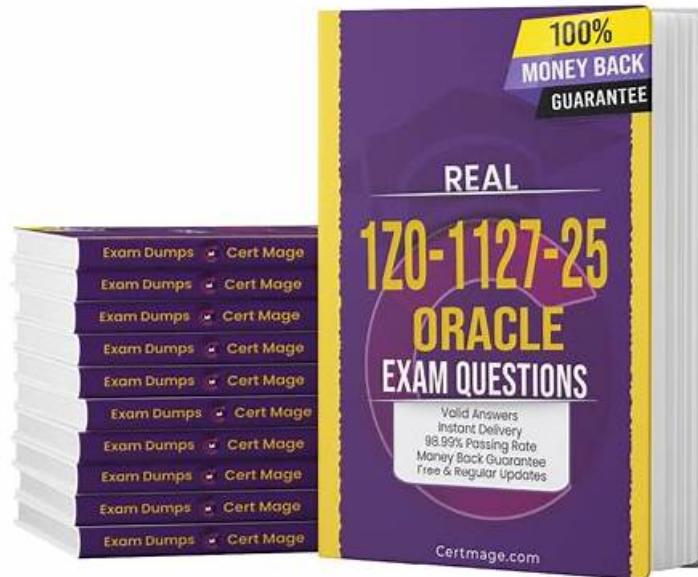


Oracle 1Z0-1127-25 Valid Exam Sample, 1Z0-1127-25 Latest Test Question



What's more, part of that DumpsTorrent 1Z0-1127-25 dumps now are free: https://drive.google.com/open?id=1YAk7wVJfiiMwzi_sPgIEspTQ6GfoCOLP

Our 1Z0-1127-25 training materials are famous for instant access to download. You can receive your downloading link and password within ten minutes, so that you can start your learning as early as possible. In order to build up your confidence for 1Z0-1127-25 exam materials, we are pass guarantee and money back guarantee, and if you fail to pass the exam, we will give you full refund. In addition, 1Z0-1127-25 test materials cover most of knowledge points for the exam, therefore you can master the major points for the exam as well as improve your professional ability in the process of learning.

Oracle 1Z0-1127-25 Exam Syllabus Topics:

Topic	Details
Topic 1	<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 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"> 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.
---------	--

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

1Z0-1127-25 Latest Test Question & Latest 1Z0-1127-25 Exam Cost

It is normally not a bad thing to pass more exams and get more certifications. In fact to a certain degree, Oracle certifications will be magic weapon for raising position and salary. Finding latest 1Z0-1127-25 valid exam questions answers is the latest and simplest method for young people to clear exam. Our exam dumps include PDF format, soft test engine and APP test engine three versions. 1Z0-1127-25 Valid Exam Questions answers will cover all learning materials of real test questions.

Oracle Cloud Infrastructure 2025 Generative AI Professional Sample Questions (Q89-Q94):

NEW QUESTION # 89

How can the concept of "Groundedness" differ from "Answer Relevance" in the context of Retrieval Augmented Generation (RAG)?

- A. Groundedness pertains to factual correctness, whereas Answer Relevance concerns query relevance.**
- B. Groundedness measures relevance to the user query, whereas Answer Relevance evaluates data integrity.
- C. Groundedness focuses on data integrity, whereas Answer Relevance emphasizes lexical diversity.
- D. Groundedness refers to contextual alignment, whereas Answer Relevance deals with syntactic accuracy.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation=

In RAG, "Groundedness" assesses whether the response is factually correct and supported by retrieved data, while "Answer Relevance" evaluates how well the response addresses the user's query. Option A captures this distinction accurately. Option B is off-groundedness isn't just contextual alignment, and relevance isn't about syntax. Option C swaps the definitions. Option D misaligns-groundedness isn't solely data integrity, and relevance isn't lexical diversity. This distinction ensures RAG outputs are both true and pertinent.

OCI 2025 Generative AI documentation likely defines these under RAG evaluation metrics.

NEW QUESTION # 90

How does a presence penalty function in language model generation?

- A. It penalizes only tokens that have never appeared in the text before.
- B. It penalizes a token each time it appears after the first occurrence.**
- C. It penalizes all tokens equally, regardless of how often they have appeared.
- D. It applies a penalty only if the token has appeared more than twice.

Answer: B

Explanation:

Comprehensive and Detailed In-Depth Explanation=

A presence penalty reduces the probability of tokens that have already appeared in the output, applying the penalty each time they reoccur after their first use, to discourage repetition. This makes Option D correct. Option A (equal penalties) ignores prior appearance. Option B is the opposite-penalizing unused tokens isn't the intent. Option C (more than twice) adds an arbitrary threshold not typically used. Presence penalty enhances output variety. OCI 2025 Generative AI documentation likely details presence penalty under generation control parameters.

NEW QUESTION # 91

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

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

Answer: C

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 Dedicated AI Clusters.

NEW QUESTION # 92

Why is normalization of vectors important before indexing in a hybrid search system?

- **A. It standardizes vector lengths for meaningful comparison using metrics such as Cosine Similarity.**
- B. It ensures that all vectors represent keywords only.
- C. It significantly reduces the size of the database.
- D. It converts all sparse vectors to dense vectors.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation=

Normalization scales vectors to unit length, ensuring comparisons (e.g., cosine similarity) reflect directional similarity, not magnitude differences, critical for hybrid search accuracy. This makes Option C correct. Option A is false—vectors represent semantics, not just keywords. Option B (size reduction) isn't the goal. Option D (sparse to dense) is unrelated—normalization adjusts length. Normalized vectors ensure fair similarity metrics.

OCI 2025 Generative AI documentation likely explains normalization under vector preprocessing.

NEW QUESTION # 93

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

- A. When the model requires continued pretraining on unlabeled data
- B. When there is a significant amount of labeled, task-specific data available
- C. When the model needs to be adapted to perform well in a domain on which it was not originally trained
- **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 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 # 94

.....

For candidates who want to evaluate and enhance their Oracle 1Z0-1127-25 Test Preparation online, the web-based practice test

is a perfect choice. You can attempt our 60 Oracle web-based practice exam whenever it suits you because it is accessible from any location with an internet connection. This Oracle Cloud Infrastructure 2025 Generative AI Professional browser-based practice exam helps you overcome exam fear as it simulates the environment of the real test.

1Z0-1127-25 Latest Test Question: <https://www.dumpstorrent.com/1Z0-1127-25-exam-dumps-torrent.html>

P.S. Free 2025 Oracle 1Z0-1127-25 dumps are available on Google Drive shared by DumpsTorrent: https://drive.google.com/open?id=1YAk7wVJfiiMwzi_sPgiEspTQ6GfoCOLP