

Hot 1Z0-1127-25 Reliable Test Cost 100% Pass | Pass-Sure 1Z0-1127-25 Latest Braindumps Book: Oracle Cloud Infrastructure 2025 Generative AI Professional



What's more, part of that ExamBoosts 1Z0-1127-25 dumps now are free: <https://drive.google.com/open?id=1wOa1MinSHt7MSBCEr9ZJBCV7WJ4Palni>

As is known to us, our company has promised that the 1Z0-1127-25 exam braindumps from our company will provide more than 99% pass guarantee for all people who try their best to prepare for the exam. If you are preparing for the exam by the guidance of the 1Z0-1127-25 study practice question from our company and take it into consideration seriously, you will absolutely pass the exam and get the related certification. So do not hesitate and hurry to buy our study materials.

Oracle 1Z0-1127-25 Exam Syllabus Topics:

Topic	Details
Topic 1	<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.
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 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.

Free PDF Quiz 2026 Newest Oracle 1Z0-1127-25 Reliable Test Cost

ExamBoosts web-based practice exam is compatible with all browsers and operating systems. Whereas the 1Z0-1127-25 PDF file is concerned this file is the collection of real, valid, and updated Oracle 1Z0-1127-25 exam questions. You can use the Oracle 1Z0-1127-25 PdfFormat on your desktop computer, laptop, tabs, or even on your smartphone and start Oracle Cloud Infrastructure 2025 Generative AI Professional (1Z0-1127-25) exam questions preparation anytime and anywhere.

Oracle Cloud Infrastructure 2025 Generative AI Professional Sample Questions (Q21-Q26):

NEW QUESTION # 21

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

Answer: D

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 days × 24 hours/day = 240 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 # 22

Which statement is true about the "Top p" parameter of the OCI Generative AI Generation models?

- **A. "Top p" limits token selection based on the sum of their probabilities.**
- B. "Top p" selects tokens from the "Top k" tokens sorted by probability.
- C. "Top p" determines the maximum number of tokens per response.
- D. "Top p" assigns penalties to frequently occurring tokens.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation=

"Top p" (nucleus sampling) selects tokens whose cumulative probability exceeds a threshold (p), limiting the pool to the smallest set meeting this sum, enhancing diversity-Option C is correct. Option A confuses it with "Top k." Option B (penalties) is unrelated.

Option D (max tokens) is a different parameter. Top p balances randomness and coherence.

OCI 2025 Generative AI documentation likely explains "Top p" under sampling methods.

Here is the next batch of 10 questions (81-90) 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 # 23

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

Answer: C

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 DedicatedAI Cluster pricing.

NEW QUESTION # 24

Which statement best describes the role of encoder and decoder models in natural language processing?

- A. Encoder models take a sequence of words and predict the next word in the sequence, whereas decoder models convert a sequence of words into a numerical representation.
- B. Encoder models are used only for numerical calculations, whereas decoder models are used to interpret the calculated numerical values back into text.
- C. Encoder models and decoder models both convert sequences of words into vector representations without generating new text.
- **D. Encoder models convert a sequence of words into a vector representation, and decoder models take this vector representation to generate a sequence of words.**

Answer: D

Explanation:

Comprehensive and Detailed In-Depth Explanation=

In NLP (e.g., transformers), encoders convert input text into a vector representation (encoding meaning), while decoders generate text from such vectors (e.g., in translation or generation). This makes Option C correct. Option A is false-decoders generate text. Option B reverses roles-encoders don't predict next words, and decoders don't encode. Option D oversimplifies-encoders handle text, not just numbers. This is the foundation of seq2seq models.

OCI 2025 Generative AI documentation likely explains encoder-decoder roles under model architecture.

NEW QUESTION # 25

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

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

Answer: D

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 # 26

.....

Thus, we come forward to assist them in cracking the Oracle 1Z0-1127-25 examination. Don't postpone purchasing Oracle 1Z0-1127-25 exam dumps to pass the crucial examination. ExamBoosts study material is available in three versions: Oracle 1Z0-1127-25 Pdf Dumps, desktop practice exam software, and a web-based Oracle 1Z0-1127-25 practice test.

