

KillTest

높은 품질, 더 나은 서비스



연습 문제

<https://www.killtest.kr>

1년간 무료 업데이트 제공

Exam : HPE0-S62

Title : HPE Compute Architect

Version : DEMO

1.A hospital is evaluating an on-premises AI inference environment for radiology images. Images originate at two sites, patient data must remain local, and the team expects the pilot to grow if it meets clinical targets.

Before selecting an HPE architecture, which TWO discovery inputs are most important to document? Select TWO answers.

- A. The inference-latency target and the network or connectivity constraints between the sites
- B. The current and forecast AI workload, including expected growth and success criteria
- C. The preferred desktop monitor vendor for radiologists
- D. The final GPU clock speed before the team has described the models to be used

Answer: A and B

Explanation:

An architect should first establish the business outcome, workload characteristics, growth expectation, data-locality requirement, and relevant connectivity constraints. Those inputs determine whether a proposed solution can meet the intended clinical outcome. A display-vendor preference does not drive the compute architecture. Selecting a detailed GPU characteristic before qualifying the workload reverses the design process.

Blueprint objective: 1 — Determine the customer business needs, requirements, and desired outcomes.

Sub-objective: Determine customer requirements before selecting an HPE solution.

Bloom level: Analyze.

References: HPE Compute Solutions, Rev. 26.21 — design and customer-requirements objectives

2.An organization wants a common cloud-based operating model for a fleet that includes supported HPE servers and several third-party servers. The architect is documenting the onboarding approach described in the HPE Compute Solutions, Rev. 26.21 course.

Which TWO statements are correct? Select TWO answers.

- A. Supported HPE servers can be onboarded for direct management by HPE Compute Ops Management.
- B. Third-party servers can be onboarded for credential-based management.
- C. Third-party servers must be converted to HPE hardware before they can be monitored.
- D. HPE Synergy Image Streamer is required before any server can be onboarded.

Answer: A and B

Explanation:

The course distinguishes direct management onboarding for HPE servers from credential-based management onboarding for third-party servers. That distinction lets the architect plan a mixed environment without claiming identical lifecycle capabilities for every system. Replacing third-party hardware is not an onboarding requirement, and Image Streamer is a Synergy deployment component rather than a prerequisite for Compute Ops Management onboarding.

Blueprint objective: 2 — Describe the benefits of HPE Compute Ops Management integrations.

Sub-objective: Identify management approaches for HPE and third-party servers.

Bloom level: Apply.

References: HPE Compute Solutions, Rev. 26.21 — Compute Ops Management onboarding; HPE Compute Ops Management overview

3.A customer is designing a production HPE Synergy environment that uses Image Streamer for

stateless server boot volumes. The customer requires high availability for boot-volume access and does not accept a non-production design.

Which Image Streamer design best meets the requirement? Select ONE answer.

- A. One Image Streamer appliance in a single frame
- B. Two Image Streamer appliances in the same frame, with no second frame
- C. Local boot disks on each compute module with no Image Streamer appliances
- D. A pair of Image Streamer appliances in different frames in a multi-frame active-active configuration

Answer: D

Explanation:

HPE documents a pair of Image Streamer appliances in different frames as a multi-frame, active-active HA configuration suitable for production. A single-frame configuration provides non-HA boot-volume access and is intended for development, test, or proof-of-concept use. Local boot disks do not provide the requested Image Streamer stateless-server design.

Blueprint objective: 3 — Position solutions with HPE Synergy and HPE Morpheus VM Essentials Software.

Sub-objective: Describe the benefits of HPE Synergy components.

Bloom level: Apply.

References: HPE Synergy Image Streamer — HA and stateless-server behavior

4.A manufacturer asks an architect to recommend an HPE AI solution for a new internal knowledge assistant. The sponsor says that confidential engineering documents must remain under company control, but has not yet defined the expected users, model workflow, data sources, or service-level target.

What should the architect do first? Select ONE answer.

- A. Select the largest GPU configuration to avoid a later redesign.
- B. Assess the customer's AI maturity, intended workload and use case, data requirements, and desired outcome.
- C. Begin fine-tuning a model with publicly available data.
- D. Choose a virtualization platform before identifying how the assistant will be used.

Answer: B

Explanation:

HPE's course objectives require the architect to assess AI maturity, workloads, and use cases before positioning an HPE AI solution. The stated confidentiality requirement is important, but it does not replace workload, data-source, scale, and outcome discovery. Selecting the largest configuration or a platform before that assessment risks an unsuitable design.

Blueprint objective: 4 — Position HPE AI solutions.

Sub-objective: Assess customer AI maturity, workload, and use case.

Bloom level: Analyze.

References: HPE Compute Solutions, Rev. 26.21 — AI requirements and solution-positioning objectives

5.An architect is using the HPE Private Cloud AI 1.5 solution-configuration table. The customer needs an environment for inferencing plus retrieval-augmented generation (RAG), and does not plan to fine-tune models in this phase.

Which HPE Private Cloud AI configuration is aligned to that use case? Select ONE answer.

A. Developer System

B. Small

C. Medium

D. Large

Answer: C

Explanation:

In the cited HPE Private Cloud AI 1.5 configuration table, the Medium configuration is positioned for inferencing plus RAG. The Large configuration adds fine-tuning to the stated use case. The Developer System is an AI sandbox, and Small is positioned for inferencing rather than the stated RAG requirement.

Blueprint objective: 5 — Describe the benefits of NVIDIA with HPE solutions.

Sub-objective: Distinguish HPE Private Cloud AI configuration sizes.

Bloom level: Apply.

References: HPE Private Cloud AI 1.5 Administration Guide — solution configurations