

Pass Guaranteed HP - HPE2-B08 - HPE Private Cloud AI Solutions Authoritative Exam Sims



The HPE Private Cloud AI Solutions exam dumps are designed efficiently and pointedly, so that users can check their learning effects in a timely manner after completing a section. Good practice on the success rate of HPE2-B08 quiz guide is not fully indicate that you have mastered knowledge is skilled, therefore, the HPE2-B08 test material let the user consolidate learning content as many times as possible, although the practice seems very boring, but it can achieve the result of good consolidate knowledge.

Many customers may doubt the quality of our HP HPE2-B08 learning quiz since they haven't tried them. But our HPE2-B08 training engine is reliable. What you have learnt on our HPE Private Cloud AI Solutions HPE2-B08 Exam Materials are going through special selection. The core knowledge of the real exam is significant.

>> HPE2-B08 Exam Sims <<

HP HPE2-B08 Exam Topics Pdf - HPE2-B08 Free Download

A second format is a HP HPE2-B08 web-based practice exam that can take for self-assessment. However, it differs from desktop-based HPE2-B08 practice exam software as it can be taken via any browser, including Chrome, Firefox, Safari, and Opera. This HP HPE2-B08 web-based practice exam does not require any other plugins. It also includes all of the functionalities of desktop HPE2-B08 software and will assist you in passing the HPE2-B08 certification test.

HPE Private Cloud AI Solutions Sample Questions (Q33-Q38):

NEW QUESTION # 33

An architect is in a discovery call with a customer who describes their project: "Our primary goal is to take our massive, proprietary dataset of chemical compound interactions and continuously update our foundational AI model's internal parameters to create a new, specialized model for drug discovery. This process runs 24/7 on a large GPU cluster." How should the architect classify this primary AI workload?

- A. Edge Computing
- B. Retrieval-Augmented Generation (RAG)
- C. AI Inferencing
- D. AI Model Training / Fine-tuning

Answer: D

NEW QUESTION # 34

What is the primary function of hidden layers in a deep learning model?

- A. To format the final output of the network into a human-readable prediction or classification.
- B. To directly receive the initial, raw input data for the model.
- C. To store the complete training dataset for real-time retrieval during inference.
- **D. To apply weights and transformations to the outputs of preceding layers, allowing the network to learn progressively more complex features.**

Answer: D

NEW QUESTION # 35

You are in a meeting with the VP of Manufacturing for a large industrial company. She describes their AI program as follows: "We have a formal Center of Excellence for AI that reports directly to the CTO. Our strategy is to scale our successful predictive maintenance models across all 25 of our global plants. We have standardized on a consistent cloud and data governance strategy for all AI projects and have clear ROI targets for reducing downtime." Based on this description, which AI maturity level best describes this customer?

- A. AI Curious
- B. AI Beginner
- C. Early AI user
- **D. AI Pro**

Answer: D

NEW QUESTION # 36

A development team reports that their custom-trained Large Language Model (LLM) is "hallucinating" - generating factually incorrect or nonsensical information, especially when asked questions outside the scope of its training data. The model was created by fine-tuning a foundation model on a large but static internal dataset. The team wants to improve the model's factual accuracy and reliability without embarking on a new, large-scale training project. Which are the most appropriate strategies to mitigate this issue? (Choose 2.)

- A. Retrain the model from scratch using a much larger and more diverse public dataset.
- B. Apply stricter content moderation and safety guardrails to the model's output.
- **C. Implement a Retrieval-Augmented Generation (RAG) framework to provide the model with verifiable, external context at inference time.**
- D. Increase the number of hidden layers in the model's architecture.
- **E. Reduce the "temperature" setting during inference to make the model's output less random and more focused.**

Answer: C,E

NEW QUESTION # 37

A customer is expanding their HPE Private Cloud AI "Medium" configuration to support a new generative AI inferencing workload. They are concerned about network congestion and latency, as the AI workload is known to generate large, sudden bursts of traffic between the compute nodes and the storage system.

The solution uses NVIDIA Spectrum SN4700M switches for the AI interconnect.

Which feature of these switches is specifically designed to handle bursty traffic and prevent packet loss in a lossless Ethernet fabric?

- **A. A fully shared buffer architecture that can dynamically absorb traffic bursts from any port.**
- B. Support for Fibre Channel over Ethernet (FCoE) encapsulation.
- C. Integrated Silicon Root of Trust to validate the switch firmware integrity.
- D. The ability to route traffic based on application-layer metadata.

Answer: A

NEW QUESTION # 38

.....

