

KillTest

Higher Quality, Better Service!



Q&A

<http://www.killtest.com>

We offer free update service for one year.

Exam : **HPE7-V01**

Title : Advanced HPE Edge-to-
Cloud Solutions

Version : DEMO

1.Scenario: A regional healthcare provider is evaluating an infrastructure refresh for their 5,000-seat VDI environment. Current storage telemetry indicates 85,000 IOPS with 12ms latency during morning boot storms, causing severe application timeouts. The Chief Information Security Officer mandates that all patient data remain physically on-premises to comply with strict data sovereignty regulations. Furthermore, the CFO requires a transition from capital expenditures to a scalable, pay-per-use financial model.

A. Migrate the VDI workloads to a public cloud Desktop-as-a-Service model to eliminate upfront capital expenditures and guarantee flexible scalability.

B. Propose an HPE GreenLake for VDI solution utilizing HPE Alletra Storage to ensure on-premises compliance and pay-per-use economics.

C. Implement an HPE Synergy architecture combined with HPE Primera storage, purchased through a standard 36-month HPE Financial Services lease.

D. Deploy HPE SimpliVity hyperconverged infrastructure in a colocation facility, utilizing a traditional procurement model to maximize total control.

Answer: B

Explanation:

Option B perfectly addresses all constraints. HPE GreenLake provides a true pay-per-use (OpEx) financial model while keeping the physical infrastructure on-premises to satisfy the CISO's data sovereignty requirements, and Alletra can easily handle the IOPS/latency demands.

Option A is incorrect because migrating to the public cloud directly violates the CISO's strict on-premises data sovereignty mandate.

Option C is incorrect because a standard HPE Financial Services lease is typically a fixed payment structure (CapEx-like/fixed OpEx), not a scalable pay-per-use consumption model.

Option D is incorrect as it utilizes a traditional procurement (CapEx) model, which directly violates the CFO's financial requirement.

2.Scenario: A global logistics enterprise needs to upgrade its core supply chain database, which operates as a mission-critical workload requiring 100% data availability. The database must replicate synchronously to a secondary site located 25 kilometers away with an RPO of zero. The customer desires an operational expense model but has strict policies against deploying any third-party software agents directly onto the database host operating systems for replication purposes.

A. Deploy HPE Alletra 6000 storage at both sites via GreenLake, utilizing HPE Zerto for agentless synchronous replication and zero RPO.

B. Utilize HPE SimpliVity federations at both sites via GreenLake, configuring native backup policies to achieve the required zero RPO metric.

C. Position HPE Alletra 9000 arrays at both sites via HPE GreenLake, leveraging Active Peer Persistence for host-transparent synchronous replication.

D. Implement HPE Synergy with traditional SAN storage on a 5-year lease, using database-level native mirroring to ensure zero RPO and 100% uptime.

Answer: C

Explanation:

Option C is the only correct architectural fit. The HPE Alletra 9000 natively guarantees 100% availability for mission-critical workloads and supports Active Peer Persistence (APP) for transparent, array-based synchronous replication (RPO=0) without requiring host-level agents, all delivered via the GreenLake

OpEx model.

Option A is incorrect because HPE Zerto utilizes continuous data protection (CDP) which is near-synchronous (RPOs in seconds), not true synchronous RPO=0, and Alletra 6000 natively offers 99.9999% availability, not 100%.

Option B is incorrect as SimpliVity's native backup policies cannot achieve a true RPO of zero for active synchronous clustering.

Option D is incorrect because a standard lease does not guarantee an OpEx classification, and native database mirroring relies on host-level configurations, failing the optimal infrastructure-level requirement.

3.Scenario: An enterprise customer is migrating a massive, unstructured data repository (currently 2PB, growing at 30% annually) to a new architecture. The primary requirements are highly scalable object storage, immutability to protect against ransomware, and seamless integration with their existing Kubernetes container environment for AI data analytics. They require a scale-out architecture that can add capacity nondisruptively without degrading read/write performance.

A. Architect a scale-out cluster using HPE Alletra 4000 series servers running scalable object storage software with S3 object lock enabled.

B. Design a solution utilizing HPE Alletra 9000 block storage arrays with Fibre Channel connectivity to ensure maximum throughput for AI workloads.

C. Architect an HPE SimpliVity multi-node federation, utilizing its native inline deduplication and compression to accommodate unstructured data.

D. Design an HPE GreenLake for Block Storage architecture, leveraging HPE Backup and Recovery Service to provide immutable data protection.

Answer: A

Explanation:

Option A is the correct design for massive unstructured data. The HPE Alletra 4000 (formerly Apollo 4000 series) paired with software-defined object storage (such as Scality or Cloudian) provides limitless scale-out S3 capabilities, native S3 object lock for immutability against ransomware, and ideal integration for Kubernetes/AI pipelines.

Option B is incorrect because Alletra 9000 is a scale-up block storage array designed for structured, mission-critical relational databases, not 2PB of unstructured object data.

Option C is incorrect as hyperconverged infrastructure (SimpliVity) is optimized for VM-centric workloads, not petabyte-scale unstructured object storage repositories.

Option D is incorrect because block storage architectures do not natively address the customer's requirement for a scalable object/S3 repository.

4.Scenario: A customer is operating a 6-node VMware vSphere cluster on traditional rackmount servers connected to an aging iSCSI SAN. They are experiencing severe performance degradation during nightly backup windows due to I/O bottlenecks. The customer wants to upgrade to an architecture that provides predictive analytics to proactively identify VM-level bottlenecks, reduces data footprint through always-on data reduction, and allows independent scaling of compute and storage.

A. Upgrade the environment to HPE Synergy with an integrated D3940 storage module to localize backup traffic and eliminate the iSCSI network bottlenecks.

B. Migrate the workloads to HPE SimpliVity to utilize its hardware-accelerated deduplication and compression, and scale compute/storage symmetrically.

- C. Migrate the environment to public cloud IaaS instances, utilizing native cloud monitoring tools to dynamically scale resources during backup windows.
- D. Deploy an HPE Nimble dHCI solution to enable independent resource scaling, while utilizing HPE InfoSight for full stack cross-stack predictive analytics.

Answer: D

Explanation:

Option D perfectly matches the requirements. HPE Nimble dHCI (disaggregated HCI) allows compute and storage to scale independently, utilizes always-on deduplication and compression, and features HPE InfoSight, which provides unparalleled cross-stack predictive analytics to pinpoint exact VM-level I/O bottlenecks.

Option A is incorrect as Synergy with a D3940 DAS module does not provide the advanced VM-level predictive analytics or the always-on data reduction expected in a modernized storage platform.

Option B is incorrect because, while SimpliVity offers excellent data reduction, it scales symmetrically (adding nodes increases both compute and storage simultaneously), directly violating the customer's requirement for independent scaling.

Option C is incorrect because migrating to the public cloud is a complete architectural overhaul that ignores the request to upgrade their on-premises virtualization architecture, and native cloud tools lack InfoSight's deep infrastructure-level AI telemetry.

5.Scenario: A research university is experiencing intermittent application timeouts during AI model training tasks. Telemetry data indicates sudden, massive traffic microbursts overwhelming the current 10GbE top-of-rack switches, though average bandwidth utilization remains below 40%. The customer requires a modern network architecture capable of deep packet inspection, real-time microburst visibility, and automated root-cause analysis without deploying external aggregation taps.

- A. Upgrade the host server network interface cards to dual-port 100GbE adapters to brute-force the bandwidth bottlenecks during model training.
- B. Deploy traditional HPE FlexFabric 5940 switches and configure standard SNMP polling at one-minute intervals to monitor the traffic spikes.
- C. Implement HPE Aruba Networking CX switches utilizing the embedded Network Analytics Engine (NAE) for real-time telemetry and microburst detection.
- D. Migrate the AI workloads to a public cloud provider to utilize their natively provided software-defined network monitoring toolsets.

Answer: C

Explanation:

Option C is the correct architectural choice. HPE Aruba Networking CX switches feature the Network Analytics Engine (NAE), which runs directly on the switch operating system. NAE provides real-time, time-series telemetry and can automatically run Python scripts to capture data during microburst events that standard polling misses.

Option A is incorrect because upgrading host NICs does not provide the required deep packet inspection or network visibility, and the existing switches would still drop packets during microbursts.

Option B is incorrect because legacy FlexFabric switches rely on standard SNMP polling, which typically samples data at 1- to 5-minute intervals, completely missing sub-second microbursts.

Option D is incorrect as it bypasses the requirement to modernize the on-premises network architecture and introduces massive data egress costs for AI workloads.