

KillTest

Higher Quality, Better Service!



Q&A

<http://www.killtest.com>

We offer free update service for one year.

Exam : HPE2-B08

**Title : HPE Private Cloud AI
Solutions**

Version : DEMO

1. During a discovery call, a customer from a telecommunications company explains their primary goal: "We need to analyze network traffic patterns in real-time to detect anomalies that could indicate a security threat or a network outage."

Which key use case for HPE Private Cloud AI does this represent?

- A. AI Cybersecurity
- B. Code Generation
- C. AI Recommender System
- D. Document Chat

Answer: A

2. A data analytics team is running workloads on an HPE Private Cloud AI solution. They observe that a data ingestion job is not meeting performance expectations, suspecting a CPU bottleneck. They believe the application is not correctly leveraging GPUDirect Storage (GDS), forcing data to be copied through the server's main memory before reaching the GPU.

Which are valid reasons why GDS might not be functioning correctly? (Choose 3.)

- A. The network switches are not configured for lossless operation (e.g., PFC is disabled).
- B. The NVIDIA GPUs have been configured with Multi-Instance GPU (MIG), which enhances GDS performance.
- C. The application is using a standard TCP/IP socket for data transfer instead of an RDMA-based library.
- D. The HPE GreenLake for File Storage array is using SATA SSDs instead of NVMe SSDs.
- E. The NVIDIA peer memory driver has not been installed on the guest VM.

Answer: A, C, E

3. A customer's ML Engineer states, "We need to deploy our trained models, and our top priority is simplifying the process. We want to treat our models like cattle, not pets—packaging them into standardized, optimized containers that we can deploy and scale easily via an API."

This statement describes the primary benefit of which software component in the HPE Private Cloud AI stack?

- A. NVIDIA NIM (NVIDIA Inference Microservices)
- B. HPE Data Fabric
- C. Apache Airflow
- D. JupyterLab

Answer: A

4. An architect has used the HPE Intelligent Configurator and determined that a "Medium - Standard" configuration is required.

When creating the final quote in One Config Advanced (OCA), what is the mandatory prerequisite that the customer is responsible for providing for the solution to be valid?

- A. A valid, bring-your-own-license (BYOL) for VMware vSphere Foundation (VVF) or VMware Cloud Foundation (VCF).
- B. A rack elevation diagram for the data center.
- C. A list of all end-users who will access the system.
- D. A subscription to the NVIDIA NGC catalog for downloading AI models.

Answer: A

5. A customer is considering the HPE Private Cloud AI solution. They need to run a moderately sized RAG application for 150 users. They do not have any fine-tuning requirements.

Why would an architect recommend a "Medium" configuration over a "Large" configuration for this customer?

- A. The Large configuration does not support Retrieval-Augmented Generation (RAG).
- B. The L40S GPUs in the Medium configuration are more cost-effective and power-efficient for this specific inference-heavy workload.
- C. The Medium configuration is the only one that includes the HPE AI Essentials software.
- D. The Medium configuration has more storage capacity than the Large configuration.

Answer: B