

Exam 1Z0-1127-25 Sample & Exam 1Z0-1127-25 Topics



DOWNLOAD the newest Lead2PassExam 1Z0-1127-25 PDF dumps from Cloud Storage for free:
https://drive.google.com/open?id=1Q0QqK6gGuemhcRoGVryhl3_HrpmKxeLj

The pass rate is 98.65% for 1Z0-1127-25 learning materials, and if you choose us, we can ensure you that you can pass the exam just one time. In addition, 1Z0-1127-25 exam dumps are edited by skilled experts, who have the professional knowledge for 1Z0-1127-25 exam dumps, therefore the quality and accuracy can be guaranteed. We also pass guarantee and money back guarantee for 1Z0-1127-25 Learning Materials, and if you fail to pass the exam, we will give you full refund, and no other questions will be asked.

Oracle 1Z0-1127-25 Exam Syllabus Topics:

Topic	Details
Topic 1	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 2	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 3	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 4	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.
---------	--

>> Exam 1Z0-1127-25 Sample <<

High efficient 1Z0-1127-25 Guide Torrent Practice Materials: Oracle Cloud Infrastructure 2025 Generative AI Professional - Lead2PassExam

The Lead2PassExam offers three formats for applicants to practice and prepare for the 1Z0-1127-25 exam as per their needs. The pdf format of Lead2PassExam is portable and can be used on laptops, tablets, and smartphones. Print real Oracle Cloud Infrastructure 2025 Generative AI Professional (1Z0-1127-25) exam questions in our PDF file. The pdf is user-friendly and accessible on any smart device, allowing applicants to study from anywhere at any time.

Oracle Cloud Infrastructure 2025 Generative AI Professional Sample Questions (Q74-Q79):

NEW QUESTION # 74

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" determines the maximum number of tokens per response.
- C. "Top p" selects tokens from the "Top k" tokens sorted by probability.
- 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 # 75

Which is a distinguishing feature of "Parameter-Efficient Fine-Tuning (PEFT)" as opposed to classic "Fine-tuning" in Large Language Model training?

- A. PEFT modifies all parameters and uses unlabeled, task-agnostic data.
- B. PEFT involves only a few or new parameters and uses labeled, task-specific data.
- C. PEFT does not modify any parameters but uses soft prompting with unlabeled data.
- D. PEFT modifies all parameters and is typically used when no training data exists.

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation=

PEFT (e.g., LoRA, T-Few) updates a small subset of parameters (often new ones) using labeled, task-specific data, unlike classic fine-tuning, which updates all parameters-Option A is correct. Option B reverses PEFT's efficiency. Option C (no modification) fits soft prompting, not all PEFT. Option D (all parameters) mimics classic fine-tuning. PEFT reduces resource demands.

OCI 2025 Generative AI documentation likely contrasts PEFT and fine-tuning under customization methods.

NEW QUESTION # 76

In which scenario is soft prompting appropriate compared to other training styles?

- A. When the model needs to be adapted to perform well in a domain on which it was not originally trained
- B. When the model requires continued pretraining on unlabeled data
- **C. When there is a need to add learnable parameters to a Large Language Model (LLM) without task-specific training**
- D. When there is a significant amount of labeled, task-specific data available

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation=

Soft prompting adds trainable parameters (soft prompts) to adapt an LLM without retraining its core weights, ideal for low-resource customization without task-specific data. This makes Option C correct. Option A suits fine-tuning. Option B may require more than soft prompting (e.g., domain fine-tuning). Option D describes pretraining, not soft prompting. Soft prompting is efficient for specific adaptations.

OCI 2025 Generative AI documentation likely discusses soft prompting under PEFT methods.

NEW QUESTION # 77

An AI development company is working on an AI-assisted chatbot for a customer, which happens to be an online retail company. The goal is to create an assistant that can best answer queries regarding the company policies as well as retain the chat history throughout a session. Considering the capabilities, which type of model would be the best?

- A. A pre-trained LLM model from Cohere or OpenAI.
- **B. An LLM enhanced with Retrieval-Augmented Generation (RAG) for dynamic information retrieval and response generation.**
- C. An LLM dedicated to generating text responses without external data integration.
- D. A keyword search-based AI that responds based on specific keywords identified in customer queries.

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation=

For a chatbot needing to answer policy queries (requiring up-to-date, specific data) and retain chat history (context awareness), an LLM with RAG is ideal. RAG integrates external data (e.g., policy documents) via retrieval and supports memory for session-long context, making Option B correct. Option A (keyword search) lacks reasoning and context retention. Option C (standalone LLM) can't dynamically fetch policy data. Option D (pre-trained LLM) is too vague and lacks RAG's capabilities. RAG meets both requirements effectively.

OCI 2025 Generative AI documentation likely highlights RAG for dynamic, context-aware applications.

NEW QUESTION # 78

How does the integration of a vector database into Retrieval-Augmented Generation (RAG)-based Large Language Models (LLMs) fundamentally alter their responses?

- **A. It shifts the basis of their responses from pretrained internal knowledge to real-time data retrieval.**
- B. It enables them to bypass the need for pretraining on large text corpora.
- C. It transforms their architecture from a neural network to a traditional database system.
- D. It limits their ability to understand and generate natural language.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation=

RAG integrates vector databases to retrieve real-time external data, augmenting the LLM's pretrained knowledge with current, specific information, shifting response generation to a hybrid approach-Option B is correct. Option A is false-architecture remains neural; only data sourcing changes. Option C is incorrect-pretraining is still required; RAG enhances it. Option D is wrong-RAG improves, not limits, generation. This shift enables more accurate, up-to-date responses.

OCI 2025 Generative AI documentation likely details RAG's impact under response generation enhancements.

• • • • •

Exam 1Z0-1127-25 Topics: <https://www.lead2passexam.com/Oracle/valid-1Z0-1127-25-exam-dumps.html>

- 2025 Latest Lead2PassExam 1Z0-1127-25 PDF Dumps and 1Z0-1127-25 Exam Engine Free Share:
https://drive.google.com/open?id=1Q0QqK6gGuemhcRoGVryhl3_HrpmKxeLj

https://drive.google.com/open?id=1Q0QqK6gGuemhcRoGVryhl3_HrpmKxeLj