

KillTest

Higher Quality, Better Service!



Q&A

<http://www.killtest.com>

We offer free update service for one year.

Exam : HPE0-J82

Title : HPE Storage Architect

Version : DEMO

1.A Backup and Recovery Engineer is attempting to right-size the storage capacity for an aging hybrid array. The primary application owner claims their database only grows by 100 GB per week, yet the storage array telemetry indicates the volume's physical footprint is expanding by 1.5 TB per week.

[Telemetry Logs - Volume: Fin_DB_01]

Mon 02:00 - Array Snap Created (Retention: 30 Days)

Tue 02:00 - Array Snap Created (Retention: 30 Days)

...

[Analytics Profile]

Host Data Change Rate: 2% Daily (Highly random overwrites)

Inline Deduplication: Disabled (Encrypted Payload)

Space Reclamation (UNMAP): Enabled & Active

Which TWO diagnostic conclusions accurately explain the massive discrepancy between the application owner's perceived growth and the actual physical capacity consumption? (Choose 2.)

- A. The application's 2% daily random overwrite rate forces the array's 30-day snapshot retention policy to permanently lock massive amounts of modified physical blocks
- B. The host application is transmitting an encrypted payload, which mathematically neutralizes the storage array's ability to deduplicate the snapshot differentials
- C. The application owner's calculation of 100 GB/week represents pure logical net-new data insertion, ignoring the storage-level retention overhead of their own highly volatile overwrites
- D. The storage controllers are suffering from severe CPU contention, which artificially inflates the physical capacity metrics reported to the telemetry engine
- E. The Space Reclamation (UNMAP) protocol is malfunctioning, trapping deleted database tables as "active" physical blocks on the backend storage

Answer: A, C

2.A Customer Success Manager is explaining the financial benefits of modern HPE storage arrays to a client. The client is confused by the terminology used in the capacity sizing proposal, specifically the difference between "usable capacity" and "effective capacity."

How does the concept of "effective capacity" mathematically differ from "usable capacity" in modern HPE storage sizing methodologies?

- A. Usable capacity is the physical space available after RAID overhead, while effective capacity is the logical space available after applying deduplication and compression ratios.
- B. Usable capacity includes the buffer reserved for system snapshots, whereas effective capacity is strictly dedicated to host-written volumes.
- C. Effective capacity represents the raw, unformatted hardware space before RAID penalties, while usable capacity represents the space after RAID is applied.
- D. Effective capacity guarantees a 4:1 data reduction ratio universally across all workloads, while usable capacity guarantees storage performance SLAs.

Answer: A

3.A Data Protection Specialist is reviewing the system event logs of an HPE Nimble storage array. The array is hosting a critical Oracle database volume.

[System Event Log - 04:00:00 AM]

Warning: Snapshot creation skipped for Volume Collection 'Oracle_DB_VolCol'.

Reason: Maximum snapshot retention limit reached (10,000 snapshots).

The administrator investigating the issue reveals they configured an extremely aggressive policy: Take a snapshot every 5 minutes and retain all of them locally for 1 year.

How does the storage array's architecture fundamentally handle this configuration flaw, and what is the required remediation?

- A. The array automatically converts the oldest snapshots into an S3 archive format, which takes 24 hours to process; the specialist must simply wait for the deduplication queue to clear before the schedule resumes
- B. The array relies on the host hypervisor (VMware vCenter) to execute a scheduled task to delete old snapshots; the specialist must reboot the vCenter server to force the hypervisor to clear the snapshot cache
- C. The storage array automatically and silently deletes the base production volume to make room for new snapshots; the specialist must instantly restore the production LUN from tape
- D. The array's operating system imposes a hard architectural limit on the maximum number of retained snapshots per volume (or globally) to protect the controller's metadata processing performance; the specialist must drastically reduce the retention schedule to stay within the array's supported maximums

Answer: D

4.A Storage Administrator is troubleshooting a performance issue on an HPE Alletra array. A specific Linux database server is generating massive I/O, causing latency across the array.

The administrator pulls the performance telemetry log.

[Alletra Performance Telemetry - Node 0]

Volume Name: DB_Prod_Vol

Host Protocol: NVMe/TCP

Read IOPS: 120,000

Write IOPS: 80,000

Average Latency: 0.8ms

Backend NVMe SSD Utilization: 45%

Controller CPU Utilization: 92% (Saturated)

Despite the backend flash drives being underutilized, the array controller's CPU is saturated.

Which inherent architectural characteristic of the NVMe/TCP protocol is the primary contributor to this specific controller CPU bottleneck, compared to FC-NVMe or RoCEv2?

- A. NVMe/TCP requires the storage controller to constantly poll the centralized Discovery Controller (DC) appliance over the management network for every single I/O request
- B. NVMe/TCP forces the storage controller's CPU to process the heavy encapsulation and decapsulation of NVMe commands into standard TCP/IP packets, consuming massive compute cycles at high IOPS rates
- C. NVMe/TCP inherently disables the array's inline data reduction ASIC, forcing the main x86 CPUs to calculate deduplication hashes via software emulation
- D. NVMe/TCP utilizes a dynamic, variable block size that constantly forces the controller to fragment and reassemble payload stripes before committing them to the backend NVMe drives

Answer: B

5.A Storage Procurement Specialist is reviewing an HPE GreenLake consumption bill generated via the

Data Ops Manager portal. The customer is furious because their monthly OPEX bill spiked by 300% over the weekend.

[Data Ops Manager - GreenLake Billing Telemetry]

Friday:

Reserved Capacity Consumed: 100%

Buffer Capacity Consumed: 0%

Monday:

Reserved Capacity Consumed: 100%

Buffer Capacity Consumed: 80%

Event Log (Saturday 02:00 AM): System-wide Antivirus Scan Initiated by SecOps team.

The SecOps team insists that an Antivirus scan only *reads* files and therefore cannot consume physical storage capacity.

How does the analyst explain the architectural mechanics of the storage array that caused this massive billing spike during a "read-only" scan?

- A. The storage array controllers detected the intensive sequential read workload and automatically disabled the deduplication engine to free up CPU cycles, physically hydrating all the data and instantly consuming the buffer capacity
- B. The Data Ops Manager billing engine is misconfigured; it mathematically calculates OPEX costs based strictly on host network bandwidth (MB/s) rather than physical disk capacity consumed
- C. The customer's backup software was configured to take array-based snapshots prior to the scan; because the Antivirus software inadvertently modifies the "Last Accessed" timestamp metadata on every single file, the storage array recorded those metadata changes as a massive differential write payload, blooming the physical snapshot size into the buffer tier
- D. The array's Automated Sub-LUN Tiering engine panicked and migrated all 100TB of cold data from the mechanical HDDs directly into the NVMe Buffer tier to satisfy the sequential read request, triggering the billing event

Answer: C