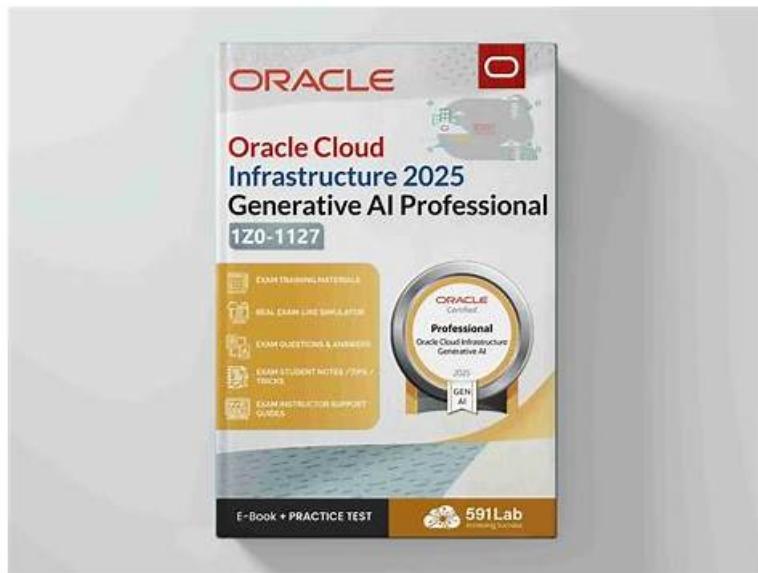


1Z0-1127-25 Visual Cert Test - 100% Pass Quiz 1Z0-1127-25 Oracle Cloud Infrastructure 2025 Generative AI Professional First-grade Pass Leader Dumps



2026 Latest 2Pass4sure 1Z0-1127-25 PDF Dumps and 1Z0-1127-25 Exam Engine Free Share: https://drive.google.com/open?id=1wjskhmlYaEneYwO0LzdZO_6KH_uLKpW

A few crops of practice materials are emerging in the market these days, with undecided quality to judge from customers' perspective. If you choose the wrong 1Z0-1127-25 practice material, it will be a grave mistake. Their behavior has not been strictly ethical and irresponsible to you, which we will never do. We know making progress and getting the certificate of 1Z0-1127-25 Training Materials will be a matter of course with the most professional experts in command of the newest and the most accurate knowledge in it. That's why our Oracle Cloud Infrastructure 2025 Generative AI Professional exam prep has taken up a large part of market.

Oracle 1Z0-1127-25 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none">Using OCI Generative AI Service: This section evaluates the expertise of Cloud AI Specialists and Solution Architects in utilizing Oracle Cloud Infrastructure (OCI) Generative AI services. It includes understanding pre-trained foundational models for chat and embedding, creating dedicated AI clusters for fine-tuning and inference, and deploying model endpoints for real-time inference. The section also explores OCI's security architecture for generative AI and emphasizes responsible AI practices.
Topic 2	<ul style="list-style-type: none">Fundamentals of Large Language Models (LLMs): This section of the exam measures the skills of AI Engineers and Data Scientists in understanding the core principles of large language models. It covers LLM architectures, including transformer-based models, and explains how to design and use prompts effectively. The section also focuses on fine-tuning LLMs for specific tasks and introduces concepts related to code models, multi-modal capabilities, and language agents.
Topic 3	<ul style="list-style-type: none">Implement RAG Using OCI Generative AI Service: This section tests the knowledge of Knowledge Engineers and Database Specialists in implementing Retrieval-Augmented Generation (RAG) workflows using OCI Generative AI services. It covers integrating LangChain with Oracle Database 23ai, document processing techniques like chunking and embedding, storing indexed chunks in Oracle Database 23ai, performing similarity searches, and generating responses using OCI Generative AI.

Topic 4	<ul style="list-style-type: none"> Using OCI Generative AI RAG Agents Service: This domain measures the skills of Conversational AI Developers and AI Application Architects in creating and managing RAG agents using OCI Generative AI services. It includes building knowledge bases, deploying agents as chatbots, and invoking deployed RAG agents for interactive use cases. The focus is on leveraging generative AI to create intelligent conversational systems.
---------	--

>> 1Z0-1127-25 Visual Cert Test <<

Pass Leader 1Z0-1127-25 Dumps | Practice 1Z0-1127-25 Online

Reliable 1Z0-1127-25 1Z0-1127-25 exam questions pdf, exam questions answers and latest test book can help customer success in their field. Oracle offers 365 days updates. Customers can download Latest 1Z0-1127-25 Exam Questions pdf and exam book. And Oracle Cloud Infrastructure 2025 Generative AI Professional 1Z0-1127-25fee is affordable. It is now time to begin your preparation by downloading the free demo of Oracle Cloud Infrastructure 2025 Generative AI Professional 1Z0-1127-25 Exam Dumps.

Oracle Cloud Infrastructure 2025 Generative AI Professional Sample Questions (Q81-Q86):

NEW QUESTION # 81

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

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

Answer: D

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 # 82

Which statement accurately reflects the differences between these approaches in terms of the number of parameters modified and the type of data used?

- A. Fine-tuning and continuous pretraining both modify all parameters and use labeled, task-specific data.
- B. Parameter Efficient Fine-Tuning and Soft Prompting modify all parameters of the model using unlabeled data.
- C. Soft Prompting and continuous pretraining are both methods that require no modification to the original parameters of the model.
- D. Fine-tuning modifies all parameters using labeled, task-specific data, whereas Parameter Efficient Fine-Tuning updates a few, new parameters also with labeled, task-specific data.**

Answer: D

Explanation:

Comprehensive and Detailed In-Depth Explanation=

Fine-tuning typically involves updating all parameters of an LLM using labeled, task-specific data to adapt it to a specific task, which is computationally expensive. Parameter Efficient Fine-Tuning (PEFT), such as methods like LoRA (Low-Rank Adaptation), updates only a small subset of parameters (often newly added ones) while still using labeled, task-specific data, making it more efficient. Option C correctly captures this distinction. Option A is wrong because continuous pretraining uses unlabeled data and isn't

task-specific. Option B is incorrect as PEFT and Soft Prompting don't modify all parameters, and Soft Prompting typically uses labeled examples indirectly. Option D is inaccurate because continuous pretraining modifies parameters, while SoftPrompting doesn't.

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

NEW QUESTION # 83

In the simplified workflow for managing and querying vector data, what is the role of indexing?

- A. To compress vector data for minimized storage usage
- B. To convert vectors into a non-indexed format for easier retrieval
- **C. To map vectors to a data structure for faster searching, enabling efficient retrieval**
- D. To categorize vectors based on their originating data type (text, images, audio)

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation=

Indexing in vector databases maps high-dimensional vectors to a data structure (e.g., HNSW, Annoy) to enable fast, efficient similarity searches, critical for real-time retrieval in LLMs. This makes Option B correct. Option A is backwards-indexing organizes, not de-indexes. Option C (compression) is a side benefit, not the primary role. Option D (categorization) isn't indexing's purpose—it's about search efficiency. Indexing powers scalable vector queries.

OCI 2025 Generative AI documentation likely explains indexing under vector database operations.

NEW QUESTION # 84

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 hours?

- A. 30 unit hours
- B. 40 unit hours
- C. 25 unit hours
- **D. 20 unit hours**

Answer: D

Explanation:

Comprehensive and Detailed In-Depth Explanation=

In OCI, unit hours typically equal actual hours of cluster activity unless specified otherwise (e.g., per GPU scaling). For 10 hours of activity, it's $10 \text{ hours} \times 1 \text{ unit/hour} = 10 \text{ unit hours}$, but options suggest a multiplier (common in cloud pricing). Assuming a standard 2-unit/hour rate (e.g., for GPU clusters), it's $10 \times 2 = 20 \text{ unit hours}$. Option C fits best. Options A, B, and D imply inconsistent rates (2.5, 4, 3).

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

NEW QUESTION # 85

What does a higher number assigned to a token signify in the "Show Likelihoods" feature of the language model token generation?

- A. The token will be the only one considered in the next generation step.
- B. The token is less likely to follow the current token.
- **C. The token is more likely to follow the current token.**
- D. The token is unrelated to the current token and will not be used.

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation=

In "Show Likelihoods," a higher number (probability score) indicates a token's greater likelihood of following the current token, reflecting the model's prediction confidence—Option B is correct. Option A (less likely) is the opposite. Option C (unrelated) misinterprets likelihood ties tokens contextually. Option D (only one) assumes greedy decoding, not the feature's purpose. This helps users understand model preferences.

OCI 2025 Generative AI documentation likely explains "Show Likelihoods" under token generation insights.

NEW QUESTION # 86

All contents of 1Z0-1127-25 training guide are being explicit to make you have explicit understanding of this exam. Their contribution is praised for their purview is unlimited. None cryptic contents in 1Z0-1127-25 learning materials you may encounter. And our 1Z0-1127-25 Exam Questions are easy to understand and they are popular to be sold to all over the world. Just look at the comments on the website, then you will know that we have a lot of loyal customers.

Pass Leader 1Z0-1127-25 Dumps: <https://www.2pass4sure.com/Oracle-Cloud-Infrastructure/1Z0-1127-25-actual-exam-braindumps.html>

P.S. Free 2026 Oracle 1Z0-1127-25 dumps are available on Google Drive shared by 2Pass4sure: https://drive.google.com/open?id=1wjskhmLYaEneYwO0LzjdZO_6KH_uLKpW