

2026 Oracle 1Z0-1127-25: Accurate Valid Oracle Cloud Infrastructure 2025 Generative AI Professional Test Sample



BTW, DOWNLOAD part of Dumpexams 1Z0-1127-25 dumps from Cloud Storage: https://drive.google.com/open?id=1FLqTKiqIfuWHSrCieXe_W5uqpDHyl2Ic

The exam time is coming, while you are not prepared well for 1Z0-1127-25 real test. Please do not be tense and worried, you can pass your 1Z0-1127-25 actual exam very simply and easily with Dumpexams 1Z0-1127-25 free pdf dumps. With the help of Oracle 1Z0-1127-25 free pdf practice, you can not only get high score in your actual test, but also can get more technology knowledge and be more professional.

We attract customers by our fabulous 1Z0-1127-25 certification material and high pass rate, which are the most powerful evidence to show our strength. We are so proud to tell you that according to the statistics from our customers' feedback, the pass rate among our customers who prepared for the exam with our 1Z0-1127-25 Test Guide have reached as high as 99%, which definitely ranks the top among our peers. Hence one can see that the Oracle Cloud Infrastructure 2025 Generative AI Professional learn tool compiled by our company are definitely the best choice for you.

>> Valid 1Z0-1127-25 Test Sample <<

Practice Oracle 1Z0-1127-25 Questions | 1Z0-1127-25 Latest Test Sample

Dumpexams's Oracle 1Z0-1127-25 exam training materials is the best training materials. If you are an IT staff, it will be your indispensable training materials. Do not take your future betting on tomorrow. Dumpexams's Oracle 1Z0-1127-25 exam training materials are absolutely trustworthy. We are dedicated to provide the materials to the world of the candidates who want to participate in IT exam. To get the Oracle 1Z0-1127-25 Exam Certification is the goal of many IT people & Network professionals. The pass rate of Dumpexams is incredibly high. We are committed to your success.

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"> 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 3	<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 4	<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.

Oracle Cloud Infrastructure 2025 Generative AI Professional Sample Questions (Q71-Q76):

NEW QUESTION # 71

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

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

Answer: B

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 # 72

Which statement is true about string prompt templates and their capability regarding variables?

- A. They can only support a single variable at a time.
- B. They are unable to use any variables.
- **C. They support any number of variables, including the possibility of having none.**
- D. They require a minimum of two variables to function properly.

Answer: C

Explanation:

Comprehensive and Detailed In-Depth Explanation=

String prompt templates (e.g., in LangChain) are flexible frameworks that can include zero, one, or multiple variables (placeholders) to customize prompts dynamically. They can be static (no variables) or complex (many variables), making Option C correct. Option A is too restrictive. Option B is false-variables are a core feature. Option D is incorrect, as no minimum is required. This flexibility aids prompt engineering.

OCI 2025 Generative AI documentation likely covers prompt templates under LangChain or prompt design.

NEW QUESTION # 73

What issue might arise from using small datasets with the Vanilla fine-tuning method in the OCI Generative AI service?

- A. Data Leakage
- B. Model Drift
- C. Underfitting
- **D. Overfitting**

Answer: D

Explanation:

Comprehensive and Detailed In-Depth Explanation=

Vanilla fine-tuning updates all model parameters, and with small datasets, it can overfit-memorizing the data rather than generalizing-leading to poor performance on unseen data. Option A is correct. Option B (underfitting) is unlikely with full updates-overfitting is the risk. Option C (data leakage) depends on data handling, not size. Option D (model drift) relates to deployment shifts, not training. Small datasets exacerbate overfitting in Vanilla fine-tuning.

OCI 2025 Generative AI documentation likely warns of overfitting under Vanilla fine-tuning limitations.

NEW QUESTION # 74

How are chains traditionally created in LangChain?

- **A. Using Python classes, such as LLMChain and others**
- B. Declaratively, with no coding required
- C. By using machine learning algorithms
- D. Exclusively through third-party software integrations

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation=

Traditionally, LangChain chains (e.g., LLMChain) are created using Python classes that define sequences of operations, such as calling an LLM or processing data. This programmatic approach predates LCEL's declarative style, making Option C correct.

Option A is vague and incorrect, as chains aren't ML algorithms themselves. Option B describes LCEL, not traditional methods.

Option D is false, as third-party integrations aren't required. Python classes provide structured chain building.

OCI 2025 Generative AI documentation likely contrasts traditional chains with LCEL under LangChain sections.

NEW QUESTION # 75

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

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

Answer: B

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 # 76

