Download Pure Storage FlashArray-Storage-Professional Exam Dumps Free

Shared by Davidson on 17-06-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

An X20R4 array containing 10 x 4.5TB DirectFlash Modules is running out of capacity. The customer found a data pack scheduled for a FlashArray//C array and has inserted it into the array. The customer is unable to admit the new capacity.

What is a possible reason for this?

Options:

A- The new capacity is SAS, which is NOT compatible with an X20R4.
B- The new capacity is QLC, which is NOT compatible with an X20R4.
C- The new capacity is TLC, which is NOT compatible with an X20R4.

Answer:

B

Explanation:

Hardware Architecture (X vs. C): Pure Storage maintains two primary FlashArray lines: the FlashArray//X (performance-oriented) and the FlashArray//C (capacity-oriented).

Flash Types (TLC vs. QLC):

FlashArray//X (like the X20R4 mentioned in the question) uses TLC (Triple-Level Cell) DirectFlash Modules (DFMs). TLC provides high performance and high endurance, which is necessary for latency-sensitive mission-critical workloads.

FlashArray//C uses QLC (Quad-Level Cell) DirectFlash Modules. QLC provides significantly higher density at a lower cost per GB, but it has different performance and endurance profiles compared to TLC.

Compatibility Constraints: Purity//FA is designed to manage specific flash geometries. QLC modules are not compatible with the //X series arrays. The controller logic and software-defined flash management in an X20R4 are tuned for the voltage and timing characteristics of TLC flash.

The Admission Process: When a new data pack is inserted, the array performs a 'handshake.' If the controller detects a module type that it is not hardware-qualified to support (in this case, QLC in an //X chassis), it will refuse to admit the capacity to prevent system instability or data integrity issues.

Why Option A is incorrect: Modern FlashArrays (since the //M series) use NVMe over a PCIe

backplane for DirectFlash Modules. Pure moved away from SAS (Serial Attached SCSI) for its primary data drives years ago to eliminate the performance bottlenecks associated with the SAS protocol.

Why Option C is incorrect: An X20R4 uses TLC flash. If the data pack were TLC, it would likely be compatible (provided it met the minimum module count and Purity version requirements).

Question 2

Question Type: MultipleChoice

A FlashArray administrator is configuring new hosts. There is an option in the personality settings for the target OS.

When is the best time to configure the personality for a host in Purity?

Options:

A- When a host is initially created and before volumes are connected to the host.
B- Host personalities can be configured at anytime except for the ESXi operating system.
C- After the host has been created and volumes are connected to the host.

Answer:

A

Explanation:

Definition of Host Personality: In Purity//FA, a Host Personality is a setting applied to a host object that modifies how the FlashArray communicates with that specific initiator. It ensures the array sends the correct SCSI responses that the target Operating System (OS) expects. Common personalities include ESXi, AIX, HP-UX, and Hitachi-VSP.

The Importance of Timing: The best practice is to set the personality during the host creation phase, before any volumes are attached or I/O has commenced. This ensures that from the very first 'Inquiry' command sent by the host, the FlashArray responds with the appropriate settings (such as specific VAAI primitives for ESXi or specific ALUA behaviors for other Unix variants).

Risks of Changing Later: While Purity allows you to change a host personality later, doing so while volumes are connected and I/O is active can be disruptive. For many operating systems, a change in personality requires the host to be rebooted or the storage paths to be 'rescanned' to recognize the change in device capabilities.

Default Behavior: If no personality is selected, the FlashArray uses a 'Generic' personality suitable for standard Windows and Linux distributions. However, for specialized hypervisors like ESXi, failing to set the personality correctly from the start can lead to performance issues or lack of support for hardware acceleration features.

Why Option C is incorrect: Changing the personality after volumes are connected is reactive rather than proactive. It increases the risk of the host misinterpreting the storage device's capabilities, potentially leading to mount failures or path instability.

Question 3

Question Type: MultipleChoice

What is unified storage for Pure?

Options:

A- FlashArray runs both NFS and SMB protocols.
B- FlashArray runs both Block and File level protocols.
C- FlashArray runs both iSCSI and Fibre Channel (FC) protocols.

Answer:

B

Explanation:

Definition of Unified Storage: In the storage industry, 'Unified Storage' refers to a platform that can natively serve both Block-level storage (accessed via protocols like Fibre Channel, iSCSI, or NVMe-oF) and File-level storage (accessed via protocols like NFS or SMB) from a single pool of capacity and under a single management interface.

Pure Storage Implementation (FA File): Pure Storage achieved unified storage on the FlashArray through the introduction of Purity//FA File Services. Unlike traditional unified storage that often required a 'gateway' or separate hardware 'heads,' Pure's implementation runs natively on the FlashArray controllers.

Shared Resources: On a unified FlashArray, the global storage pool is shared between volumes (Block) and file systems (File). All of Pure's core data services---such as deduplication, compression, and SafeMode snapshots---apply globally across both block and file data.

Protocol Diversity: While Option A mentions NFS and SMB, those are strictly File protocols. Option

C mentions iSCSI and FC, which are strictly Block protocols. Only Option B correctly identifies the combination of Block and File, which defines the 'Unified' architecture of the FlashArray.

Question 4

Question Type: MultipleChoice

A new array is directly connected to a host with Direct Attach Copper (DAC) cables. The link does not come up.

Which document can be used to help identify the issue?

Options:

A- The FlashArray User Guide
B- FlashArray Transceiver and Cable Support article
C- The Port Usage and Definitions article

Answer:

B

Explanation:

When physical links fail to establish---especially when using Direct Attach Copper (DAC) cables or Twinax---the most common culprit is a hardware compatibility mismatch. Pure Storage arrays have specific requirements for optics and cabling to ensure optimal signal integrity and performance.

The FlashArray Transceiver and Cable Support article (available on the Pure Storage Support portal) is the authoritative, verified resource for this scenario. It provides a comprehensive, constantly updated compatibility matrix detailing exactly which vendor DAC cables (e.g., Cisco, Brocade, Arista) and transceivers are officially validated and supported for use with specific FlashArray models and port types. If an unsupported DAC cable is used, the switch or host bus adapter (HBA) on the array might simply refuse to bring the link up.

Here is why the other options are incorrect for this specific issue:

The FlashArray User Guide (A): This guide is excellent for day-to-day administration (volume creation, host grouping, etc.) but is too broad to contain granular, constantly updating hardware compatibility matrices for specific cables.

The Port Usage and Definitions article (C): This document explains the logical and physical purpose of the ports on the back of the controllers (e.g., defining which ports are used for management, replication, or host connectivity), but it does not dictate hardware transceiver or cable interoperability.

Question 5

Question Type: MultipleChoice

During a test failover using ActiveDR, what content will be presented to the target pod?

Options:

A- The content from the last periodic refresh
B- The content from the last real fail-over
C- The content from the undo pod

Answer:

C

Explanation:

ActiveDR is Pure Storage's continuous, near-sync replication solution. It differs fundamentally from standard asynchronous replication because it uses a continuous stream of data rather than snapshot-based 'periodic refreshes' (which eliminates Option A).

When you perform a test failover in ActiveDR, you do so by promoting the target pod. The target pod becomes writable, allowing your hosts and applications to run against the replicated data without disrupting the ongoing continuous replication from the source array in the background.

When the test is completed, you demote the target pod. To ensure that the data generated during your test failover isn't accidentally lost forever, ActiveDR automatically creates an undo pod at the exact moment of demotion.

If you need to resume that exact test failover scenario or recover the test data, you can re-promote the target pod and instruct ActiveDR to present the content from the undo pod. This unique mechanism allows storage administrators to seamlessly non-disruptively test, pause, and resume DR environments without affecting production protection.

Question 6

Question Type: MultipleChoice

A storage administrator has presented VMFS datastores from a FlashArray with 10TB of raw capacity.

Why would the administrator see system space when logging in to the FlashArray GUI?

Options:

A- Virtual machines have not yet issued an unmap command.
B- There is more than 2TB of reclaimable space on the FlashArray.
C- More than 2TB of volume snapshots were destroyed.

Answer:

B

Explanation:

On a Pure Storage FlashArray, 'System Space' is a specific GUI-reported metric. Purity has a predefined, hidden internal space budget---typically around 20% of the raw mapped capacity (which would be 2TB on a 10TB array)---reserved for internal array operations. This budget covers RAID/parity overhead, metadata, and reclaimable space (data from deleted volumes, snapshots, or overwritten blocks that are waiting for the backend garbage collection process to fully erase them from the flash chips).

Normally, this internal overhead stays below the 20% budget, and 'System Space' displays as 0.00 in the GUI. However, if an administrator deletes a massive amount of data at once, causing the reclaimable space to exceed that 2TB budget, the overflow is prominently displayed in the GUI as 'System Space.'

Here is why the other options are incorrect:

Virtual machines have not yet issued an unmap command (A): If a VMware VM deletes a file but the OS hasn't issued an UNMAP/TRIM command, the FlashArray is completely unaware that the data was deleted. Therefore, the array continues to report that capacity as standard Volume Space, not System Space.

More than 2TB of volume snapshots were destroyed (C): While destroying snapshots leads to reclaimable space, 'reclaimable space' (Option B) is the specific, correct Purity architectural term and metric that the system uses to calculate the internal budget threshold.

Question 7

Question Type: MultipleChoice

How is SAN Time measured?

Options:

A- Average time, measured in milliseconds, required to transfer data between the initiator and the array.
B- Average time, measured in milliseconds, that an IO request spends waiting to synchronize to the peer array.
C- Average time, measured in milliseconds, that an I/O request spends in the array waiting to be served.

Answer:

A

Explanation:

Understanding Total Latency: In a FlashArray environment, total latency as seen by the host application is the sum of several components. Pure Storage breaks this down into Array Time and SAN Time to help administrators pinpoint where performance bottlenecks exist.

SAN Time Definition: SAN Time represents the latency introduced by the network infrastructure between the host (initiator) and the FlashArray (target). This includes the time spent traveling across Fibre Channel or Ethernet switches, cables, and host bus adapters (HBAs). It is calculated by taking the total round-trip time measured by the host and subtracting the time the FlashArray spent processing the I/O.

Metric Breakdown: * Array Time: The time the FlashArray takes to process the I/O once it hits the front-end ports (Option C describes internal array time).

SAN Time: The transit time for the request to reach the array and the response to return to the host (Option A).

Wait Time: In ActiveCluster environments, there is also 'Mirror Latency,' which is the time spent synchronizing data to a peer array (Option B).

Troubleshooting Value: If a user reports high latency but the FlashArray GUI shows very low Array Time, the administrator can look at the SAN Time metric. A high SAN Time indicates an issue with the fabric, such as a failing SFP, a congested switch port, or oversubscribed ISLs (Inter-Switch

Links).

Question 8

Question Type: MultipleChoice

What happens when you demote the original source pod?

Options:

A- It saves a temporary copy of the source pod content in the eradication bin.
B- Replication is reversed.
C- Replication is paused.

Answer:

B

Explanation:

ActiveCluster and Pod Roles: In a Pure Storage ActiveCluster or ActiveDR environment, a Pod is a management container for volumes. To move workloads or perform a planned failover between two arrays, you use the Promote and Demote commands.

The Reversal Process: When you have two pods in a replication relationship (Source and Target), data flows from the Promoted (Active/Source) pod to the Demoted (Passive/Target) pod.

When you Demote the current source, it transitions from a 'read-write' state to a 'read-only' (passive) state.

If the other pod in the pair is then Promoted, Purity automatically intelligently reverses the direction of replication. The array that was previously receiving data now begins sending incremental updates back to the original source.

Continuous Protection: This design ensures that you don't have to manually tear down and recreate replication schedules every time you switch production sites. The system tracks the metadata changes and ensures that only the delta (changed blocks) are sent in the new direction.

Why Option C is incorrect: If replication were simply paused, the two sites would quickly drift out of sync, making it impossible to fail back without a full baseline resync.

Why Option A is incorrect: Demoting a pod does not delete any data; it simply changes the

access characteristics and replication role. The data remains fully intact on the storage media.

Question 9

Question Type: MultipleChoice

The SNMP protocol can be used to monitor what statistics of the array's performance?

Options:

A- Latency, bandwidth, IOPS
B- Host and volume performance
C- Available and used capacity

Answer:

A

Explanation:

SNMP Capabilities on FlashArray: Pure Storage FlashArrays support SNMP (Simple Network Management Protocol) versions v2c and v3. The primary use case for SNMP on the FlashArray is for integration with third-party monitoring tools (like SolarWinds, Nagios, or Zabbix) to provide a real-time health and performance heartbeat.

Performance Metrics (MIBs): The Pure Storage Management Information Base (MIB) specifically exposes high-level array performance metrics. These include the 'Big Three' of storage performance: Latency (response time in microseconds), Bandwidth (throughput in bytes per second), and IOPS (I/O operations per second).

Scope of Monitoring: While the FlashArray GUI and CLI provide deep granularity (per host, per volume, or per pod), standard SNMP queries typically focus on array-wide performance and health status. Detailed 'Host and volume performance' (Option B) is generally more effectively monitored via the Pure Storage REST API or the Pure1 VM Analytics, as SNMP is less efficient for pulling large tables of per-object data.

Capacity vs. Performance: While SNMP can report on capacity (Option C) through specific OIDs, the question specifically asks about performance statistics. In the context of Pure Storage monitoring documentation, the core performance metrics provided via SNMP traps and polling are latency, bandwidth, and IOPS.