KillTest

Higher Quality, Better Service!

Q&A

http://www.killtest.com

We offer free update service for one year.

Exam : **HPE3-CL12**

Title : HPE Artificial Intelligence (AI)

Version : DEMO

1.A cybersecurity team is implementing a machine learning algorithm to group network traffic and identify anomalous patterns. The historical data provided to the algorithm does not contain any labels indicating whether previous traffic was malicious or benign.
Which type of machine learning is being used?
A. Supervised Learning
B. Unsupervised Learning
C. Reinforcement Learning
D. Transfer Learning
Answer: B
Explanation:
Unsupervised learning is used when the training dataset is unlabelled. The algorithm must discover inherent groupings or anomalies (clustering) in the data on its own, which is standard practice for zero-day threat detection.

2.A healthcare provider requires a robust platform to develop deep learning models for analyzing patient records. Due to strict data sovereignty and compliance regulations, the data absolutely cannot leave the hospital's on-premises data center. However, the IT director requires the automated provisioning, self-service portals, and pay-per-use financial model typically found in public clouds.
Which HPE solution should be recommended?
A. Deploying the workloads on a public multi-tenant SaaS ML platform.
B. HPE GreenLake for Private Cloud AI.
C. Purchasing standalone desktop workstations for each data scientist.
D. HPE StoreOnce with Catalyst enabled.
Answer: B
Explanation:
HPE GreenLake for Private Cloud AI provides a fully managed, consumption-based hybrid cloud experience directly within the customer's on-premises environment. This addresses the constraint of keeping data local for compliance while delivering a cloud-like operational model.

3.An enterprise wants to build a generative AI chatbot to assist engineers with troubleshooting. The equipment manuals change weekly. The enterprise requires the chatbot to provide answers based strictly on the latest documentation without hallucinating.
Which architectural approach is the most efficient?
A. Performing full-parameter fine-tuning on the LLM every weekend.
B. Utilizing a Retrieval-Augmented Generation (RAG) architecture.
C. Applying model quantization to compress the weekly manuals.
D. Relying exclusively on zero-shot prompt engineering.
Answer: B
Explanation:
For data that changes frequently, fine-tuning is computationally expensive and slow. RAG separates the knowledge base from the model by searching a vector database for the most up-to-date documents and injecting them into the LLM's prompt at runtime, ensuring factual and current responses.

4.A global manufacturing company has machine log data distributed across multiple regional data

centers and edge locations. The data science team is manually copying massive datasets to a central server for model training, violating data governance and consuming WAN bandwidth. They need a solution that provides a unified global namespace to access this distributed data in place.
Which HPE software solution addresses this constraint?
A. HPE OneView
B. HPE InfoSight
C. HPE Ezmeral Data Fabric
D. HPE Compute Ops Management
Answer: C
Explanation:
HPE Ezmeral Data Fabric unifies distributed data across edge, on-premises, and cloud environments into a single logical data store. It allows data scientists to access and analyze data where it resides, eliminating the need to physically move massive datasets.

5.A financial institution deployed a machine learning model to predict housing prices. The model performed well initially, but after a major shift in national interest rates, the model's accuracy dropped significantly, even though the infrastructure is functioning normally.
What MLOps phenomenon has occurred?
A. Catastrophic Forgetting
B. Data/Concept Drift
C. Hardware Throttling
D. Overfitting
Answer: B
Explanation:
Concept drift (or data drift) occurs when the statistical relationship between the input variables and the target variable changes in the real world after deployment. The model's underlying assumptions are no longer valid for the new economic environment, requiring retraining.