementing the N-to-1 Remote Copy (RC) feature, what is the minimum number of storage arrays the storage administrator needs to plan for at the disaster recovery site?

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

Answer: A

Explanation:


The HPE Alletra 9000 (and its predecessor, HPE Primera) supports various Remote Copy (RC) topologies to meet different disaster recovery and data distribution requirements. These include 1-to-1, 1-to-N (fan-out), and N-to-1 (fan-in) configurations.

In an N-to-1 Remote Copy configuration, multiple source storage systems (represented by 'N') replicate their data to a single, centralized target system at a disaster recovery (DR) or secondary site. This architecture is particularly efficient for organizations with multiple regional or branch offices that wish to centralize their backup and DR operations into a single data center to reduce hardware costs and simplify management. In the scenario described, the company has three primary sites ($N = 3$), each with its own storage array. To implement an N-to-1 strategy, the administrator only needs to provide one storage array at the DR site. This single target array must be sized appropriately to handle the combined capacity and performance requirements (IOPS and throughput) of the incoming replication streams from all three source systems.

Architecturally, the Alletra 9000 uses Remote Copy Groups to manage these relationships. Each group on the source systems is mapped to a corresponding group on the single target system. It is important to note that while the hardware requirement is a single array, the administrator must ensure the target array has sufficient Remote Copy ports (RCIP or RCFC) and licensed capacity to accommodate the fan-in ratio. The Alletra 9000 management interface and HPE GreenLake Data Services Cloud Console (DSCC) provide the orchestration necessary to monitor these multiple inbound streams and ensure that the Recovery Point Objectives (RPOs) are met across all sites simultaneously.

QUESTION 2

An administrator needs to create an FCIP trunk connection between two data centers to interconnect their Brocade fibre data fabrics. Refer to the exhibit.



Based on this configuration, which statement is correct?

- A. This is an invalid configuration. FOP trunks must be configured with an even number of circuits.
- B. This is an invalid configuration. FCIP trunks must be configured using fibre channel ports.
- C. This is a valid configuration. The trunk will have one active connection and two standby connections.
- D. This is a valid configuration. The trunk will have two active connections and one standby

connection.

Answer: D

Explanation:

In a Brocade FCIP (Fibre Channel over IP) environment, an extension trunk (or tunnel) can be composed of multiple circuits to provide both increased bandwidth and high availability. The operational state of these circuits--whether they are active or standby--is determined by the Metric assigned to each individual circuit.

According to the Brocade Fabric OS Extension Configuration Guide, all circuits within a tunnel or trunk have a metric of either 0 or 1.

Metric 0: This is the default value and indicates an active circuit. If multiple circuits are configured with Metric 0, they will operate in an active-active mode, and the load will be balanced across them. Metric 1: This indicates a standby (or passive) circuit. Standby circuits with Metric 1 are not used for data transmission unless all Metric 0 circuits within that tunnel/failover group fail. In the provided exhibit, there is a single VE_Port (Virtual E_Port) trunking three individual IP circuits:

ge0 is configured with Metric 0 (Active).

ge1 is configured with Metric 0 (Active).

ge2 is configured with Metric 1 (Standby).

Therefore, this is a valid configuration where the system will utilize the two Metric 0 circuits (ge0 and ge1) simultaneously for data traffic, providing an active-active load-balanced connection. The third circuit (ge2) will remain in a standby state, only becoming active to maintain the link if both primary circuits go offline.

QUESTION 3

On which object is snapshot locking with Qumulo implemented?

- A. File
- B. Volume
- C. Directory
- D. Storage pool

Answer: C

Explanation:

The HPE Solutions for Qumulo utilize a modern, distributed file system designed specifically for the era of multi-petabyte unstructured data management. Unlike traditional block storage systems that organize data into LUNs or volumes, Qumulo uses a single, unified namespace where all data is organized into a hierarchical structure of Directories.

According to the HPE Qumulo Administration Guide, all advanced data services--including snapshots, quotas, and replication--are applied at the Directory level. When a storage administrator wants to protect a dataset, they create a snapshot of a specific directory (and all its subdirectories). The Snapshot Locking feature is an extension of this capability, designed to provide "immutable" data protection against ransomware or accidental deletion.

When snapshot locking is implemented, it is associated with a snapshot policy that targets a specific Directory path. Once a snapshot is locked, the metadata associated with that directory at that specific point in time becomes immutable; it cannot be deleted, modified, or shortened in its expiration period until the lock period has expired.

QUESTION 4

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.

QUESTION 5

What is a dependency to keep in mind regarding trunking, cable lengths, and deskew units when calculating RTT for fibre channel Brocade ISLs for optimal performance?

- A. The shortest ISL is set to a deskew value that depends on the switch hardware platform generation.
- B. Trunks can be a mixture of cable lengths, as long as all cables in the ISL use the same transceiver type.
- C. Deskew units represent the time difference for traffic to travel over a single connection of the ISL.
- D. A 20-meter difference is approximately equal to one deskew unit.

Answer: D

Explanation:

In Brocade Fibre Channel fabrics, ISL Trunking allows multiple physical links to behave as a single logical entity. For this to work efficiently, the switch must synchronize the delivery of frames across all physical links to ensure they arrive in the correct order. This process is managed by the Deskew mechanism.

"Skew" refers to the difference in time it takes for a signal to travel across the different physical cables within a trunk, often caused by slight variations in cable lengths. According to the Brocade Fabric OS Administration Guide, the switch hardware automatically measures these differences and applies "deskew units" to the faster (shorter) links to delay them, effectively matching the speed of the slowest (longest) link in the trunk.

A critical rule in SAN design is the distance limitation between cables in a trunk. While Brocade switches are highly capable of compensating for skew, the maximum supported difference in cable length within a single trunk is usually around 30 meters. For calculation purposes, one deskew unit is approximately equal to 20 meters of cable length. If the physical length difference between the shortest and longest cable exceeds the hardware's deskew buffer capacity (which varies by ASIC generation but is measured against this 20m/unit metric), the trunk will fail to

initialize or will experience significant performance degradation.

QUESTION 6

A customer intentionally removes all three drives from a JBOF from an HPE Alletra MP X10000 used in an HPE GreenLake for File Storage solution. What is the correct description of the result of this action?

- A. This results in the start of a rebuild on an integrated spare drive. After the rebuild/resync of the spare is complete, data services will be available and I/O will continue.
- B. This results in a catastrophic failure with an I/O outage, as no data service is available and the system is not in RO mode. To recover, the customer should reinsert the same three drives, which will recover the system.
- C. This results in a catastrophic failure, with an I/O outage as no data service is available and the system is not in RO mode. To recover, HPE support/engineering needs to be involved and try the recovery steps.
- D. This results in the start of a rebuild on an integrated spare drive. After the rebuild/resync of the spare is complete, the customer must contact HPE support/engineering to complete the process.

Answer: C

Explanation:

The HPE Alletra MP X10000, which powers HPE GreenLake for File Storage, utilizes a disaggregated shared-everything (DASE) architecture based on VAST Data software. Unlike traditional RAID, this architecture uses highly advanced locally decodable erasure coding. While the system is designed to be incredibly resilient--often surviving multiple concurrent drive failures across the cluster--the removal of three drives simultaneously from a single JBOF (Just a Bunch of Flash) chassis can exceed the immediate "vertical" stripe protection thresholds, especially in smaller cluster configurations. In the Alletra MP File architecture, the metadata and data are distributed with specific redundancy parameters. Intentionally pulling three drives at once is treated as a multi-point catastrophic failure rather than a standard drive wear-out event. When such an event occurs, the system enters a "Fail-Stop" state to protect data integrity and prevent file system corruption. Because the system cannot guarantee the consistency of the data stripes or the underlying V-Trees (metadata structures), it will cease I/O services.

QUESTION 7

An HPE Partner is designing a software-defined storage (SDS) solution that includes HPE Alletra 4000 storage servers and the HPE Ezmeral Data Fabric software solution. The customer wants to manage the HPE Alletra 4000 storage servers using HPE GreenLake. Which component in HPE GreenLake should the customer use?

- A. Compute Ops Manager
- B. File Storage
- C. Block Storage
- D. Data Ops Manager

Answer: A

Explanation:

The HPE Alletra 4000 series (specifically the Alletra 4110 and 4120) are technically classified as Storage Servers. Unlike traditional "closed" storage arrays like the Alletra 6000 or 9000, the Alletra 4000s are open platforms derived from the HPE Apollo lineage, designed to run Software-Defined Storage (SDS) stacks such as HPE Ezmeral Data Fabric, Scality, or Qumulo. Because these systems are fundamentally high-density servers, their lifecycle management--including firmware updates (BIOS, iLO, controllers), health monitoring, and remote configuration--is integrated into the HPE GreenLake for Compute Ops Management (COM) service. COM

provides a cloud-native console designed specifically for server administrators to manage fleets of ProLiant and Alletra 4000 servers from a single pane of glass.

QUESTION 8

An administrator needs to use a plug-in to provide monitoring of the performance of applications, servers, databases, tools, and services through an SaaS-based data analytics platform. Which Morpheus plug-in should the administrator install?

- A. Morpheus Agent
- B. Datadog
- C. HPE OpsRamp
- D. Morpheus Instances

Answer: B

Explanation:

HPE Morpheus Enterprise provides extensive native monitoring for provisioned resources, but for deep, full-stack observability into applications, databases, and third-party services via a specialized SaaS-based analytics platform, it integrates with external providers through plugins. The Datadog Integration Plugin is a primary example of this capability. Datadog is an industry-leading SaaS data analytics platform that monitors performance across every layer of the stack--from infrastructure metrics to application traces and logs. By installing the Datadog plugin in Morpheus, the administrator bridges the gap between the orchestration engine (Morpheus) and the observability platform (Datadog). This plugin adds a contextual "Datadog" tab to instances within the Morpheus UI, pulling real-time performance telemetry directly from the SaaS backend into the Morpheus management pane. This allows the administrator to view workload health and deep-dive analytics without switching consoles.

QUESTION 9

What is a prerequisite for a successful Fibre Channel storage array Peer Motion migration?

- A. An N-Port ID Virtualization (NPIV) capable FC fabric between the source and destination.
- B. IP connectivity to initiate and control the data migration.
- C. A direct connection between storage arrays via their FC ports.
- D. The configuration of a maximum of eight peer link pairs.

Answer: A

Explanation:

The HPE Peer Motion Utility (PMU) and its integrated counterpart in HPE GreenLake and SSMC are designed for the non-disruptive migration of data between storage systems, such as from an HPE 3PAR to an HPE Alletra 9000 or Primera. A core requirement for the "Online" (non-disruptive) version of this migration is that the storage fabric must support and have N-Port ID Virtualization (NPIV) enabled. Architecturally, Peer Motion relies on the destination array's ability to "masquerade" as the source array during the transition. When a volume is migrated, the destination array creates virtual ports using NPIV to inherit the identity (WWNs) of the source array's ports. This allows the host's multipathing software to see the new storage paths as if they were additional paths to the original volume, enabling a seamless transition without a server reboot or I/O interruption. According to the HPE Peer Motion Utility User Guide, if the SAN fabric (the switches) does not support NPIV or if NPIV is disabled on the specific ports, the migration utility will default to a Minimally Disruptive Migration (MDM) or an offline migration, both of which involve host-side downtime.

Furthermore, the fabric must be zoned such that the source and destination arrays can "see" each other to establish the Peer Motion relationship and handle the data orchestration.

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.

View list of all certification exams: <http://www.lead2pass.com/all-products.html>



Microsoft



ORACLE



CITRIX



JUNIPER
NETWORKS



EMC²
where information lives



10% Discount Coupon Code: ASTR14