

HPE Storage Solutions

HP HPE0-J68

Version Demo

Total Demo Questions: 30

Total Premium Questions: 302

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Describe storage architectures, technologies, and solutions	91
Topic 2, Identify and position HPE storage products and solutions	134
Topic 3, Gather and analyze customer requirements	9
Topic 4, Design HPE storage solutions	33
Topic 5, Validate and optimize HPE storage solutions	35
Total	302

QUESTION NO: 1

You have proposed a Primera array to a customer for primary storage. The customer currently has an older HPE 3PAR array and will continue to use it moving forward. The customer would like high availability configured.

Which RMC Peer Copy pre-requisites do you need to point out to the customer? (Choose two.)

- A. 4-port 1 Gb Ethernet Adapter (NIC) is required in each array
- B. Target drive type must match source drive type
- C. 2-port 10 Gb Ethernet Adapter (NIC) is required in each array
- D. Check that the 3PAR OS version is supported
- E. Required license must be provided
- F. RMC 6.2

ANSWER: A E

Explanation:

RMC Peer Copy depends on a supported Remote Copy over IP connectivity design between the HPE Primera and HPE 3PAR systems. A 4-port 1 Gb Ethernet adapter in each array provides the required IP replication connectivity for this Peer Copy configuration. This infrastructure must be planned on both arrays so that the Peer Copy relationship can provide the intended availability and recovery capability.

The required software licensing must also be available. The older HPE 3PAR array must be appropriately licensed for the replication functionality used by the configuration; licensing is therefore an essential item to identify during solution design, before deployment and testing. Confirming both the replication-network hardware and the necessary licenses early avoids deploying an RMC configuration that cannot establish or maintain its protected-copy relationship.

HPE's Primera documentation resources provide platform and software guidance at [HPE Primera documentation](#). HPE's 3PAR documentation catalog is available at [HPE 3PAR documentation](#); use the applicable RMC and Remote Copy compatibility documentation for the exact array software and hardware revisions in the customer environment.

QUESTION NO: 2

You are proposing a new Nimble solution to a customer that recently experienced data loss in their existing storage array.

Which features of the Nimble array should you emphasize to your customer to address their concerns? (Choose two.)

- A. quick RAID rebuild
- B. deduplication
- C. compression
- D. encryption
- E. Triple+ Parity

ANSWER: A E

Explanation:

Quick RAID rebuild and Triple+ Parity directly address a customer's concern about data loss caused by drive failures. HPE Nimble Storage uses an architecture designed to maintain protection while minimizing the exposure window that can occur after a disk failure. Its fast RAID rebuild capability rebuilds only the data that was actually written, rather than rebuilding the full capacity of a replacement disk. This significantly reduces rebuild time and restores the array's intended protection level sooner.

Triple+ Parity provides strong fault tolerance by protecting data against multiple simultaneous drive failures within a RAID set. This level of resilience is particularly relevant when the customer has already experienced a data-loss event and needs confidence that storage media failures will not compromise application data. Together, rapid rebuilds and multi-drive parity protection provide both immediate recovery following a failed drive and robust ongoing availability protection. HPE identifies built-in data resiliency and efficient RAID protection as core capabilities of the Nimble Storage platform. See the [HPE Nimble Storage product page](#) and [HPE Nimble Storage documentation](#).

QUESTION NO: 3

Your customer would like to implement a data protection technology that will offer the lowest possible RPO and will protect from entire array failure. Which technology should the be used?

- A. Snapshots
- B. Disk-based backup
- C. Asynchronous replication
- D. Synchronous replication

ANSWER: D

Explanation:

Synchronous replication is the appropriate technology because it provides a recovery point objective (RPO) of zero, or as close to zero as the application and infrastructure can support. For each host write, synchronous replication requires the data to be committed at both the production array and a remote replica array before the write is acknowledged to the host. Consequently, if the production array suffers a complete failure, the remote array holds a current, consistent copy of the protected data and can be used for recovery or failover with no acknowledged-write data loss.

This protection model is specifically intended for business-critical workloads where even a small interval of lost data is unacceptable. It requires a suitably low-latency, highly available connection between sites, since replication acknowledgement is in the application write path. HPE storage platforms support synchronous remote-copy capabilities to provide disaster recovery across separate arrays and locations. HPE describes Remote Copy as a replication feature for maintaining copies on remote systems, while its storage data-services material identifies synchronous replication as the approach used when zero-data-loss protection is required. See [HPE Primera and Alletra 9000 Remote Copy documentation](#) and [HPE storage data services](#).

QUESTION NO: 4

A customer is looking for a centralized file-sharing solution that can provide print services in addition to 15TB of storage.

Which solution component should you recommend?

- A. FC switch
- B. NAS
- C. FC-based array
- D. Disk enclosure with integrated SAS switch

ANSWER: B

Explanation:

NAS is the appropriate solution component because Network Attached Storage is designed to provide centralized, network-accessible file storage to multiple users and systems. It presents storage through file-sharing protocols such as SMB/CIFS and NFS, making it well suited to a customer requirement centered on shared files rather than block storage for individual servers. A 15 TB capacity requirement is readily addressed by an appropriately sized NAS platform and can generally be expanded as requirements grow.

In an HPE environment, HPE StoreEasy is an example of a NAS-oriented solution built for simple, efficient file storage. Its Windows Server foundation also enables organizations to use Windows file-server capabilities alongside print-server services, allowing centralized management of shared folders and network printers. This aligns directly with the stated need for both file sharing and print services in one solution. HPE describes StoreEasy as purpose-built for file and application storage, while Microsoft documents the Print Server role used to share and manage network printers. See [HPE StoreEasy](#) and [Microsoft Print Server documentation](#).

QUESTION NO: 5

A customer is designing iSCSI connectivity from VMware hosts to an HPE MSA array. They require resilient storage access and want to avoid contention with normal user LAN traffic.

Which design practices should you recommend? (Choose two.)

- A. Dedicated network paths for iSCSI traffic
- B. MPIO on each host
- C. One iSCSI path per host
- D. Use the production user LAN as the only storage path
- E. Disable redundant controller ports

ANSWER: A B

Explanation:

Dedicated network paths for iSCSI traffic should be used to separate storage I/O from general user and application LAN traffic. This can be implemented with dedicated switches or appropriately isolated Ethernet networks and helps provide more predictable storage performance.

MPIO on each host is also required for a resilient iSCSI design. Multipathing combines multiple physical paths to the same MSA volume, presents the volume as one device to the operating system, and provides path failover if a NIC, cable, switch, controller port, or controller path fails. HPE MSA product and management information is available at [HPE MSA Storage](#).

QUESTION NO: 6 - (SIMULATION)

Match the data access type in a SAN fabric to the correct description.

ANSWER: See the explanation for the answer

Explanation:

Correct matches:

Block access — The storage array presents a LUN/volume as a set of raw logical blocks. The host operating system or application creates and manages the file system on that LUN. This is the standard access type for a SAN (for example, Fibre Channel or iSCSI).

File access — Data is accessed as named files and directories through a file-sharing protocol such as NFS or SMB. The NAS system, rather than the host, owns and manages the file system.

Object access — Data is stored and retrieved as objects, each with an object identifier and metadata, typically through an HTTP-based API. It is not accessed as a mounted block LUN or conventional shared file system.

Therefore, for a SAN fabric, select **block access**: hosts access storage through LUNs and manage the file system themselves.

QUESTION NO: 7

What is the benefits of HPE synergy(Select 2)

- A. It can present block storage services from HPE 3par as part of the resource pool
- B. It support up to 3 HPE store Virtual VSA nodes for storage RAID
- C. It can use HPE 3PAR file persona for objet storage
- D. It can compose and reclaim DAS storage
- E. It support HPE blade system C-Class Flex Fabrics Interconnect module (need to check)

ANSWER: D

Explanation:

“It can compose and reclaim DAS storage” is a core HPE Synergy benefit. HPE Synergy is designed as composable infrastructure, meaning that physical infrastructure resources can be programmatically assigned to a workload through templates and profiles, then returned to the available resource pool when that workload no longer requires them. With supported direct-attached storage, this enables administrators to compose storage capacity for a server configuration and reclaim that capacity during decommissioning or reprovisioning. The result is faster deployment, more consistent configuration, and improved utilization of locally attached storage resources without requiring each server to remain permanently tied to a fixed storage configuration. This capability is particularly valuable for environments where application requirements change frequently and infrastructure must be repurposed quickly. HPE describes Synergy as a platform for fluid pools of compute, storage, and fabric resources, managed through software-defined templates and APIs. The Synergy D3940 storage module provides the direct-attached storage capacity used in these composable configurations. See the [HPE Synergy overview](#) and the [HPE Synergy D3940 Storage Module product information](#).

QUESTION NO: 8

You are required to provide key sizing values for a customer solution for HCI sizing in CloudPhysics. Which HPE tool can you utilize to create the required output?

- A. SAF
- B. HPE SimpliVity Sizing Tool
- C. OCA
- D. Ninja Stars

ANSWER: B

Explanation:

HPE SimpliVity Sizing Tool is the appropriate tool because it is designed to translate customer workload and infrastructure requirements into the sizing information needed to design an HPE SimpliVity hyperconverged solution. Its output provides the key capacity, performance, resiliency, and node-sizing values that can be used as the basis for an HCI sizing exercise in CloudPhysics. This creates a repeatable sizing workflow: gather the customer’s workload characteristics, model the required SimpliVity configuration, and use the resulting sizing data when evaluating the proposed HCI solution. The tool therefore supports solution architects in moving from assessment data to a defensible HPE SimpliVity design rather than relying on informal estimates. HPE positions SimpliVity as an integrated hyperconverged infrastructure platform that combines compute, storage, and data services, making accurate workload-driven sizing essential to the proposed configuration. See HPE’s [HPE SimpliVity solution overview](#) and the [HPE hyperconverged infrastructure portfolio](#) for product and solution context.

QUESTION NO: 9

A customer has deployed an HPE 3PAR StoreServ 8400 with two nodes and three additional disk enclosure .They want to expand the solution with two additional nodes and three

additional disk enclosure ..What must you do to expand the customers solution?

- A. The new disks must be imported before the new nodes can be added
- B. The existing nodes need to be take Offline
- C. The system should be updated to the latest supported OS before the hardware expansion
- D. The new nodes and enclosure must be placed in an adjacent racks

ANSWER: C

Explanation:

The system should be updated to the latest supported OS before the hardware expansion. A two-node-to-four-node expansion on an HPE 3PAR StoreServ 8400 is a controller-node upgrade procedure, and the installed HPE 3PAR OS level must support both the existing array configuration and the additional controller nodes, drive enclosures, and their associated firmware. Confirming and installing the required supported HPE 3PAR OS release before starting the physical upgrade establishes the compatibility baseline required for the expansion. It also ensures that the system has the code functions, service procedures, and firmware interoperability necessary to recognize and configure the new hardware correctly after installation. The precise target release should be selected using HPE's current compatibility and release guidance for the particular node and enclosure hardware, rather than assuming that any arbitrary newer release is appropriate. The expansion should then be performed under the applicable HPE installation and upgrade procedures, including the required health checks and configuration validation. HPE publishes product documentation and software release material through its [Support Center](#); platform lifecycle and support information is also available from the [HPE 3PAR storage page](#).

QUESTION NO: 10

Which software feature should an HPE 3PAR Storeserv 8000 customer use to optimize data protection and recovery, Including rapid online recovery for his VMware environment?(select two)

- A. HPE 3PAR Virtual Copy software
- B. HPE 3PAR Remote copy software
- C. HPE 3PAR Security software suite
- D. HPE 3PAR File persona Software suite
- E. HPE StoreOnce recovery manager central for VMware

ANSWER: A E

Explanation:

HPE 3PAR Virtual Copy software is appropriate because it creates highly space-efficient, pointer-based virtual copies of volumes. These snapshots provide a rapid recovery point on the array, allowing data to be restored or made available quickly without requiring a full-volume physical copy. In a VMware environment, this capability supports fast operational recovery of virtual-machine data and can be integrated into storage-aware protection workflows. HPE describes Virtual Copy as a 3PAR data-protection capability for creating space-efficient snapshots and clones; see [HPE 3PAR Virtual Copy Software documentation](#).

HPE StoreOnce recovery manager central for VMware is also designed specifically to improve VMware data protection when used with 3PAR and StoreOnce. Recovery Manager Central orchestrates VMware-aware, application-consistent snapshot protection and moves protected data efficiently to StoreOnce backup storage. This combination reduces backup impact on production resources while retaining the rapid online recovery benefit of 3PAR snapshots, with StoreOnce providing efficient longer-term retention and recovery copies. HPE's StoreOnce portfolio information describes its purpose as efficient backup and recovery storage, including integration with HPE data-protection workflows: [HPE StoreOnce Backup and Recovery](#). Together, these capabilities address both fast local recovery and optimized backup-based protection for VMware workloads.

QUESTION NO: 11

You need to deploy a Recovery Manager Central appliance and integrate it with a customer's 3PAR system.

Which information must be collected from the customer in order to proceed with the installer? (Choose two.)

- A. StoreOnce store name
- B. vCenter IP address or hostname
- C. HPE 3PAR used capacity
- D. IP address of the RMC
- E. NFS path

ANSWER: D

Explanation:

The IP address of the RMC must be collected because deployment of the Recovery Manager Central appliance requires its network identity to be defined before the appliance can communicate with the surrounding infrastructure. The installer uses the appliance network configuration to establish management connectivity, allowing administrators and integrated systems to reach the RMC service reliably after deployment. In practice, the deployment team should obtain a reserved, routable address together with the related network parameters, such as subnet mask, default gateway, DNS configuration, and hostname, so the appliance is placed correctly on the customer network and does not conflict with another device. This network information is foundational to subsequent configuration of storage-system access and data-protection workflows. HPE describes Recovery Manager Central as software that integrates storage-based copy services with application and virtualization environments, making dependable appliance connectivity an essential deployment prerequisite. See HPE's [data protection portfolio information](#) and the [HPE 3PAR StoreServ documentation portal](#) for product documentation and deployment resources.

QUESTION NO: 12

A customer has several virtual machines distributed across two HPE SimpliVity clusters. They need to protect selected VMs at the primary site and retain backup copies at the remote site.

Which SimpliVity capabilities should you recommend? (Choose two.)

- A. Policy-based VM backups
- B. Remote backup to a SimpliVity cluster
- C. Fibre Channel LUN masking
- D. Storage tier affinity rules
- E. Virtual tape library emulation

ANSWER: A B

Explanation:

Policy-based VM backups are appropriate because HPE SimpliVity protection policies can be assigned to individual virtual machines. The policy defines the backup frequency and retention settings, allowing administrators to apply consistent VM-centric protection without separately configuring storage snapshots and backup jobs.

Remote backup to a SimpliVity cluster is also correct. SimpliVity can send deduplicated backup copies from a local cluster to another SimpliVity cluster, supporting disaster-recovery protection at a remote location. The remote cluster stores the backup data and can be used to restore or recover protected virtual machines if the primary site is unavailable. HPE describes SimpliVity integrated data protection capabilities at [HPE SimpliVity](#).

QUESTION NO: 13

What is the purpose of WWN in FC SAN? (Choose two.)

- A. Identification of zone members
- B. Identification of entities denied access to a volume
- C. Identification of entities via Storage Name Server (SNS)
- D. Load balanced frame delivery
- E. Forward translation of DNS entries

ANSWER: A C

Explanation:

A World Wide Name (WWN) is a persistent, globally unique identifier assigned to a Fibre Channel node or port. In an FC SAN, port WWNs are commonly used to identify the host bus adapter ports and storage target ports that are permitted to communicate. Therefore, WWNs provide the **identification of zone members**: an administrator can create a zone containing the specific initiator and target port WWNs that require connectivity. WWN-based zoning remains consistent even when a device is connected through a different physical switch port, which makes it a common and portable zoning practice.

WWNs are also used for the **identification of entities via Storage Name Server (SNS)**. The FC name server maintains registration and discovery information for devices in the fabric. Devices register attributes such as their port WWN, node WWN, FC address, and supported capabilities. Other fabric participants can query this information to discover and identify eligible FC devices. This identity-based discovery is fundamental to SAN device visibility and subsequent zoning or storage-presentation workflows. See the [Fibre Channel Industry Association technical information](#) and the [Cisco Fibre Channel configuration documentation](#) for Fibre Channel naming and fabric-service concepts.

QUESTION NO: 14

You have just completed a performance analysis of an HPE 3PAR StoreServ 8400 disk array with different tiers at a customer's site. After reviewing the performance metrics, you determine that they should install a new tier of SSD drives.

Which actions should you recommend they take to fully optimize the array? (Choose two.)

- A. Configure Adaptive Flash Cache.
- B. Configure scheduled tunesys tasks.
- C. Rebalance all flash drives between nodes 0 and 2.
- D. Configure Adaptive Optimization.
- E. Configure dedicated drives to enable Adaptive Flash Cache.

ANSWER: B D

Explanation:

Configuring Adaptive Optimization is appropriate after adding an SSD tier because Adaptive Optimization uses tiering policies and performance measurements to relocate regions of virtual volumes to the most suitable Common Provisioning Group tier. This allows frequently accessed, latency-sensitive data to benefit from the SSD tier while less active data can remain on lower-cost tiers. The SSD capacity must be incorporated into the relevant tiering configuration so that Adaptive Optimization can use it as a destination tier.

Configuring scheduled tunesys tasks is also appropriate. The `tunesys` operation balances chunklets and system resources after capacity changes, helping distribute the newly added physical drives effectively across the array. Scheduling the operation allows the array to maintain a balanced layout as its configuration and workload evolve. Together, tier-aware data

movement through Adaptive Optimization and layout balancing through scheduled `tunesys` operations help the 3PAR StoreServ system realize the performance benefit of the additional SSD tier.

For HPE storage product documentation and software guidance, consult the [HPE Support Center](#) and the [HPE Storage documentation portal](#).

QUESTION NO: 15

Which actions can be done within InfoSight Cross-Stack Analytics for VMware and Hyper-V? (Choose two.)

- A. View replication policies
- B. Find noisy neighbors
- C. Locate inactive VMs
- D. Modify administrative privileges
- E. Review backup schedules

ANSWER: B C

Explanation:

HPE InfoSight Cross-Stack Analytics correlates infrastructure telemetry across the virtual-machine, hypervisor, compute, network, and storage layers. This shared, VM-centric visibility lets an administrator **find noisy neighbors** by identifying virtual machines that contend for shared resources or create disproportionate demand and then viewing the resulting performance impact across the stack. The capability is intended to reduce troubleshooting time by connecting VM behavior to the underlying infrastructure rather than requiring separate manual investigations in each management tool.

Cross-Stack Analytics can also **locate inactive VMs**. Identifying VMs with little or no activity helps administrators recognize candidates for cleanup, reclamation, or further review, improving resource efficiency and avoiding unnecessary consumption of virtualized infrastructure capacity. These analytics use the relationship between VM-level activity and the supporting HPE storage and compute environment to provide actionable operational insight. HPE describes the Hyper-V implementation as providing end-to-end visibility and analysis from the VM through the infrastructure stack. See HPE's [Cross-Stack Analytics for Hyper-V overview](#) and the [HPE InfoSight product page](#).

QUESTION NO: 16

Select a correct statement regarding the expandability of the HPE Alletra 9080 with 4 controllers and maximum capacity requirements.

- A. Additional drive enclosure increases the performance of the solution.
- B. Eight drive enclosures increase capacity efficiency with 82 RAID 6 set.
- C. HPE Alletra 9000 supports drive enclosure daisy chaining
- D. Sixteen drive enclosures increase capacity efficiency with 82 RAID 6 set.

ANSWER: D

Explanation:

Sixteen drive enclosures increase capacity efficiency with 82 RAID 6 set. For a four-controller HPE Alletra 9080 configuration built to its maximum drive-enclosure capacity, the appropriate RAID-6 geometry is an 8+2 layout (written as "82" in the option). This geometry places eight data chunklets and two parity chunklets in each RAID set. It provides RAID-6 dual-parity protection while using 80% of the set's raw capacity for data, which is more capacity-efficient than a 6+2 RAID-6 layout. The four-controller architecture and sixteen-enclosure scale provide sufficient drive population and controller resources to use this wider RAID-set geometry effectively, while maintaining the platform's distributed storage design and resilience characteristics. This capacity-planning guidance is relevant when the design objective is maximum usable capacity rather

than simply adding enclosures for a smaller configuration. HPE identifies the Alletra 9000 family as a mission-critical, scale-out storage platform; the detailed supported configurations and sizing guidance are maintained in its product documentation and support resources. See the [HPE Alletra 9000 product page](#) and the [HPE Support Center](#) for current QuickSpecs and configuration documentation.

QUESTION NO: 17

Which statement is correct when comparing functionalities of the AF,CF, and SF series in the nimble storage portfolio?

- A. Zero copy clone are unique feature of all flash array within the nimble portfolio
- B. SF series supports storage snapshots and replication using veeeam backup & replication CS and AF series support storage snapshot only.
- C. Inline deduplication across the complete array capacity, in combination with compression, is offered by the AF and SF series only.
- D. 99.9999% availability measured and guaranteed for AF and CF series arrays, SF series array offer 99.999% availability.

ANSWER: C

Explanation:

Inline deduplication across the complete array capacity, combined with compression, is offered only by the AF and SF series. In the legacy HPE Nimble Storage portfolio, the AF (All Flash) platform and the SF (Secondary Flash) platform provide always-on inline data reduction capabilities that include variable-block deduplication and compression. Applying these functions inline means that duplicate blocks can be eliminated and data can be compressed before it consumes capacity on the array. This improves effective usable capacity without requiring a separate post-process deduplication activity. Complete-array-capacity support is important because the data-reduction benefit is not confined to a small cache tier or a limited portion of the installed flash media. AF systems apply this capability while serving performance-oriented primary workloads, whereas SF systems use it to deliver capacity-efficient flash storage for secondary use cases. HPE describes Nimble Storage data services and efficiency capabilities as part of its storage platform information at [HPE Nimble Storage](#). Product documentation and technical resources for the Nimble portfolio are also available through the [HPE Support Center](#).

QUESTION NO: 18

What is the maximum number of host-facing PCIe expansion slots available in a HPE Alletra 6030 system?

- A. 2
- B. 4
- C. 3
- D. 6

ANSWER: B

Explanation:

The maximum is 4 host-facing PCIe expansion slots. The HPE Alletra 6030 is a dual-controller storage array, and its architecture provides two host-facing PCIe expansion slots in each controller. Across the complete high-availability system, this produces a maximum of four available slots. These slots support installation of supported host-connectivity adapters, allowing the array's front-end connectivity to be selected for the SAN or Ethernet environment in which it is deployed. Counting the slots at the system level, rather than only on one controller, is essential for this specification because both controllers participate in serving host I/O and maintaining controller redundancy. HPE identifies the Alletra 6000 family as an all-flash, business-critical storage platform with flexible connectivity intended for block-storage deployments. Product

specifications and current platform documentation are available from HPE's [Alletra 6000 storage page](#) and the [HPE Support Center](#), where the applicable QuickSpecs and installation documentation can be located by model.

QUESTION NO: 19

Which HPE storage solutions offer a two-tier backup solution for hybrid cloud with optimal recovery time? (Choose two.)

- A. Simplivity 380
- B. Nimble SF series
- C. 3PAR 8000 series
- D. StoreOnce with Catalyst
- E. Nimble CS series

ANSWER: C D

Explanation:

3PAR 8000 series and **StoreOnce with Catalyst** provide the complementary storage and data-protection capabilities used in a two-tier hybrid-cloud backup design. 3PAR StoreServ 8000 provides the production-storage tier, with application-consistent snapshots and Remote Copy replication features that enable rapid local recovery and support tight recovery-time objectives. Its integration with HPE data-protection software and StoreOnce allows protected data to be moved beyond the primary array while retaining efficient recovery options.

StoreOnce with Catalyst provides the backup tier. Catalyst enables backup applications and HPE storage platforms to write optimized, deduplicated backup data directly to StoreOnce. StoreOnce can retain recent recovery points locally for fast restores and extend protected copies to cloud object storage through Cloud Bank Storage, creating the hybrid-cloud tier without requiring all recoveries to originate from the cloud. This combination keeps frequently needed recovery data close to workloads while using cloud capacity for scalable off-site retention.

HPE describes StoreOnce as a deduplicating backup appliance with cloud integration and Catalyst-based backup workflows, while its 3PAR documentation details snapshot and replication data-protection capabilities. See [HPE StoreOnce](#) and [HPE 3PAR StoreServ Storage Concepts Guide](#).

QUESTION NO: 20

You are designing an HPE MSA iSCSI for a customer.

Which components are typically part of the solution? (Choose two.)

- A. Drives
- B. SFPs
- C. Service Processor
- D. All-inclusive Multi-system software
- E. Switch zones

ANSWER: A B

Explanation:

Drives are a fundamental part of an HPE MSA iSCSI storage solution because the MSA enclosure provides shared block storage by housing and managing the physical SSDs and/or HDDs that supply capacity and performance. Drive selection is part of solution design, including capacity, performance tiering, resiliency, and the required drive carriers for the selected

MSA platform. HPE describes MSA systems as entry storage arrays designed to use flash and disk media for shared storage workloads.

SFPs are also typically included when the iSCSI host connectivity uses optical Ethernet. The storage controllers' host ports require compatible transceivers to connect through fiber-optic Ethernet cabling to the customer network switches. The precise SFP type must match the controller port speed and media requirements, such as an appropriate short-range optical transceiver. Therefore, a complete design normally identifies the needed SFPs along with the selected iSCSI controllers, host network infrastructure, and cabling. HPE MSA documentation and product information provide supported connectivity and hardware configuration details: [HPE MSA Storage](#) and [HPE MSA Storage Management Guide](#).

QUESTION NO: 21 - (HOTSPOT)

HOTSPOT

You are handing over the installed array to customer. The customer wants to review the amount of HDDs and SSDs in the system.

Click the relevant menu option where the front panel of the head shelf can be visualized.

Hot Area:



ANSWER:



Explanation:

The correct selection is the **Hardware** menu in the top navigation bar. The customer is asking for a physical inventory review: specifically, the number and type of installed HDDs and SSDs. That information is obtained from the array hardware view rather than from the performance dashboard or an object in the group tree. Opening Hardware provides access to the installed storage shelves and the graphical shelf representation, including the front-panel layout and populated drive bays. From that shelf view, an administrator can visually identify the drives installed in the head shelf and review their media types, allowing the customer to confirm the HDD and SSD configuration during handover.

This is the appropriate operational path because the front-panel display represents the physical enclosure and its individual drive locations. It is designed for hardware inventory, component identification, and validation of the array's installed configuration. Performance pages show operational statistics such as latency, IOPS, throughput, and capacity-related information, but they do not provide the front-facing shelf visualization requested here. Selecting Hardware first leads to the shelf and disk details needed for a clear, customer-facing review of the installed media. HPE's Nimble Storage documentation portal is available at [HPE Nimble Storage documentation](#), which includes the relevant administration and hardware-management guidance.

QUESTION NO: 22

While demonstrating an HPE 3PAR StoreServ 9450 array, the customer asks you to demonstrate the File Persona features.

Which objects do you need to create to demonstrate these features? (Choose two.)

- A. File share
- B. Storage container
- C. Host
- D. Remote copy group
- E. Virtual file server

ANSWER: A E

Explanation:

To demonstrate HPE 3PAR File Persona, create a virtual file server and a file share. A virtual file server is the logical file-serving entity that provides the NAS identity and services used by clients, including network configuration and protocol access. It is the object through which File Persona presents file services such as SMB/CIFS and NFS. After the virtual file server is available, a file share can be created and associated with the relevant file store and protocol configuration. The share is the client-facing export that users access over the network, making it the most direct object for demonstrating file-level access, permissions, and shared-data behavior. Together, these objects allow a demonstration to show both the File Persona service endpoint and actual client consumption of file storage on the 3PAR array. HPE documents virtual file servers and file shares as central File Persona configuration and administration objects in its File Persona guidance. See the [HPE 3PAR File Persona User Guide](#) and the [HPE Storage documentation portal](#) for File Persona configuration documentation.

QUESTION NO: 23

A customer has a large amount of the data on their Primera that is mission critical. The business needs RPO to be under 5 minutes.

Which protection technology should you recommend?

- A. synchronous replication
- B. asynchronous replication
- C. archiving
- D. backup to tape

ANSWER: A

Explanation:

synchronous replication is the appropriate recommendation because the requirement is for mission-critical data with a recovery point objective of less than five minutes. HPE Primera Remote Copy synchronous mode commits a host write only after the write has been acknowledged by both the primary and remote storage systems. Consequently, each acknowledged write exists at both sites, providing a zero-data-loss recovery point objective during a planned or unplanned failover, assuming the replication link and remote system are available.

This behavior is specifically suited to applications for which data currency and transactional consistency take priority, such as critical databases and core business services. The application should be designed with the latency implications of synchronous remote writes in mind, and the two sites must have a network connection capable of supporting the required response time and bandwidth. Nevertheless, where the stated business objective is an RPO below five minutes, synchronous replication provides the necessary continuously current remote copy rather than relying on scheduled or periodic updates.

HPE documents Remote Copy synchronous mode as a disaster-recovery configuration that writes data to both systems before acknowledging completion. See the [HPE Primera OS Remote Copy User Guide](#) and the [HPE Primera Storage documentation](#).

QUESTION NO: 24 - (DRAG DROP)

DRAG DROP

Select the correct initiation option for each HPE Primera function.

Select and Place:

Answer Area

Initiation Option	Function
User Initiated	Stripe data across drives
Autonomic	Create Virtual Volumes
	Create Common Provisioning Groups
	Create RAID-6 Virtual Volumes
	Optional Performance Management
	Export LUNs

ANSWER:

Answer Area

User Initiated	Autonomic	Stripe data across drives
Autonomic	User Initiated	Create Virtual Volumes
	User Initiated	Create Common Provisioning Groups
	User Initiated	Create RAID-6 Virtual Volumes
	Autonomic	Optional Performance Management
	User Initiated	Export LUNs

Explanation:

HPE Primera combines administrator-controlled provisioning with autonomous system services. The correct initiation mapping is: **Stripe data across drives — Automatic; Create Virtual Volumes — User Initiated; Create Common Provisioning Groups — User Initiated; Create RAID-6 Virtual Volumes — User Initiated; Optional Performance Management — Automatic; and Export LUNs — User Initiated.**

The automatic entries represent behavior that is built into the array's architectural and operational intelligence. Primera uses wide striping across available physical drives so that data placement is distributed at the system level rather than requiring an administrator to manually select individual disks for each volume. This approach supports balanced use of storage resources and is part of the platform's autonomous operation. HPE describes Primera as delivering an AI-driven, self-managed experience; see the [HPE Primera overview](#). Its performance capabilities similarly use system intelligence to manage and optimize performance behavior once that optional capability is available and enabled, rather than making the administrator manually perform ongoing placement actions.

The user-initiated entries are explicit storage-administration actions. An administrator defines a Common Provisioning Group because it establishes the storage characteristics and capacity resources that will be used. The administrator then creates virtual volumes, choosing the intended provisioning and protection characteristics, including RAID-6 where that is required. Finally, exporting a LUN is a deliberate host-presentation operation: the administrator presents the selected virtual volume to the appropriate host or host set with the required access controls. These actions express application requirements and host connectivity intent, so they are initiated through management configuration rather than being created autonomously. HPE's [Primera OS documentation](#) describes the platform's virtual-volume, provisioning, and host-presentation administration model.

QUESTION NO: 25

A customer needs to add another HPE 3PAR StoreServ system and plans to use the federated storage approach.

Which advantages can the customer expect? (Choose two.)

- A. load balancing
- B. reduced power, cooling, and environmental impacts
- C. allows multiple management tools
- D. easier migration
- E. federated deduplication

ANSWER: A D

Explanation:

HPE 3PAR StoreServ federation is designed to let independently operating arrays work together while retaining the ability to move workloads between systems without disrupting host access. **load balancing** is an expected advantage because federation enables workloads and virtual volumes to be redistributed across participating arrays. This supports more effective use of the performance and capacity available across the storage estate, rather than requiring a workload to remain permanently tied to the array on which it was originally provisioned.

easier migration is also a core federation benefit. HPE 3PAR Peer Motion technology enables online, non-disruptive movement of virtual volumes between compatible 3PAR StoreServ systems. The host can continue accessing the data during the movement, which helps customers perform technology refreshes, array expansions, and workload rebalancing with substantially less application interruption. Federation therefore provides operational flexibility as storage requirements change, while preserving the enterprise availability characteristics expected of shared storage. HPE describes 3PAR storage capabilities at [HPE 3PAR StoreServ Storage](#); relevant official documentation can also be located through the [HPE Support search for 3PAR Peer Motion](#).

QUESTION NO: 26

Your customer has very low deduplication ratio of their Linux VMs farm which is hosted on a HPE Alletra 9000 Which filesystem type could cause this?

- A. XFS
- B. EXT4
- C. Btrfs
- D. EXT3

ANSWER: A

Explanation:

XFS is the filesystem type that can contribute to a very low deduplication ratio for a Linux virtual-machine farm on HPE Alletra 9000. Array deduplication operates on repeated block patterns, rather than on the logical identity of files. Although many virtual machines may contain similar operating-system files, the way those filesystems allocate blocks, distribute metadata, and lay out data inside virtual disks affects how much identical block content reaches the storage array. XFS is designed primarily for scalable, high-performance filesystem operation and uses allocation groups and allocation behavior that can reduce block-pattern alignment between otherwise similar VM images. Consequently, data that appears similar at the guest level may not present enough matching blocks for high array-side deduplication savings. This is particularly relevant when virtual disks are independently created, patched, expanded, or maintained over time, because their internal layouts diverge. HPE documents that data-reduction effectiveness depends on the characteristics and similarity of the data presented to the array; filesystem and VM layout are therefore material considerations. See the [HPE Alletra 9000 product information](#) and the [Linux kernel XFS documentation](#).

QUESTION NO: 27

Which HPE 3PAR StoreServ feature allows a host to continue I/O uninterrupted during a firmware upgrade?

- A. Peer Persistence
- B. Persistent Ports
- C. Persistent Memory
- D. Persistent Cache

ANSWER: B

Explanation:

Persistent Ports is the HPE 3PAR StoreServ capability designed to preserve host connectivity and enable uninterrupted I/O when a storage node is taken offline, including during nondisruptive firmware upgrades. Normally, a host-facing port is associated with a particular controller node. With Persistent Ports enabled, the surviving partner node can assume the identity of the affected target port, including its Fibre Channel WWPN or iSCSI identity. From the host's perspective, the target remains available at the same address, so the host can continue sending I/O through its configured path without requiring a fabric change, host reconfiguration, or an application outage.

This functionality is especially valuable in highly available SAN environments because it works with multipathing while reducing the disruption associated with controller maintenance. It helps maintain continuous access as the system performs a rolling upgrade: one node can be upgraded or restarted while its partner temporarily services I/O using the persistent port identity, after which service can be returned appropriately. HPE documents Persistent Ports as a feature for transparent host-path failover during node failure and planned maintenance. See the [HPE 3PAR OS 3.3.1 Concepts Guide](#) and HPE's [HPE Alletra 9000 and Primera OS Concepts Guide](#) for the Persistent Ports architecture and host-connectivity behavior.

QUESTION NO: 28

Which StoreOnce Gen4 features are included with the system software without additional licenses? (Choose two.)

- A. Encryption
- B. Cloud Bank
- C. Catalyst
- D. VTL/NAS replication
- E. Flexible I/O HBAs

ANSWER: C D

Explanation:

Catalyst and **VTL/NAS replication** are included in the StoreOnce Gen4 system-software entitlement. Catalyst is StoreOnce's application-aware data-protection interface, allowing supported backup applications to use the StoreOnce appliance's deduplication capabilities efficiently and to manage backup data through Catalyst stores. It is a core capability of the StoreOnce platform rather than an add-on required merely to begin using Catalyst-enabled backup workflows.

VTL/NAS replication is likewise supplied as part of the system software. It enables replication of backup data held in virtual tape library and NAS shares between StoreOnce systems, using StoreOnce deduplication to reduce WAN bandwidth and target capacity requirements. This capability supports remote-copy and disaster-recovery designs without a separate feature license for the basic VTL/NAS replication function.

These inclusions are reflected in HPE's StoreOnce Gen4 software entitlement information and align with the StoreOnce product architecture, which provides Catalyst and replication as foundational backup and recovery capabilities. See HPE's [StoreOnce Gen4 system software entitlement page](#) and the [HPE StoreOnce product page](#).

QUESTION NO: 29

A customer needs to upgrade their HPE storage environment, including driver and firmware levels and OS, to the latest version.

Where can you find the needed information about supported versions?

- A. HPE Storage Assessment Foundry
- B. HPE SAN Design Reference Guide
- C. HPE OneView
- D. HPE SPOCK

ANSWER: D

Explanation:

HPE SPOCK is the correct source because it is HPE's compatibility and interoperability knowledge base for supported storage configurations. When planning an upgrade, it provides the support matrices needed to validate compatible combinations of storage-system software, host operating systems, multipathing software, host bus adapters, HBA drivers, HBA firmware, SAN switches, and switch firmware. This allows the customer to identify versions that are not merely current, but are explicitly supported together for the applicable HPE storage platform and connectivity design.

Upgrade planning must use the relevant SPOCK array and host-connectivity information to confirm the exact model, protocol, operating-system release, and driver or firmware combination. This validation is important because a newer component release is not automatically qualified with every array software version or host environment. SPOCK is continuously maintained and should be consulted alongside the product-specific release notes and upgrade documentation before changes are performed. Access the HPE compatibility database at [HPE SPOCK](#). HPE also provides storage documentation, release notes, and product guidance through the [HPE Support Center](#).

QUESTION NO: 30 - (SIMULATION)

Match the feature to the appropriate C-Series DCNM license type.

Answer Area

SAN Essentials Edition

Advanced Edition

Bundled at no charge with the C-series switches

Provides configuration, real-time monitoring, and troubleshooting capabilities

Switch/server-based license

Historical performance monitoring for network traffic hotspot analysis

ANSWER: See the explanation for the answer

Explanation:

Match the features as follows:

The Essentials edition is the included DCNM functionality for basic SAN management and real-time operations. Advanced licensing adds licensed switch/server functionality and historical performance analysis.