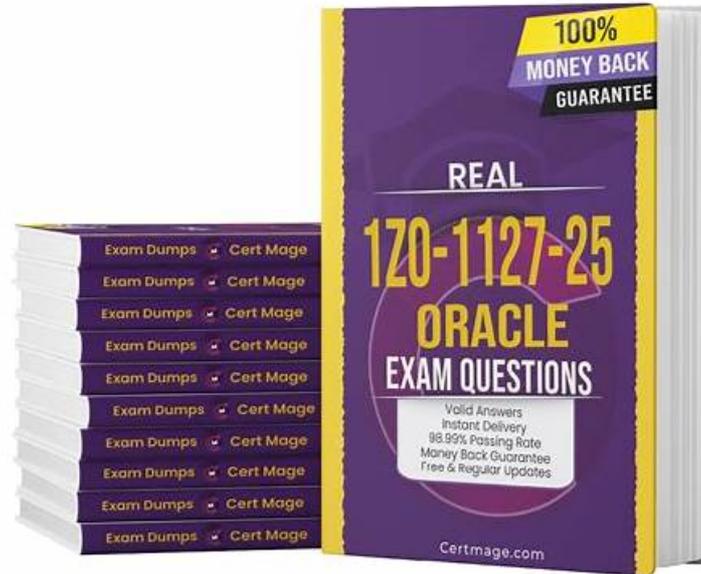


Free PDF Oracle - 1Z0-1127-25 Pass-Sure Exam Dumps Pdf



BONUS!!! Download part of ExamTorrent 1Z0-1127-25 dumps for free: https://drive.google.com/open?id=12COnE3ibt9vp5JbT6LtmD8Zi_1huhNSF

Oracle Cloud Infrastructure 2025 Generative AI Professional study questions provide free trial service for consumers. If you are interested in 1Z0-1127-25 exam material, you only need to enter our official website, and you can immediately download and experience our trial PDF file for free. Through the trial you will have different learning experience, you will find that what we say is not a lie, and you will immediately fall in love with our products. As a key to the success of your life, the benefits that 1Z0-1127-25 Exam Guide can bring you are not measured by money. 1Z0-1127-25 exam guide can not only help you pass the exam, but also help you master a new set of learning methods and teach you how to study efficiently, 1Z0-1127-25 exam material will lead you to success.

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"> 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 3	<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 4	<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.

>> Exam Dumps 1Z0-1127-25 Pdf <<

Real 1Z0-1127-25 Dumps Free & 1Z0-1127-25 Latest Test Sample

This is a Oracle 1Z0-1127-25 practice exam software for Windows computers. This 1Z0-1127-25 practice test will be similar to the actual Oracle Cloud Infrastructure 2025 Generative AI Professional (1Z0-1127-25) exam. If user wish to test the Oracle 1Z0-1127-25 study material before joining ExamTorrent, they may do so with a free sample trial. This 1Z0-1127-25 Exam simulation software can be readily installed on Windows-based computers and laptops. Since it is desktop-based Oracle 1Z0-1127-25 practice exam software, it is not necessary to connect to the internet to use it.

Oracle Cloud Infrastructure 2025 Generative AI Professional Sample Questions (Q48-Q53):

NEW QUESTION # 48

Which statement describes the difference between "Top k" and "Top p" in selecting the next token in the OCI Generative AI Generation models?

- A. "Top k" considers the sum of probabilities of the top tokens, whereas "Top p" selects from the "Top k" tokens sorted by probability.
- B. "Top k" and "Top p" both select from the same set of tokens but use different methods to prioritize them based on frequency.
- C. "Top k" and "Top p" are identical in their approach to token selection but differ in their application of penalties to tokens.
- D. "Top k" selects the next token based on its position in the list of probable tokens, whereas "Top p" selects based on the cumulative probability of the top tokens.

Answer: D

Explanation:

Comprehensive and Detailed In-Depth Explanation=

"Top k" sampling selects from the k most probable tokens, based on their ranked position, while "Top p" (nucleus sampling) selects from tokens whose cumulative probability exceeds p, focusing on a dynamic probability mass-Option B is correct. Option A is false-they differ in selection, not penalties. Option C reverses definitions. Option D (frequency) is incorrect-both use probability, not frequency. This distinction affects diversity.

OCI 2025 Generative AI documentation likely contrasts Top k and Top p under sampling methods.

NEW QUESTION # 49

Given the following code:

```
chain = prompt | llm
```

Which statement is true about LangChain Expression Language (LCEL)?

- A. LCEL is a legacy method for creating chains in LangChain.
- B. LCEL is an older Python library for building Large Language Models.
- C. LCEL is a programming language used to write documentation for LangChain.
- D. LCEL is a declarative and preferred way to compose chains together.

Answer: D

Explanation:

Comprehensive and Detailed In-Depth Explanation=

LangChain Expression Language (LCEL) is a declarative syntax (e.g., using | to pipe components) for composing chains in LangChain, combining prompts, LLMs, and other elements efficiently-Option C is correct. Option A is false-LCEL isn't for documentation. Option B is incorrect-it's current, not legacy; traditional Python classes are older. Option D is wrong-LCEL is part of LangChain, not a standalone LLM library. LCEL simplifies chain design.

OCI 2025 Generative AI documentation likely highlights LCEL under LangChain chaincomposition.

NEW QUESTION # 50

What is the purpose of Retrievers in LangChain?

- A. To break down complex tasks into smaller steps
- **B. To retrieve relevant information from knowledge bases**
- C. To train Large Language Models
- D. To combine multiple components into a single pipeline

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation=

Retrievers in LangChain fetch relevant information (e.g., documents, embeddings) from external knowledge bases (like vector stores) to provide context for LLM responses, especially in RAG setups. This makes Option B correct. Option A (training) is unrelated-Retrievers operate at inference. Option C (task breakdown) pertains to prompting techniques, not retrieval. Option D (pipeline combination) describes chains, not Retrievers specifically. Retrievers enhance context awareness.

OCI 2025 Generative AI documentation likely defines Retrievers under LangChain components.

NEW QUESTION # 51

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

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

Answer: D

Explanation:

Comprehensive and Detailed In-Depth Explanation=

Soft prompting (e.g., prompt tuning) involves adding trainable parameters (soft prompts) to an LLM's input while keeping the model's weights frozen, adapting it to tasks without task-specific retraining. This is efficient when fine-tuning or large datasets aren't feasible, making Option C correct. Option A suits full fine-tuning, not soft prompting, which avoids extensive labeled data needs.

Option B could apply, but domain adaptation often requires more than soft prompting (e.g., fine-tuning). Option D describes continued pretraining, not soft prompting. Soft prompting excels in low-resource customization.

OCI 2025 Generative AI documentation likely discusses soft prompting under parameter-efficient methods.

NEW QUESTION # 52

Which role does a "model endpoint" serve in the inference workflow of the OCI Generative AI service?

- **A. Serves as a designated point for user requests and model responses**
- B. Hosts the training data for fine-tuning custom models
- C. Updates the weights of the base model during the fine-tuning process
- D. Evaluates the performance metrics of the custom models

Answer: A

Explanation:

