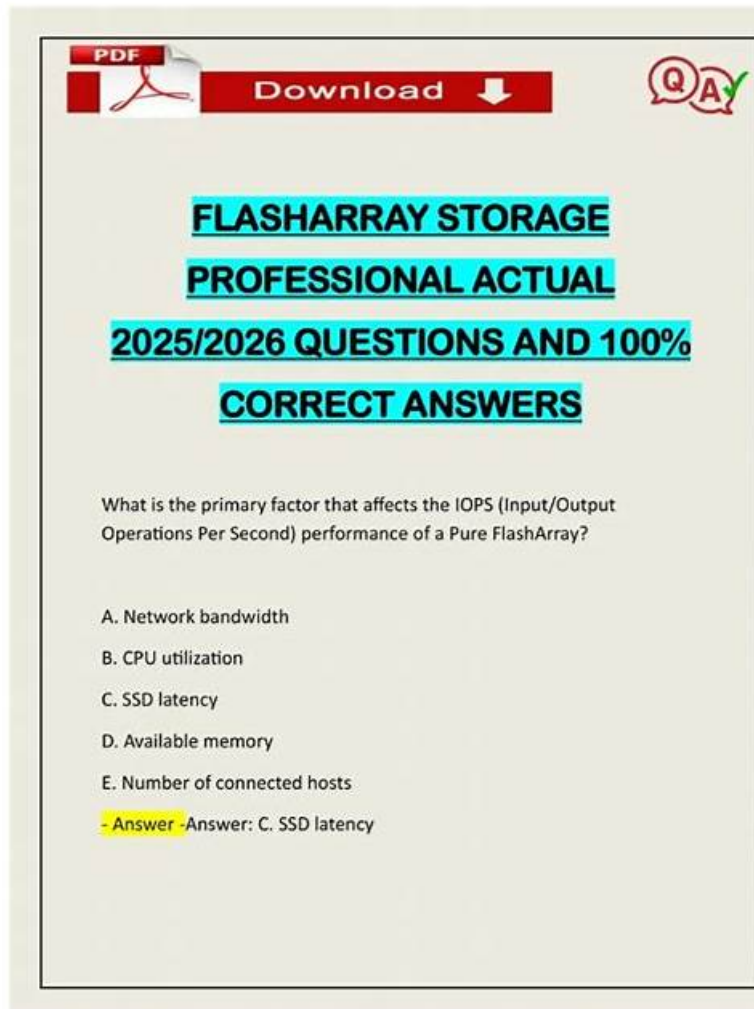


Pass-sure FlashArray-Storage-Professional Training Materials - FlashArray-Storage-Professional Quiz Torrent & FlashArray-Storage-Professional Exam Bootcamp



Preparing for the Pure Certified FlashArray Storage Professional (FlashArray-Storage-Professional) test can be challenging, especially when you are busy with other responsibilities. Candidates who don't use FlashArray-Storage-Professional dumps fail in the FlashArray-Storage-Professional examination and waste their resources. Using updated and valid FlashArray-Storage-Professional Questions; can help you develop skills essential to achieve success in the FlashArray-Storage-Professional certification exam.

Pure Storage FlashArray-Storage-Professional Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"> • FA File: Covers configuration and management of FA File services, including DNS setup, Active Directory integration, and protocol access. Focuses on enabling secure and efficient file sharing across the organization.

Topic 2	<ul style="list-style-type: none"> • Troubleshooting: Covers identification and resolution of configuration errors, performance issues, and replication problems using Pure Storage diagnostic tools and alerts. Includes port configuration and predictive support mechanisms to maintain system reliability.
Topic 3	<ul style="list-style-type: none"> • Monitoring: Covers the use of Pure1, GUI, and CLI tools to monitor array health, generate reports, and analyze performance and capacity metrics. Includes data reduction ratios, meta forecasting, and proactive capacity planning.
Topic 4	<ul style="list-style-type: none"> • Data Protection: Covers snapshot management, replication configuration, policy management, SafeMode, and advanced replication technologies such as ActiveDR. Focuses on ensuring data availability, disaster recovery, and protection against data loss.
Topic 5	<ul style="list-style-type: none"> • Administration: Covers core administrative tasks including volume configuration, array management, host connections, third-party integrations, and security protocols. Focuses on best practices for maintaining optimal performance and secure access across the storage environment.

>> **FlashArray-Storage-Professional Valid Test Voucher** <<

Pure Storage Valid FlashArray-Storage-Professional Valid Test Voucher – Pass FlashArray-Storage-Professional First Attempt

As the rapid development of the world economy and intense competition in the international, the leading status of knowledge-based economy is established progressively. A lot of people are in pursuit of a good job, a FlashArray-Storage-Professional certification, and a higher standard of life. It is very important for us to keep pace with the changeable world and update our knowledge if we want to get a good job, a higher standard of life and so on. First, we need to get a good FlashArray-Storage-Professional Quiz prep. Because we only pass FlashArray-Storage-Professional exam and get a certificate, we can have the chance to get a decent job and make more money.

Pure Storage Pure Certified FlashArray Storage Professional Sample Questions (Q48-Q53):

NEW QUESTION # 48

An administrator is running commands to verify NVME/TCP connectivity from the hosts to the FlashArray. They use the command `ping -M do -s 8972 <ip_addr>` from the initiator and it fails.

What should the administrator do to resolve the issue?

- A. Check the MTU of 9000 is set on each hop to the FA.
- B. Engage support to enable NVME/ TCP services.
- C. Run the command from the target.

Answer: A

Explanation:

When configuring NVMe/TCP (or iSCSI) for optimal performance on a Pure Storage FlashArray, configuring Jumbo Frames (an MTU of 9000) end-to-end is a standard best practice.

The command `ping -M do -s 8972 <ip_addr>` is specifically used to verify Jumbo Frame configuration across the network.

The `-M do` flag sets the "Do Not Fragment" (DF) bit, meaning the network is not allowed to break the packet into smaller pieces.

The `-s 8972` flag sets the ICMP data payload to 8972 bytes. When you add the standard 8-byte ICMP header and the 20-byte IP header, the total packet size equals exactly 9000 bytes.

If this ping command fails, it indicates that somewhere along the network path between the host (initiator) and the FlashArray (target), a switch port, router, or network interface is not configured to support an MTU of 9000. The packet is being dropped because it is too large and cannot be fragmented. The administrator must verify the MTU settings on every network hop (switches, routers, and host NICs) to resolve the issue.

Here is why the other options are incorrect:

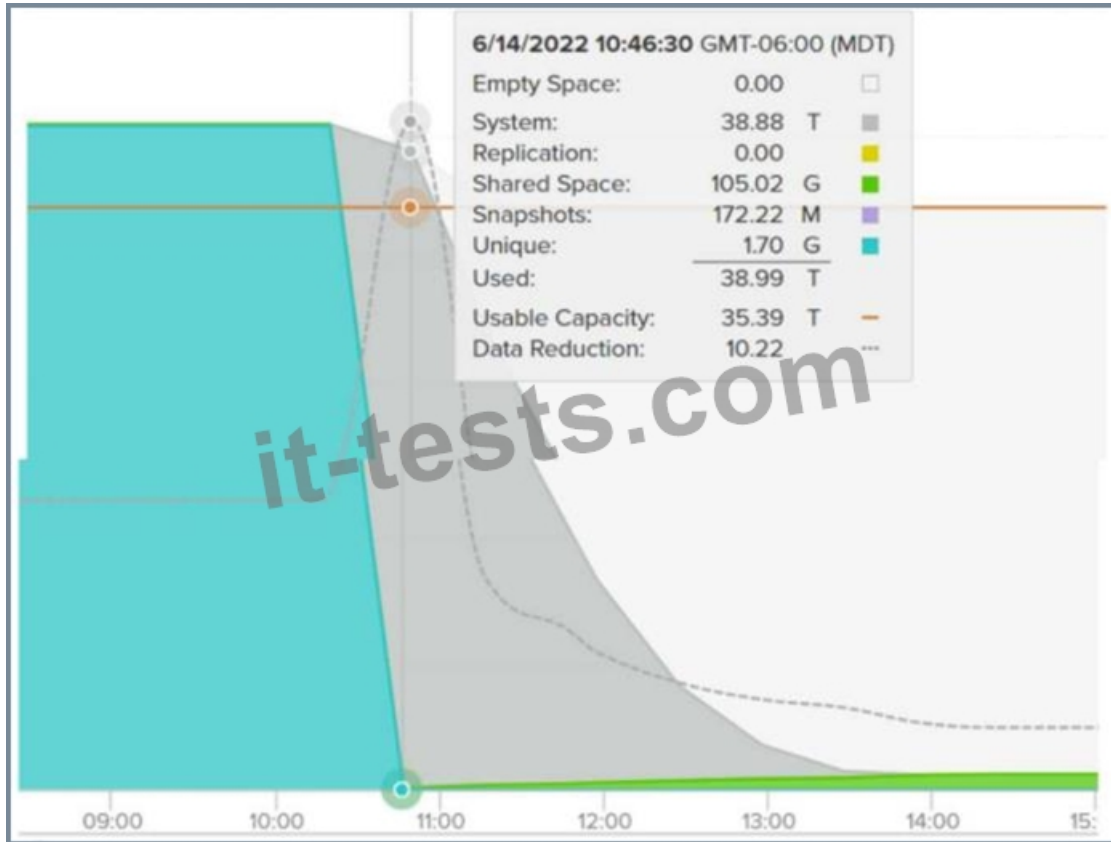
Engage support to enable NVME/ TCP services (A): The failure of a Jumbo Frame ping test is a Layer 2/Layer 3 network configuration issue, not an indicator that the NVMe/TCP storage protocol service is disabled on the array.

Run the command from the target (C): While pinging from the FlashArray back to the host is a valid secondary troubleshooting step,

it will likely also fail if the network path doesn't support Jumbo Frames. The actual resolution is to fix the MTU on the network hops.

NEW QUESTION # 49

Refer to the exhibit:



What is the most likely cause of system space in this capacity graph?

- A. SafeMode was enabled on the system.
- B. Space reclamation stopped working properly.
- C. Large volume destroyed and eradicated.

Answer: C

Explanation:

Analyzing the Graph: The provided capacity graph shows a sharp, vertical drop in Unique and Shared space (the teal/green sections) at approximately 10:45. Simultaneously, there is a massive spike in System space (the grey shaded area).

The "System" Space Spike: In Purity, when a volume is deleted and eradicated, the data isn't physically wiped from the flash modules instantly. Instead, the logical references are removed, and the blocks are moved into a "Pending Reclamation" state. In older versions of Purity reporting (or specific views), this pending work is often categorized under System space until the background Garbage Collection (GC) process can physically reclaim it.

The Decay Curve: Notice the grey "System" area gradually tapers down over the following hours (from 11:00 to 14:00). This is the visual representation of the array's Garbage Collection engine working in the background to physically free up those blocks.

Why Option C is incorrect: If SafeMode were enabled, the space would not move into "System" space in this manner; it would remain locked in the Snapshots or Destroyed bucket for the duration of the retention period, preventing the graph from dropping at all.

Why Option A is incorrect: If space reclamation stopped working, the grey "System" area would stay flat and would not decay over time. The fact that it is trending downward proves that reclamation is actively processing the deleted volume.

NEW QUESTION # 50

An administrator is preparing an array pair for ActiveDR and is trying to calculate the total minimum bandwidth requirement.

What percent of bandwidth above the incoming write rate should be allocated to accommodate for unexpected write bursts and still maintain near-sync RPO?

- A. 10%
- B. 50%
- C. 30%

Answer: C

Explanation:

ActiveDR Bandwidth Sizing: ActiveDR is a continuous, asynchronous replication technology designed to provide near-zero RPO. Because it streams data continuously rather than in discrete snapshot intervals, the bandwidth between the source and target arrays must be able to handle the application's write workload.

Handling Write Bursts: Application workloads are rarely flat; they have peaks and valleys. If you size the bandwidth exactly to the average change rate, any burst in write activity will cause the replication lag to increase, thereby increasing your RPO.

The 30% Rule: Pure Storage best practices and sizing guides recommend providing a 30% buffer (headroom) above the measured average write rate. This extra capacity ensures that during a high-IO period, the replication engine has enough "pipe" to catch up quickly and return to a near-sync state.

Calculation Example: If a workload generates an average of 100 MB/s of new unique data, the administrator should ensure at least 130 MB/s of usable, dedicated bandwidth is available between the sites.

Consequences of Under-sizing: If only 10% (Option A) is used, the array may struggle to recover from even minor bursts, leading to a consistently climbing RPO. 50% (Option B) is often considered safe but can be cost-prohibitive or overkill for standard networking budgets unless the workload is exceptionally volatile.

NEW QUESTION # 51

Which protection group cannot be ratcheted for SafeMode?

- A. Protection groups with hosts or hostgroups
- B. A protection group without a local snapshot schedule
- C. A default protection group

Answer: B

NEW QUESTION # 52

How would a FlashArray administrator view external latency for write requests for a specific volume?

- A. Analysis > Performance > Volumes > Select the appropriate volume > Deselect "Read" and "Mirrored Write"
- B. Storage > Volumes > Select the appropriate volume > Details
- C. Health > Network > Select the appropriate protocol > Select the appropriate port

Answer: A

Explanation:

In the Pure Storage FlashArray GUI, granular performance metrics (Latency, IOPS, Bandwidth) are located under the Analysis > Performance tabs. When you navigate to the Volumes sub-tab and select a specific volume, Purity displays a unified line graph tracking the performance of that volume over time.

By default, the Latency graph simultaneously plots Read, Write, and Mirrored Write (for volumes participating in an ActiveCluster or synchronous replication pod) latencies. Because these lines can overlap or compress the Y-axis (especially if one metric spikes), isolating a specific metric requires interacting with the graph's legend.

To view the exact, un-obscured latency for standard write requests to that volume, the administrator should click on "Read" and "Mirrored Write" in the chart's legend. This deselects those metrics, effectively hiding their lines from the graph and automatically rescaling the view to exclusively display the host write latency.

Here is why the other options are incorrect:

Health > Network (A): The Health tab is used to check the hardware status of the physical controller ports, including link state and errors. While you might see port-level throughput or queue depth here, it does not provide volume-specific application latency.

Storage > Volumes > Details (B): The Storage tab is primarily used for provisioning and configuration management. Clicking on a volume here will show its size, data reduction ratio, snapshot policies, and connected hosts, but it does not provide detailed interactive performance graphs.

NEW QUESTION # 53

