

HPE2-B08測試 & HPE2-B08熱門證照



你還在猶豫什麼，機不可失，失不再來。現在你就可以獲得HP的HPE2-B08考題的完整本，只要你進Fast2test網站就能滿足你這個小小的欲望。你找到了最好的HPE2-B08考試培訓資料，請你放心使用我們的考題及答案，你一定會通過的。

Fast2test為你提供真實的環境中找到的真正的HP的HPE2-B08考試的準備過程，如果你是初學者或是想提高你的專業技能，Fast2test HP的HPE2-B08考古題將提供你，一步步讓你靠近你的願望，你有任何關於考試的考題及答案的問題，我們將第一時間幫助你解決，在一年之內，我們將提供免費更新。

>> HPE2-B08測試 <<

HPE2-B08熱門證照 - HPE2-B08考試

是不是還在為怎樣有把握地通過HP HPE2-B08 認證考試而煩惱？你有想過選擇一個針對性的培訓嗎？選擇好的培訓可以有效的幫助你快速鞏固關IT方面的大量知識，讓你可以為HP HPE2-B08 認證考試做好充分的準備。Fast2test的專家團隊利用自己的經驗和知識不斷努力地研究，終於開發出了關於HP HPE2-B08 認證考試的針對性的培訓資料，可以有效的幫助你為HP HPE2-B08 認證考試做好充分的準備。Fast2test提供的培訓資料將是你的最佳選擇。

最新的 Hybrid Cloud HPE2-B08 免費考試真題 (Q81-Q86):

問題 #81

A retail customer wants to implement an AI-powered recommender system to personalize product suggestions on their e-commerce website. They are an 'Early AI user' and need a full-stack solution that simplifies deployment and management.

Which HPE AI solution is most appropriate for this use case?

- A. HPE AI Services - Transformation Workshop
- B. HPE Cray systems
- C. AI-optimized HPE ProLiant DL servers deployed individually
- **D. HPE Private Cloud AI with NVIDIA**

答案：D

問題 #82

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. Implement a Retrieval-Augmented Generation (RAG) framework to provide the model with verifiable, external context at inference time.
- B. Apply stricter content moderation and safety guardrails to the model's output.
- C. Retrain the model from scratch using a much larger and more diverse public dataset.
- 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.

答案: A,E

問題 #83

A data science team has trained a deep learning model for image classification. While the model achieves 99.8% accuracy on the training dataset, its accuracy drops to only 75% on a new, unseen validation dataset.

The team provides the following training metrics:

...

- Training Epochs: 500
- Training Dataset Size: 1,000 images
- Model Parameters: 15 million
- Training Accuracy: 99.8%
- Validation Accuracy: 75.3%

...

What is the most likely cause of this performance discrepancy?

- A. The model is overfitting to the training data and cannot generalize to new data.
- B. The model is underfitting due to an insufficient number of training epochs.
- C. The learning rate used for training was set too low.
- D. The model has too few parameters to learn the features effectively.

答案: A

問題 #84

An architect is selecting the appropriate HPE Private Cloud AI configuration for a customer. The customer's primary workload is fine-tuning a 70B parameter model, a task that is extremely sensitive to compute latency and requires the highest possible performance for multi-node distributed training.

Which specific infrastructure components included in the "Large" configuration of HPE Private Cloud AI are essential for meeting this high-performance training requirement? (Select all that apply.)

- A. NVIDIA NVLink Bridge
- B. NVIDIA H100 NVL GPUs
- C. NVIDIA SN4700M switches supporting 200GbE RoCE
- D. HPE Aruba Networking CX 6300M switches
- E. NVIDIA L40S GPUs

答案: A,B,C

問題 #85

An architect uses the HPE Intelligent Configurator for a customer with 120 concurrent users for a text generation task with RAG.

The tool recommends a "Medium - Expanded (4-node)" configuration. The customer then reveals they want to use a smaller, more efficient 7B parameter model instead of the 13B model the tool defaulted to.

How will this change in model size likely affect the sizing tool's recommendation?

- A. It will have no effect, as the number of users is the primary factor.
- B. It will require the architect to switch to the "application capacity" sizing method.
- C. It will downgrade the recommendation to a "Medium - Standard (2-node)" configuration.
- D. It will upgrade the recommendation to a "Large - Standard (4-node)" configuration.

答案: C

