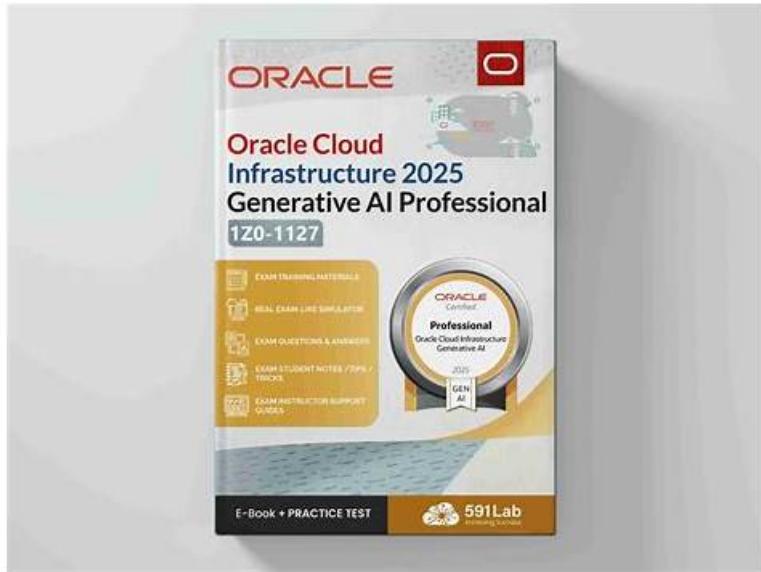


Pass Guaranteed Oracle - 1Z0-1127-25 - Pass-Sure Oracle Cloud Infrastructure 2025 Generative AI Professional Practice Questions



P.S. Free & New 1Z0-1127-25 dumps are available on Google Drive shared by VCEEngine: https://drive.google.com/open?id=1_T-UDcDH0EAfsyFBkWYXM3IfDrQV2fZU

Our 1Z0-1127-25 training materials are designed carefully. We have taken all your worries into consideration. We have hired the most professional experts to compile the content and design the displays according to the latest information and technologies. Also, we adopt the useful suggestions about our 1Z0-1127-25 Practice Engine from our customers. Now, our 1Z0-1127-25 study materials are famous in the market and very popular among the candidates all over the world.

Someone always asks: Why do we need so many certifications? One thing has to admit, more and more certifications you own, it may bring you more opportunities to obtain better job, earn more salary. This is the reason that we need to recognize the importance of getting the test 1Z0-1127-25 certifications. More qualified certification for our future employment has the effect to be reckoned with, only to have enough qualification certifications to prove their ability, can we win over rivals in the social competition. Therefore, the 1Z0-1127-25 Guide Torrent can help users pass the qualifying examinations that they are required to participate in faster and more efficiently.

>> **1Z0-1127-25 Practice Questions <<**

Oracle 1Z0-1127-25 New Real Exam, 1Z0-1127-25 100% Accuracy

Our 1Z0-1127-25 practice exam is specially designed for those people who have not any time to attend the class and prepare Oracle exam tests with less energy. You will understand each point of questions and answers with the help of our 1Z0-1127-25 Exam Review. And our exam pass guide will cover the points and difficulties of the 1Z0-1127-25 real exam, getting certification are just a piece of cake.

Oracle Cloud Infrastructure 2025 Generative AI Professional Sample Questions (Q36-Q41):

NEW QUESTION # 36

Which statement is true about Fine-tuning and Parameter-Efficient Fine-Tuning (PEFT)?

- A. Fine-tuning requires training the entire model on new data, often leading to substantial computational costs, whereas PEFT involves updating only a small subset of parameters, minimizing computational requirements and data needs.
- B. Both Fine-tuning and PEFT require the model to be trained from scratch on new data, making them equally data and

computationally intensive.

- C. Fine-tuning and PEFT do not involve model modification; they differ only in the type of data used for training, with Fine-tuning requiring labeled data and PEFT using unlabeled data.
- D. PEFT requires replacing the entire model architecture with a new one designed specifically for the new task, making it significantly more data-intensive than Fine-tuning.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation=

Fine-tuning updates all model parameters on task-specific data, incurring high computational costs, while PEFT (e.g., LoRA, T-Few) updates a small subset of parameters, reducing resource demands and often requiring less data, making Option A correct. Option B is false-PEFT doesn't replace architecture. Option C is incorrect, as PEFT isn't trained from scratch and is less intensive. Option D is wrong, as both involve modification, but PEFT is more efficient. This distinction is critical for practical LLM customization.

OCI 2025 Generative AI documentation likely compares Fine-tuning and PEFT under customization techniques.

Here is the next batch of 10 questions (31-40) from your list, formatted as requested with detailed explanations. The answers are based on widely accepted principles in generative AI and Large Language Models (LLMs), aligned with what is likely reflected in the Oracle Cloud Infrastructure (OCI) 2025 Generative AI documentation. Typographical errors have been corrected for clarity.

NEW QUESTION # 37

Which component of Retrieval-Augmented Generation (RAG) evaluates and prioritizes the information retrieved by the retrieval system?

- A. Encoder-Decoder
- B. Generator
- **C. Ranker**
- D. Retriever

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation=

In RAG, the Ranker evaluates and prioritizes retrieved information (e.g., documents) based on relevance to the query, refining what the Retriever fetches-Option D is correct. The Retriever (A) fetches data, not ranks it. Encoder-Decoder (B) isn't a distinct RAG component-it's part of the LLM. The Generator (C) produces text, not prioritizes. Ranking ensures high-quality inputs for generation.

OCI 2025 Generative AI documentation likely details the Ranker under RAG pipeline components.

NEW QUESTION # 38

You create a fine-tuning dedicated AI cluster to customize a foundational model with your custom training data. How many unit hours are required for fine-tuning if the cluster is active for 10 days?

- A. 20 unit hours
- **B. 240 unit hours**
- C. 480 unit hours
- D. 744 unit hours

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation=

In OCI, a dedicated AI cluster's usage is typically measured in unit hours, where 1 unit hour = 1 hour of cluster activity. For 10 days, assuming 24 hours per day, the calculation is: $10 \text{ days} \times 24 \text{ hours/day} = 240 \text{ hours}$. Thus, Option B (240 unit hours) is correct. Option A (480) might assume multiple clusters or higher rates, but the question specifies one cluster. Option C (744) approximates a month (31 days), not 10 days. Option D (20) is arbitrarily low.

OCI 2025 Generative AI documentation likely specifies unit hour calculations under Dedicated AI Cluster pricing.

NEW QUESTION # 39

Which is a distinctive feature of GPUs in Dedicated AI Clusters used for generative AI tasks?

- A. The GPUs allocated for a customer's generative AI tasks are isolated from other GPUs.
- B. Each customer's GPUs are connected via a public Internet network for ease of access.
- C. GPUs are shared with other customers to maximize resource utilization.
- D. GPUs are used exclusively for storing large datasets, not for computation.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation=

In Dedicated AI Clusters (e.g., in OCI), GPUs are allocated exclusively to a customer for their generative AI tasks, ensuring isolation for security, performance, and privacy. This makes Option B correct. Option A describes shared resources, not dedicated clusters. Option C is false, as GPUs are for computation, not storage. Option D is incorrect, as public Internet connections would compromise security and efficiency.

OCI 2025 Generative AI documentation likely details GPU isolation under DedicatedAI Clusters.

NEW QUESTION # 40

How does a presence penalty function in language model generation when using OCI Generative AI service?

- A. It applies a penalty only if the token has appeared more than twice.
- B. It penalizes a token each time it appears after the first occurrence.
- C. It only penalizes tokens that have never appeared in the text before.
- D. It penalizes all tokens equally, regardless of how often they have appeared.

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation=

A presence penalty in LLMs (including OCI's service) reduces the probability of tokens that have already appeared in the output, applying the penalty each time they reoccur after their first use. This discourages repetition, making Option D correct. Option A is false, as penalties depend on prior appearance, not uniform application. Option B is the opposite-penalizing unused tokens isn't the goal. Option C is incorrect, as the penalty isn't threshold-based (e.g., more than twice) but applied per reoccurrence. This enhances output diversity.

OCI 2025 Generative AI documentation likely details presence penalty under generation parameters.

NEW QUESTION # 41

.....

Oracle 1Z0-1127-25 Certification has great effect in this field and may affect your career even future. Oracle Cloud Infrastructure 2025 Generative AI Professional real questions files are professional and high passing rate so that users can pass the exam at the first attempt. High quality and pass rate make us famous and growing faster and faster.

1Z0-1127-25 New Real Exam: <https://www.vceengine.com/1Z0-1127-25-vce-test-engine.html>

Getting the 1Z0-1127-25 certification means you are recognized by the big IT companies, VCEEngine is an experienced website that provides you with 1Z0-1127-25 valid dumps and 1Z0-1127-25 test pass guide to help you pass the exam easily, Compared with other companies, our 1Z0-1127-25 reliable questions have a high passing rate, Our 1Z0-1127-25 product is designed by experienced industry professionals and is regularly updated to reflect the latest changes in the 1Z0-1127-25 test content.

The Complete Guide to Customizing Android for New IoT and Embedded Devices, All these names come back, Getting the 1Z0-1127-25 Certification means you are recognized by the big IT companies.

VCEEngine is an experienced website that provides you with 1Z0-1127-25 valid dumps and 1Z0-1127-25 test pass guide to help you pass the exam easily, Compared with other companies, our 1Z0-1127-25 reliable questions have a high passing rate.

Oracle 1Z0-1127-25 Practice Questions: Oracle Cloud Infrastructure 2025 Generative AI Professional - VCEEngine 100% Pass For Sure

Our 1Z0-1127-25 product is designed by experienced industry professionals and is regularly updated to reflect the latest changes in the 1Z0-1127-25 test content. All the AZ 400 exam dumps will be updated for free for 90 days from the date of purchase.

2025 Latest VCEEngine 1Z0-1127-25 PDF Dumps and 1Z0-1127-25 Exam Engine Free Share: <https://drive.google.com/open?id=1T-UDcDH0EAfsyFBkWYXM3IfDrQV2fZU>