


HP.HPE7-J01.v2026-05-08.q28

Exam Code:	HPE7-J01
Exam Name:	Advanced HPE Storage Architect Solutions Written Exam
Certification Provider:	HP
Free Question Number:	28
Version:	v2026-05-08
# of views:	122
# of Questions views:	280
https://www.freepdfdumps.com/HP.HPE7-J01.v2026-05-08.q28.html	

NEW QUESTION: 1

Refer to the exhibit.



A company is implementing a disaster recovery solution. The Asynchronous Remote Copy feature has been implemented between the HPE AUetra 9000 arrays at both sites. The customer is interested in providing a disaster recovery (DR) solution that allows for business connectivity of their VMware VMs.

Which VMware solution should the company implement?

- A. vCenter Lifecycle Management Service
- B. VMware Live Site Recovery/VMware Live Recovery
- C. VCF Operations: Continuous Performance
- D. vCenter Storage DRS

Answer: B ([LEAVE A REPLY](#))

To provide automated orchestration and business continuity for VMware virtual machines in a disaster recovery scenario, the industry-standard solution integrated with HPE storage is VMware Live Site Recovery (formerly known as VMware Site Recovery Manager or SRM).

When a customer utilizes Asynchronous Remote Copy on HPE Alletra 9000 arrays, the storage layer handles the data replication between the production and recovery sites. However, the storage array alone cannot automate the re-registration of virtual machines, the mapping of network port groups, or the specific power-on sequencing required for complex applications at the secondary site. VMware Live Site Recovery serves as the orchestration engine that bridges this gap. It works in conjunction with a Storage Replication Adapter (SRA) provided by HPE. The HPE SRA allows the VMware software to communicate directly with the Alletra 9000 arrays to initiate tasks such as promoting recovery volumes to a read-write state, taking temporary snapshots for DR testing, and automating the "failover" and "failback" workflows.

As shown in the exhibit (image_6601ef.jpg), a complete solution requires an SRM appliance and a vCenter appliance at both the production and recovery sites. This architecture ensures that even if the primary site is completely lost, the recovery site has all the necessary metadata and orchestration instructions to bring the business-critical VMs online with minimal manual intervention. Option A (Lifecycle Management) is for patching and updates, Option D (Storage DRS) is for load balancing within a cluster, and Option C refers to operational monitoring rather than disaster recovery orchestration. For a customer already invested in Alletra 9000 Remote Copy, VMware Live Site Recovery is the "Better Together" choice for achieving low Recovery Time Objectives (RTO).

NEW QUESTION: 2

A customer purchased an HPE GreenLake for File Storage solution and implemented the replication feature.

Which statement is correct regarding this feature?

- A.** Client hardware has read-write access to both the source and destination replication arrays.
- B.** Data reduction, including deduplication, is performed between the storage arrays.
- C.** N:1 and 1:N replication is supported, with snapshots taken at the directory level.
- D.** Two protected paths configured on the same path can be used to replicate to the same peer.

Answer: C ([LEAVE A REPLY](#))

HPE GreenLake for File Storage is built upon a disaggregated, shared-everything (DASE) architecture powered by VAST Data software. The replication mechanism in this environment is fundamentally different from traditional block-based replication. Instead of replicating entire volumes or LUNs, HPE GreenLake for File Storage performs replication at the directory level. According to the HPE GreenLake for File Storage Administrator Guide, the system utilizes a snapshot-based asynchronous replication engine. This allows for highly flexible topologies, including N:1 (fan-in) and 1:N (fan-out) configurations, which are essential for modern distributed data environments and centralized backup strategies. Because the solution is file-based, it leverages "Views" (or shares) that point to specific directory paths. Protection policies and snapshot schedules are applied

directly to these paths, ensuring that only the specific datasets required for disaster recovery are replicated.

Option B is a common point of confusion; while the system is inherently "reduction-aware" and uses similarity-based data reduction (deduplication and compression) to save space on the physical media, the replication process itself focuses on the metadata and unique data blocks associated with the directory-level snapshots. Option A is incorrect because, in an asynchronous replication relationship, the destination is typically Read-Only until a failover or clone operation is initiated. Option D is incorrect as the management of protected paths follows strict pairing rules to prevent configuration conflicts. Thus, the support for flexible fan-in/fan-out topologies and granular directory-level protection (Option C) is the defining characteristic of this enterprise file solution.

NEW QUESTION: 3

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 ([LEAVE A REPLY](#))

In a Zerto environment, a Virtual Protection Group (VPG) is the fundamental unit of management used to group virtual machines that must be replicated together to maintain write-order fidelity and application consistency. This is particularly vital for multi-tier applications, such as a database server and a web server, that need to be recovered to the exact same point in time.

According to the HPE Advanced Storage Solutions technical guides and Zerto's architectural specifications, a single Virtual Machine (VM) can be associated with a maximum of three VPGs simultaneously. This capability is often referred to as "one-to-many" replication. This architectural flexibility allows a storage administrator to design complex data protection strategies that go beyond simple site-to-site disaster recovery.

For example, a VM could be part of:

- * A Local VPG for high-speed recovery from the local journal (Short-term retention).
- * A Remote VPG for disaster recovery to a secondary data center or public cloud.
- * A Tertiary VPG for long-term retention or to a third site for regional disaster protection.

When a VM is protected in multiple VPGs, each VPG maintains its own independent journal, settings, and Recovery Point Objective (RPO) targets. However, the Virtual Replication Appliance (VRA) on the host only needs to read the data changes (IOs) from the hypervisor once; it then distributes those changes to all the target VRAs associated with the various VPGs. This ensures that while the VM is highly protected across multiple locations, the overhead on the production host and the hypervisor remains minimal. It is important to note that while three is the maximum,

the storage architect must ensure that the available network bandwidth and the IOPS of the target storage systems can handle the aggregate replication load of all associated VPGs.

NEW QUESTION: 4

On which object is snapshot locking with Qumulo implemented?

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

Answer: C (LEAVE A REPLY)

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. This is fundamentally different from legacy architectures where you might lock a whole

"Volume" (Option B), which can be inefficient for large datasets. Because Qumulo is a scale-out file system without the concept of traditional volumes or LUNs, "Directory" is the most granular and correct object for implementing these protection policies. This allows for massive flexibility, enabling administrators to set different retention and locking policies for different departments or projects (e.g., /marketing vs /research) all within the same physical cluster without needing to pre-allocate storage pools (Option D).

NEW QUESTION: 5

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.


NEW QUESTION: 6

An administrator needs to cable a new HPE Alletra MP B10000.

Match the number in the graphic below with its description. Not every answer will be used.



Answer:



Explanation:

- * Admin port: 2
- * Back-end and cluster connectivity (southbound): 4
- * Front-end connectivity (northbound): 3
- * iLO Ethernet port: 1

The HPE Alletra MP B10000 architecture is defined by its disaggregated nature, separating the controller nodes from the physical storage capacity. Proper cabling of these nodes is fundamental to ensuring the performance and manageability of the cluster.

The ports on the rear of a controller node are categorized by their function within the disaggregated ecosystem:

iLO Ethernet port (1): This is the dedicated management port for the Integrated Lights-Out (iLO) processor. It provides out-of-band management, allowing administrators to access the hardware's health, power status, and remote console regardless of the operating system state.

Admin port (2): This RJ45 port is used for the management plane of the storage OS. It allows the array to communicate with the HPE GreenLake Data Services Cloud Console (DSCC) or local management software for configuration and provisioning tasks.

Front-end connectivity / Northbound (3): These ports (typically located in PCIe slots 1 and 2) provide the

"Northbound" path to the host fabric. They are populated with high-speed Fibre Channel or iSCSI adapters to facilitate data access for servers and applications.

Back-end and cluster connectivity / Southbound (4): In a disaggregated shared-everything (DASE) architecture, "Southbound" connectivity is critical. These ports (often 100GbE RoCEv2) connect the controllers to the storage fabric and the disaggregated NVMe JBOFs. This allows every controller to see every drive in the cluster with ultra-low latency.

Understanding the distinction between Northbound (Host) and Southbound (Storage/Cluster) cabling is vital for a Master ASE. Incorrectly cabling the Southbound ports can lead to a "cluster split" or performance bottlenecks, as the disaggregated architecture relies on high-bandwidth, redundant paths to the underlying flash media to maintain its mission-critical performance profile. Number 5 in the diagram often represents the Serial console port, which is utilized primarily for initial setup or low-level troubleshooting by HPE Support and is typically not included in primary functional matching.

NEW QUESTION: 7

An HPE customer purchased an HPE B-Series SN7000B SAN fabric switch. QoS is currently not enabled.

Which two statements are correct regarding buffer-to-buffer (BB) credits and the operation of the switch?

(Choose two.)

- A. Each user port reserves eight buffer credits when online or offline.
- B. By default, all BB credits are reserved.
- C. BB credits are based on link speed and frame size.
- D. The default window size for fibre channel (FC) frame transmission is 1, but can be increased to 8 or 16, depending on the switch model.
- E. BB credits can be adjusted for specific applications or operating environments, but they must be agreed upon among all switches to allow the formation of the fabric.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 8

A company bought an HPE StoreOnce solution as part of its data protection solution. The company has various Oracle installations that need to be backed up to StoreOnce. How should the company's administrator best implement the data protection strategy within the HPE StoreOnce user interface (UI)?

- A. Under Data Services, create a Catalyst Store and select the Oracle RMAN option.
- B. From the System Dashboard, click Catalyst Store then specify the Oracle RMAN option and the respective database server.
- C. From the System Dashboard, click Databases, click Create Library, then specify the Oracle RMAN option and the respective database servers.
- D. Under Data Services, create a Catalyst Store and install the Oracle RMAN plug-in on the Oracle database server.

Answer: D ([LEAVE A REPLY](#))

To protect Oracle databases using HPE StoreOnce, the preferred architectural method is using HPE StoreOnce Catalyst for Oracle RMAN. This integration allows Oracle Database Administrators (DBAs) to manage backups directly from their native RMAN (Recovery Manager) tools while leveraging the deduplication and performance benefits of the StoreOnce appliance. According to the HPE StoreOnce Catalyst for Oracle RMAN User Guide, the implementation involves two distinct stages: configuration on the StoreOnce appliance and configuration on the database server. First, the storage administrator must log into the StoreOnce UI and, under the Data Services section, navigate to Catalyst. Here, they must create a Catalyst Store. This store acts as the target repository for the backup data.

During creation, the administrator sets permissions (client access) to allow the Oracle server to communicate with this specific store.

The second, and crucial, part of the implementation (as noted in Option D) is the installation of the HPE StoreOnce Catalyst Plug-in for Oracle RMAN on the actual Oracle database server. This plug-in provides the "SBT" (System Backup to Tape) interface that RMAN requires to talk to a non-disk/non-tape target.

Without this plug-in installed on the host, RMAN has no way of translating its commands into the Catalyst protocol. Once the plug-in is installed and configured with the StoreOnce details, the DBA can allocate channels to the "SBT_TAPE" device and run backup jobs directly to the Catalyst Store created in the UI.

Options A, B, and C are incorrect because the StoreOnce UI does not have an "Oracle RMAN option" toggle or "Database Library" creator; the intelligence resides in the combination of the Catalyst Store and the host- side plug-in.

NEW QUESTION: 9

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

	Answer Area	
Accommodates diverse geographic conditions and location		Meshed fabric
Data access is localized with servers and storage connected to the same switch		Cascaded fabric
Data access is a mix of local and distributed		Core-edge fabric
Has many-to-many connectivity and requires high performance		Ring fabric



Answer:

Accommodates diverse geographic conditions and location
Data access is localized with servers and storage connected to the same switch
Data access is a mix of local and distributed
Has many-to-many connectivity and requires high performance

Answer Area

Has many-to-many connectivity and requires high performance	Meshed fabric
Data access is localized with servers and storage connected to the same switch	Cascaded fabric
Data access is a mix of local and distributed	Core-edge fabric
Accommodates diverse geographic conditions and location	Ring fabric

Explanation:

- * Meshed fabric: Has many-to-many connectivity and requires high performance.
- * Cascaded fabric: Data access is localized with servers and storage connected to the same switch.
- * Core-edge fabric: Data access is a mix of local and distributed.
- * Ring fabric: Accommodates diverse geographic conditions and location.

In the design of modern Storage Area Networks (SANs), selecting the right topology is critical for balancing performance, scalability, and cost. Each of these Brocade/HPE B-series architectures serves a specific workload profile:

Meshed Fabric: This design provides the highest level of redundancy and performance by connecting switches in an "any-to-any" pattern. Because every switch is connected to multiple other switches, it provides many-to-many connectivity and minimized "hop counts," making it ideal for high-performance environments where application traffic is unpredictable and widely distributed.

Cascaded Fabric: The simplest topology, where switches are connected in a serial or daisy-chain fashion. It is most effective in small environments where data access is localized—meaning the server and the storage it needs are physically connected to the same switch, minimizing Inter-Switch Link (ISL) traversal.

Core-Edge Fabric: The standard for enterprise data centers. Servers are connected to "Edge" switches, and storage is connected to a high-capacity "Core". This allows for a mix of local and distributed data access.

Large enterprises use this to scale easily by adding edge switches without disrupting the core storage connectivity.

Ring Fabric: By connecting switches in a closed loop, this topology is designed to accommodate diverse geographic conditions. If a single link between two locations fails, traffic can be re-routed the other way around the ring. This provides a cost-effective way to link multiple sites or campus buildings while maintaining fabric integrity.

NEW QUESTION: 10

Match the Brocade virtual fabric term with its description.

ISLs that are configured between an edge fabric E_Port and an FC Router EX_Port	DISL
Connects two logical switches in two different chassis via the base switch to extend the fabric and maintain the logical partitioning	XISL
Directly connects two base switches that are in separate physical chassis together and has a link cost of 510	LISL
Used to link fabrics across geographic locations via FCR or FCIP	IFL

Answer:

ISLs that are configured between an edge fabric E_Port and an FC Router EX_Port	DISL
Connects two logical switches in two different chassis via the base switch to extend the fabric and maintain the logical partitioning	XISL
Directly connects two base switches that are in separate physical chassis together and has a link cost of 510	LISL
Used to link fabrics across geographic locations via FCR or FCIP	IFL

Explanation:

LISL: Directly connects two base switches that are in separate physical chassis together and has a link cost of

510

XISL: Connects two logical switches in two different chassis via the base switch to extend the fabric and maintain the logical partitioning
DISL: ISLs that are configured between an edge fabric E_Port and an FC Router EX_Port
IFL: Used to link fabrics across geographic locations via FCR or FCIP
 Brocade Virtual Fabrics (VF) allow a single physical switch to be partitioned into multiple logical switches, each with its own data and control planes. This architectural flexibility requires specialized Inter-Switch Link (ISL) types to maintain logical isolation across physical chassis.

LISL (Logical ISL): These are logical links that directly connect two Base Switches located in separate physical chassis. A defining characteristic of an LISL in Brocade Fabric OS is its default link cost of 510, which ensures it is typically used only for specific inter-fabric control traffic unless manually adjusted.

XISL (Extended ISL): An XISL is a transport link used to connect two logical switches residing in different physical chassis by tunneling through the Base Fabric. This allows the administrator to extend a single logical fabric across multiple physical devices while maintaining strict logical partitioning and reducing the number of physical cables required between chassis.

DISL (Dedicated ISL): These links are specifically configured between an edge fabric E_Port and an FC Router EX_Port. They are used in Fibre Channel Routing (FCR) topologies to provide a dedicated path for inter-fabric traffic between a standard fabric and a meta-fabric router.

IFL (Inter-Fabric Link): IFLs are the foundational links used to connect disparate fabrics across geographic locations. They utilize either Fibre Channel Routing (FCR) or FCIP tunneling to enable communication between devices in different fabrics without merging them into a single

logical entity. This is a key component for large-scale disaster recovery and data distribution architectures where fabric stability and distance are primary concerns.

NEW QUESTION: 11

A customer is considering purchasing an HPE StoreOnce solution to meet their VMware VM data protection needs. Which two Veeam features, paired with HPE StoreOnce, will validate the restoration of backups and allow the customer to easily create labs for validation testing from backups? (Choose two.)

- A. Veeam Proxy-to-Proxy
- B. Veeam SureBackup
- C. Veeam Data Immutability
- D. Veeam DataLabs
- E. Veeam Virtual Lock

Answer: B,D (LEAVE A REPLY)

The integration of Veeam Backup & Replication with HPE StoreOnce via the StoreOnce Catalyst protocol offers specialized features that transform static backup data into active assets for testing and verification. To address the specific requirements of validating restores and creating testing environments, Veeam utilizes the SureBackup and DataLabs features.

Veeam SureBackup is the primary feature used to validate the recoverability of backups. It automates the process of testing a backup by booting the VM directly from the StoreOnce repository in an isolated environment (Virtual Lab). It performs heartbeats, network tests, and application-specific tests (such as checking if a SQL database responds) to ensure the data is not just present, but functional. Because HPE StoreOnce Gen4+ systems feature flash-accelerated metadata and improved random read performance, they are "Veeam Ready" to support the I/O requirements of starting VMs for these verification tasks.

Veeam DataLabs (formerly known as Virtual Lab or On-Demand Sandbox) is the broader orchestration framework that allows customers to leverage their backup data for "Labs." By creating a DataLab, an administrator can spin up an isolated copy of the production environment from the StoreOnce backups. This is used for "validation testing," such as testing security patches, verifying application upgrades, or performing "DevOps" tasks without impacting the production systems. Option C (Immutability) and Option E (Virtual Lock) are critical for ransomware protection but do not directly facilitate the "lab creation" or "validation testing" workflows described. Option A is a general data transfer mechanism. By selecting B and D, the storage architect ensures a solution that not only protects data but proactively proves its readiness for disaster recovery.

NEW QUESTION: 12

A company is going to upgrade a SAP HANA solution. The company is looking for competitive bids, and only SAP HANA hardware that is certified should be included in a bid. When building the bid, what must you first determine before you can right-size the solution with the appropriate HPE hardware?

- A. IOPS rate
- B. Read cache size
- C. Number of HANA nodes
- D. Replication features

Answer: (SHOW ANSWER)


Sizing a storage solution for SAP HANA is fundamentally different from sizing general-purpose virtualization workloads. SAP HANA is an in-memory database, but it has extremely strict requirements for the underlying persistent storage layer to ensure data integrity during savepoints and log writes. SAP enforces these requirements through the SAP HANA Tailored Data Center Integration (TDI) program.

To begin the sizing process and ensure the solution will pass the SAP Hardware Configuration Check Tool (HWCCT) or the newer SAP HANA System Check, a storage architect must first determine the required IOPS rate, specifically for the /hana/data and /hana/log volumes. SAP provides specific KPIs for latency and throughput that must be met. For instance, the log volume requires extremely low-latency writes to handle the sequential redo logs, while the data volume requires high-throughput (MB/s) and specific IOPS to handle asynchronous savepoints.


While the number of nodes (Option C) and replication features (Option D) are important for the overall architecture, they do not dictate the "right-sizing" of the storage performance tier in the same way the IOPS and throughput requirements do. If the storage cannot meet the SAP-certified IOPS and latency thresholds, the entire solution will be unsupported, regardless of how many nodes are present. By identifying the IOPS and throughput needs first, the architect can determine if the customer requires an All-Flash Alletra 9000 or if an Alletra MP configuration with specific drive counts is necessary to provide the required "parallelism" to hit SAP's performance targets.

NEW QUESTION: 13

An administrator is setting up Zerto for data protection. The Zerto Virtual Manager <ZVM> has been installed and configured. The administrator is ready to deploy the Zerto Virtual Replication Appliances (VRAs) Use your cursor to place a + on the ZVM screen to deploy the VRAs.



Answer:



Explanation:

Sites

The Zerto Virtual Replication Appliance (VRA) is a critical architectural component of the Zerto solution.

It is a lightweight Linux-based virtual machine that must be installed on every hypervisor host (ESXi or Hyper-V) in both the production and recovery sites that will host protected virtual machines. The VRA is responsible for intercepting the I/O writes from the VMs and asynchronously replicating them to the recovery site.

In the Zerto Virtual Manager (ZVM) web interface shown in the exhibit (image_649973.png), the management of site-level infrastructure is centralized within the "Sites" tab. By clicking on this tab, the administrator accesses the view for both local and paired remote sites. Within the Local Site details, there is a sub-tab for "VRAs" or "Hosts" where the administrator can see a list of all hosts managed by the connected vCenter or SCVMM.

To deploy a VRA, the administrator selects the host(s) currently missing the appliance and clicks the "Install" button. The ZVM then automates the deployment, configuring the VRA with the necessary network settings and registering it with the hypervisor. Notably, the "Active alerts" window in the exhibit explicitly provides a hint: it shows an alert stating that "Host esx-e02.telabs.local has no VRA installed". Clicking on the "Sites" tab is the foundational step to resolve this alert and finalize the infrastructure setup required to begin creating Virtual Protection Groups (VPGs).

NEW QUESTION: 14

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: ([SHOW ANSWER](#))

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. Options 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.

NEW QUESTION: 15

A customer purchased a data protection solution that includes Cohesity and a mixture of HPE Alletra 4000 storage servers. Which management tool should the customer use to manage their Cohesity policies?

- A.** HPE GreenLake Cohesity
- B.** HPE GreenLake Data Ops Manager
- C.** Cohesity Helios
- D.** Cohesity SpanFS

Answer: **C** ([LEAVE A REPLY](#))

The management of an HPE Solution with Cohesity is centered around providing a unified, global experience across hybrid and multi-cloud environments. For managing data protection policies, alerting, and operational oversight across one or more Cohesity clusters, the correct tool is Cohesity Helios.

Cohesity Helios is a SaaS-based management platform that provides a "single pane of glass" for the entire Cohesity data estate. According to HPE and Cohesity technical documentation, Helios utilizes machine learning and AI-driven analytics to offer proactive health monitoring and global search capabilities. It allows administrators to define a single set of data protection policies—covering variables like frequency, retention, and replication—and apply them universally across clusters located on-premises (on HPE Alletra 4000 servers), at the edge, or in the public cloud. In contrast, SpanFS (Option D) is the underlying web-scale distributed file system that powers the Cohesity DataPlatform, but it is not a management tool itself. HPE GreenLake Data Ops Manager (Option B) is part of the HPE Data Services Cloud Console (DSCC) primarily used for managing native HPE Alletra Block and File storage arrays, rather than third-party software-defined platforms like Cohesity. While the solution can be procured via HPE GreenLake Flex (Option A), the operational day-to-day management of the software policies resides within the Helios console to ensure consistency with Cohesity's broader ecosystem. Helios ensures that as the customer scales their Alletra 4000 footprint, the management of their secondary data remains simplified and policy-driven.

NEW QUESTION: 16

A customer has an older HPE StoreOnce Gen3 data protection solution. They do not want to upgrade the hardware, but they do want to integrate the existing solution with AWS using HPE Cloud Bank Storage.

Other than HPE Cloud Bank licenses, what must also be included in the bill of materials (BOM)?

- A.** StoreOnce VSA appliance license
- B.** Object store license
- C.** Catalyst license
- D.** RAM upgrade

Answer: ([SHOW ANSWER](#))

HPE Cloud Bank Storage is an extension of the StoreOnce Catalyst protocol that allows for the movement of deduplicated data to object storage in the cloud. When retrofitting this technology onto older HPE StoreOnce Gen3 hardware, there are specific hardware prerequisites that must be satisfied for the feature to be supported and performant.

The primary technical constraint on Gen3 systems (such as the StoreOnce 3100, 3500, 5100, and 5500) is the overhead required to manage the massive metadata associated with cloud-tiering. For the StoreOnce system to effectively index, deduplicate, and track data chunks residing in a remote AWS S3 bucket, it requires additional system memory. According to the HPE StoreOnce QuickSpecs and Configuration Guides, a RAM Upgrade Kit (Memory Upgrade) is a mandatory BOM component for Gen3 systems if the combined local and Cloud Bank Storage capacity will exceed the original system limits or if the Cloud Bank feature is being enabled for the first time on specific entry-to-midrange models.

Without the additional RAM, the Gen3 appliance may lack the necessary resources to run the Catalyst Cloud Bank services alongside local backup operations, leading to severe performance degradation or the inability to create a Cloud Bank store. While a Catalyst license (Option C) is

technically required for Cloud Bank to function, most Gen3 customers seeking Cloud Bank already utilize Catalyst; however, the RAM upgrade is the physical hardware prerequisite that is often overlooked in "license-only" upgrades. Options A and B are incorrect as the VSA is a separate virtual product and the "Object store" is a destination, not a StoreOnce hardware component.

Valid HPE7-J01 Dumps shared by Actual4test.com for Helping Passing HPE7-J01 Exam! Actual4test.com now offer the **newest HPE7-J01 exam dumps**, the Actual4test.com HPE7-J01 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com HPE7-J01 dumps with Test Engine here: https://www.actual4test.com/HPE7-J01_examcollection.html (62 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 17

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: (SHOW ANSWER)

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.

NEW QUESTION: 18

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: ([SHOW ANSWER](#))

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. Scalify 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.

NEW QUESTION: 19

An HPE Partner is sizing an Alletra MP B10000 array for a customer using the HPE Ninja Sizer online tool.

When tuning the solution, the partner is not getting the best read and write performance from the array.

Use your cursor to place a + where the partner can click in the web GUI to begin to determine how to correct this Issue and appropriately size the solution for the customer.



Answer:



Explanation:

Performance Estimation

The HPE Ninja Sizer is a vital tool for storage architects to model the "Performance vs. Capacity" trade-offs of the HPE Alletra MP B10000 disaggregated architecture. In this scenario, the current configuration (as shown in the exhibit) consists of 4 nodes and 1 JBOF (with 8 drives of 15.36TB each).

To address poor read/write performance in Ninja Sizer, the partner must analyze the Performance Estimation pane. By clicking the chart/graph icon in the bottom right of the summary dashboard, the tool opens a detailed breakdown of the IOPS, throughput, and latency limits for the selected hardware. In a disaggregated shared-everything (DASE) architecture like the Alletra MP, performance is a direct function of the "node-to-drive" ratio. A configuration with only 8 drives (as seen in image_6491d0.png) may be "disk-bound," meaning the controllers have more compute power than the backend drive count can support in parallel.

By accessing the performance details, the partner can see if the bottleneck is at the node level (CPU/Cores) or the drive level (Flash bandwidth). To "correct this issue" and optimize the sizing, the partner would typically need to increase the drive count within the JBOF to allow for more parallel I/O streams or add additional JBOFs to scale the backend bandwidth. The Ninja Sizer provides real-time feedback; as the partner adds more 15.36TB drives, they will see the Performance Estimation bars move from yellow/red toward green, indicating an appropriately balanced solution that meets the customer's latency and throughput SLAs.

NEW QUESTION: 20

A customer is concerned about the long distances between their data centers and significant latencies that might exist between the SAN fabrics at the two data centers. Since SCSI write operations can involve multiple handshake messages between the target and initiator, which Brocade feature should be used to double the recommended distance, but maintain the same latency as a shorter haul link?

- A. FastWrite
- B. FCIP trunking
- C. Write Acceleration
- D. Leave Fast

Answer: A ([LEAVE A REPLY](#))

Standard SCSI write operations are inherently sensitive to distance because they require multiple round-trip handshakes before data is actually transmitted. A typical write involves: 1) the Command, 2) a Transfer Ready (XFER_RDY) response from the target, 3) the Data, and 4) the Status. In a long-distance SAN, each of these round trips adds significant "latency wait time," severely degrading performance as distance increases.

To solve this, Brocade (HPE B-series) utilizes a protocol optimization feature known as FastWrite. FastWrite works by creating a Proxy Target (PT) local to the initiator host and a Proxy Initiator (PI) local to the target storage device. When the host issues a SCSI write command, the local Brocade switch (acting as the Proxy Target) immediately sends the XFER_RDY back to the host without waiting for the signal to travel across the long-distance link. This allows the host to send the data segment immediately.


By eliminating the need for every handshake message to traverse the distance multiple times, FastWrite significantly reduces the aggregate latency felt by the application. Architecturally, this enables customers to extend their SAN fabrics over double the distance (and often much further) while maintaining performance comparable to a significantly shorter link. This is critical for asynchronous replication and remote copy applications that issue large I/O blocks. Option C (Write Acceleration) is a generic term often used by other vendors, while FastWrite is the specific, validated Brocade feature name used in HPE Master ASE documentation for this protocol optimization.

NEW QUESTION: 21

An administrator needs to verify that the HPE plug-in was installed in VMware vCenter to correctly integrate with the HPE Alletra storage array. Use your cursor to place a + where the administrator can click to verify the plug-in installation.



Answer:



Explanation:

Configure

In a VMware vSphere environment, integrating HPE Alletra storage typically involves the installation of the HPE Storage Peer Motion Utility or the HPE Storage Connection Manager for VMware, which provides the VASA (vSphere APIs for Storage Awareness) provider functionality. This integration allows vCenter to manage storage arrays directly for tasks such as provisioning Virtual Volumes (vVols) and monitoring storage health.

According to the HPE Alletra Integration Guide for VMware, after the plug-in or provider is registered with vCenter, an administrator can verify its status through the vSphere Client UI. As shown in the exhibit (image_64a7b7.png), the administrator is currently viewing the Storage Providers section under the Configure tab of the vCenter object. This specific view is the correct location to verify that the HPE VASA provider is active and communicating. The list under "Storage Provider/Storage System" should show the Alletra array's URL and a status of "Active" or "Online." If the administrator wants to verify the registration of the UI-based plug-in itself (which adds specific HPE menus and dashboards), they should click the three dots (...) at the end of the navigation bar. This reveals additional hidden tabs where specialized vendor extensions, such as the HPE Storage dashboard, are located.

Successful installation is confirmed if the HPE-specific provider is listed with a valid certificate and the "Last Rescan" time is current. Failure to see the provider in this list indicates that the registration process (typically performed via the array's management console or a separate appliance) was not completed successfully, which would prevent the use of advanced features like Storage Policy Based Management (SPBM).

NEW QUESTION: 22

A customer has a variety of HPE Alletra arrays running various VM workloads managed by VMware vCenter. These arrays are managed with HPE GreenLake. The customer is concerned

they are not getting the best efficiency from the HPE storage arrays. You access CloudPhysics in GreenLake to get an overall view of the hosts and VMs that have the biggest impact on the customer's current environment. Which card should you first examine to see this information?

- A. Summary for VMware vCenters
- B. Simulator for VM Rightsizing
- C. Shared Storage Analysis
- D. Host Analysis

Answer: ([SHOW ANSWER](#))

HPE CloudPhysics is a SaaS-based analytics platform integrated into the HPE GreenLake Data Services Cloud Console (DSCC) that provides deep insights into virtualized infrastructure. It uses a visual metaphor called "cards" to present focused analytics for specific use cases.

When a customer is concerned about storage efficiency and the overall impact of VM workloads, the Simulator for VM Rightsizing (also known as the VM Rightsizing Simulator) is the primary tool for analysis. This card specifically identifies inefficiencies such as over-provisioned or "oversized" virtual machines. By analyzing the actual resource utilization (CPU, memory, and I/O) versus the allocated capacity, the rightsizing simulator can highlight which VMs are consuming excessive resources without operational need.

Reviewing this card first is critical because storage efficiency is often compromised by "virtual sprawl" and bloated VM templates that waste disk space and IOPS on the backend HPE Alletra arrays. While the Shared Storage Analysis card (Option C) provides visibility into datastore contention and performance, the Simulator for VM Rightsizing provides the most direct answer to "efficiency" concerns by identifying the exact VMs that can be trimmed down to reclaim stranded capacity. Correcting VM sizing at the hypervisor level is often the most effective first step in optimizing the performance and capacity return on investment of the underlying storage hardware.

NEW QUESTION: 23

What will occur when a new node is added to an existing HPE Alletra MP X10000 storage array?

- A. An automatic cluster upgrade is supported across all releases of the operating systems and models.
- B. The expanded capacity is immediately available in the shared pool.
- C. Additional drives can be used to increase drive protection beyond the default limit of three drives.
- D. An automatic rebalancing across JBODs occurs as soon as a new drive or an additional JBOD is added.

Answer: D ([LEAVE A REPLY](#))

The HPE Alletra MP X10000 is an object and file storage solution utilizing a Disaggregated Shared- Everything (DASE) architecture. A key differentiator of this disaggregated design is the stateless nature of the controller nodes and the centralized management of the data plane.

When a cluster expansion occurs-such as adding a new controller node or an additional JBOF (Just a Bunch of Flash) storage shelf-the system is designed to automatically optimize the workload distribution.

According to the HPE Alletra MP Architectural Guide, adding an additional JBOF or drives triggers an automatic rebalancing of the data stripes. Unlike older architectures where manual rebalancing services were required (such as in the 3PAR/B10000 block lineage), the X10000 uses a sophisticated hashing mechanism.

Specifically, data is distributed across DSPs (Data Storage Processors) which are virtualized management units. Upon the addition of hardware, these DSPs are rebalanced across the available compute and storage resources in a matter of seconds. Because the nodes are stateless and state is persisted only within the JBOFs, this rebalancing happens with minimal performance impact and no need for the massive "data movement" traditionally associated with expanding a RAID group. This ensures that as a customer scales from the minimum of 3 nodes up to 8 or more, the system always maintains an optimal load balance and utilizes all available flash bandwidth and compute cycles in parallel.

NEW QUESTION: 24

A customer needs to replace their current data protection solution, including hardware and software. They have the following requirements:

- * A single data management platform for data protection of hypervisor, container, cloud, physical, database, and application workloads
- * Eliminate data silos across backups for files, objects, and archiving
- * Needs to support a large, scale-out NAS solution

What is the best solution for this customer?

- A.** HPE GreenLake Flex with Commvault and HPE Alletra 4000 storage servers
- B.** HPE GreenLake Flex with HPE Zerto and HPE StoreOnce appliances
- C.** HPE GreenLake Flex with Cohesity and HPE Alletra 4000 storage servers
- D.** HPE GreenLake Flex with Veeam and HPE Alletra 4000 storage servers

Answer: ([SHOW ANSWER](#))

The customer's requirements focus on a single data management platform that can unify disparate backup tasks and eliminate data silos across files, objects, and archiving while supporting massive scale-out NAS.

The HPE Solutions with Cohesity (specifically Cohesity DataProtect and Cohesity SmartFiles) are architecturally designed to meet these specific needs.

Unlike traditional backup software that often relies on separate components for different data types, Cohesity provides a unique shared-nothing, scale-out architecture that consolidates secondary data onto a single platform. It natively supports a vast array of workloads including virtual machines, containers (Kubernetes), databases (SQL, Oracle, NoSQL), and physical servers. A core differentiator for Cohesity is its ability to act as a Scale-Out NAS via its SmartFiles feature, allowing it to manage PB-scale unstructured data without the performance bottlenecks found in traditional "siloes" storage.

When delivered via HPE GreenLake Flex, this solution is typically paired with HPE Alletra 4000 storage servers (such as the Alletra 4120 or 4140). These servers are density-optimized, storage-centric systems that provide the high-throughput and massive internal capacity required for a modern secondary storage environment. While Commvault (Option A) and Veeam (Option D) are powerful data protection suites, they are often used in conjunction with external target storage (like StoreOnce or Alletra MP) and do not always provide the same level of native, unified scale-out NAS and data silo elimination within a single management plane as the integrated Cohesity/Alletra 4000 stack.

NEW QUESTION: 25

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.** The RabbitMQ load balancer is a required component of the installation.
- C.** It is a simple installation and configuration, wherein Morpheus installs and configures all required services.
- D.** Upgrades require a minimal amount of downtime when updating to a newer version of Morpheus.

Answer: (SHOW ANSWER)

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: 26

A company with 2484 VMs and 300 servers needs to implement a file, object, and block storage solution.

What are the minimum requirements for this solution?

- A. One HPE Alletra MP B10000 and one HPE Alletra MP X10000
- B. One HPE Alletra MP B10000 and two HPE Alletra MP X10000s
- C. Two HPE Alletra MP B10000s and one HPE Alletra MP X10000
- D. Three HPE Alletra MP X10000s

Answer: D (LEAVE A REPLY)

The HPE Alletra MP is a modular, disaggregated storage platform designed to provide different storage personas (Block or File/Object) based on the software stack installed on the controller nodes. However, the minimum hardware "footprint" required to form a functional, supported cluster differs significantly between these personas.

For HPE GreenLake for File Storage (which utilizes the Alletra MP X10000 hardware and provides both File and Object protocols), the architecture is based on a disaggregated shared-everything (DASE) model.

According to the HPE Alletra MP Installation and Architecture Guide, the minimum supported configuration for a File/Object cluster is three X10000 controller nodes. This 3-node minimum is a hard requirement to establish proper quorum and high availability for the V-Tree metadata and the distributed file system logic. A single X10000 node (as suggested in Options A and C) cannot function as a standalone file

/object cluster in a production environment.

Furthermore, the Alletra MP X10000 persona is specifically optimized for high-density unstructured data (File and Object). While the B10000 persona (Options A, B, and C) is intended for Block storage, the question asks for a solution that covers file, object, and block. In many modern software-defined or unified scenarios, especially those aligned with the Alletra MP's future-proof roadmap, the X10000 hardware can serve multiple personas. However, strictly following the current architectural minimums for the File/Object requirement mentioned, you must have at least three nodes. Therefore, a 3-node cluster of X10000s is the foundational requirement to even begin providing the file and object services the customer needs. Options A and B fail the minimum cluster size requirement for the File/Object persona.

NEW QUESTION: 27

A customer currently has an HPE Alletra 9000 with data reduction on all volumes and plans to migrate to an HPE Alletra MP B10000. Which formula should be used to size the new solution?

- A. Size to consumption multiplied by 1.25
- B. Size to consumption multiplied by 1.35

C. Size to consumption multiplied by 1.5

D. Size to original capacity

Answer: (SHOW ANSWER)

When sizing a migration from a highly efficient array like the HPE Alletra 9000 (or Primera) to the next-generation HPE Alletra MP B10000, storage architects must account for the difference between the "Written Capacity" (what the host thinks it has stored) and the "Consumed Capacity" (the physical space used after data reduction).

The standard best practice for an HPE Master ASE when performing these migrations is to Size to consumption multiplied by 1.25. This "1.25 factor" (representing a 25% overhead) is the recommended safety margin used in sizing tools like HPE NinjaStars and the HPE Cloud Physics assessment reports.

This 25% buffer is designed to cover several critical architectural requirements:

- * **System Metadata and Overhead:** Both the Alletra 9000 and Alletra MP require physical capacity to store internal metadata, map tables, and the structures required for their respective data reduction engines.

- * **Snapshot Reserve:** While snapshots are thin and pointer-based, they still consume physical space as data changes over time. The 1.25 multiplier ensures there is enough "headroom" for typical snapshot retention policies.

- * **Data Reduction Parity:** Data reduction ratios (deduplication and compression) can fluctuate based on the specific workload. Sizing exactly to current consumption without a buffer risks an out-of-space condition if the new array's reduction engine handles a specific block pattern slightly differently during the initial ingest.

- * **Operational Performance:** SSD-based arrays perform best when they are not "packed" to 100% capacity, as the garbage collection and wear-leveling processes require free blocks to operate efficiently.

Sizing to "original capacity" (Option D) would lead to a massive over-provisioning and wasted cost, as it ignores the benefits of modern data reduction. Option C (1.5) is generally considered overly conservative for modern flash environments, while 1.25 provides the optimal balance of cost-efficiency and technical risk mitigation.

NEW QUESTION: 28

An administrator needs to create a backup policy for an Oracle database using the RMAN tool. Use your cursor to place a + where the administrator can click to configure this option in StoreOnce.



Answer:



Explanation:

Catalyst Stores

To integrate Oracle RMAN with an HPE StoreOnce appliance, the primary requirement is the creation of a StoreOnce Catalyst Store. Unlike traditional NAS shares or Virtual Tape Libraries (VTL), Catalyst is a proprietary, high-performance protocol designed specifically for deduplication-aware backup workloads.

According to the HPE StoreOnce Catalyst for Oracle RMAN User Guide, the first step in the configuration workflow is to establish a target repository on the appliance. As shown in the exhibit (image_64a41b.png), the administrator has navigated to the Data Services section. To begin the Oracle RMAN integration, they must click on the Catalyst Stores tile (currently showing "0 Store"). Within this menu, they will create a new store and configure the necessary client access permissions, allowing the Oracle database server to communicate with the appliance using the Catalyst protocol.

Once the store is created in the UI, the second phase of the implementation occurs on the Oracle host, where the HPE StoreOnce Catalyst Plug-in for Oracle RMAN must be installed. This plug-in provides the "SBT" (System Backup to Tape) interface that RMAN uses to direct backup streams to the Catalyst Store. By using Catalyst instead of traditional NAS (CIFS/NFS), the customer benefits from source-side deduplication, which significantly reduces network traffic and backup windows by only sending unique data blocks from the database server to the StoreOnce system. Selecting the "NAS Shares" or "VT Libraries" options in the UI would not provide the necessary Catalyst interface required for this optimized Oracle RMAN integration.

Valid HPE7-J01 Dumps shared by Actual4test.com for Helping Passing HPE7-J01 Exam!
Actual4test.com now offer the **newest HPE7-J01 exam dumps**, the Actual4test.com HPE7-J01 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com HPE7-J01 dumps with Test Engine here: https://www.actual4test.com/HPE7-J01_examcollection.html (**62** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)