


New 1Z0-1127-25 Test Camp, 1Z0-1127-25 Valid Exam Notes

Oracle Cloud Infrastructure 2025 Generative AI Professional - 1Z0-1127-25

Q 1. category



Q 2. Which is a distinguishing feature of "Parameter-Efficient fine-tuning (PEFT)"

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

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



Co
nfi
de
nfi

P.S. Free & New 1Z0-1127-25 dumps are available on Google Drive shared by Prep4SureReview: <https://drive.google.com/open?id=1qICd-vVIYK8hN0xGckX5ftjqSr6ZKxHV>

Fantasy can make people to come up with many good ideas, but it can not do anything. So when you thinking how to pass the Oracle 1Z0-1127-25 Exam, It's better open your computer, and click the website of Prep4SureReview, then you will see the things you want. Prep4SureReview's products have favorable prices, and have quality assurance, but also to ensure you to 100% pass the exam.

It is believe that employers nowadays are more open to learn new knowledge, as they realize that Oracle certification may be conducive to them in refreshing their life, especially in their career arena. 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.

>> New 1Z0-1127-25 Test Camp <<

Exam Questions for the Oracle 1Z0-1127-25 - Master Your Certification

Journey

This kind of polished approach is beneficial for a commendable grade in the Oracle Cloud Infrastructure 2025 Generative AI Professional (1Z0-1127-25) exam. While attempting the exam, take heed of the clock ticking, so that you manage the Oracle Cloud Infrastructure 2025 Generative AI Professional (1Z0-1127-25) questions in a time-efficient way. Even if you are completely sure of the correct answer to a question, first eliminate the incorrect ones, so that you may prevent blunders due to human error.

Oracle 1Z0-1127-25 Exam Syllabus Topics:

Topic	Details
Topic 1	<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 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">• 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 Cloud Infrastructure 2025 Generative AI Professional Sample Questions (Q71-Q76):

NEW QUESTION # 71

What does the Loss metric indicate about a model's predictions?

- **A. Loss is a measure that indicates how wrong the model's predictions are.**
- B. Loss describes the accuracy of the right predictions rather than the incorrect ones.
- C. Loss indicates how good a prediction is, and it should increase as the model improves.
- D. Loss measures the total number of predictions made by a model.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation=

Loss is a metric that quantifies the difference between a model's predictions and the actual target values, indicating how incorrect (or "wrong") the predictions are. Lower loss means better performance, making Option B correct. Option A is false-loss isn't about prediction count. Option C is incorrect-loss decreases as the model improves, not increases. Option D is wrong-loss measures overall error, not just correct predictions. Loss guides training optimization.

OCI 2025 Generative AI documentation likely defines loss under model training and evaluation metrics.

NEW QUESTION # 72

How are fine-tuned customer models stored to enable strong data privacy and security in the OCI Generative AI service?

- **A. Stored in Object Storage encrypted by default**
- B. Shared among multiple customers for efficiency
- C. Stored in Key Management service
- D. Stored in an unencrypted form in Object Storage

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation=

In OCI, fine-tuned models are stored in Object Storage, encrypted by default, ensuring privacy and security per cloud best practices-Option B is correct. Option A (shared) violates privacy. Option C (unencrypted) contradicts security standards. Option D (Key Management) stores keys, not models. Encryption protects customer data.

OCI 2025 Generative AI documentation likely details storage security under fine-tuning workflows.

NEW QUESTION # 73

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

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

Answer: D

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 # 74

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

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 # 75

Which statement is true about Fine-tuning and Parameter-Efficient Fine-Tuning (PEFT)?

- **A. Fine-tuning requires training the entire model on new data, often leading to substantial computational costs, whereas PEFT involves updating only a small subset of parameters, minimizing computational requirements and data needs.**
- B. Fine-tuning and PEFT do not involve model modification; they differ only in the type of data used for training, with Fine-tuning requiring labeled data and PEFT using unlabeled data.
- C. PEFT requires replacing the entire model architecture with a new one designed specifically for the new task, making it significantly more data-intensive than Fine-tuning.
- D. Both Fine-tuning and PEFT require the model to be trained from scratch on new data, making them equally data and computationally intensive.

Answer: A

Explanation:

Comprehensive and Detailed In-Depth Explanation=

Fine-tuning updates all model parameters on task-specific data, incurring high computational costs, while PEFT (e.g., LoRA, T-Few) updates a small subset of parameters, reducing resource demands and often requiring less data, making Option A correct. Option B is false-PEFT doesn't replace architecture. Option C is incorrect, as PEFT isn't trained from scratch and is less intensive. Option D is wrong, as both involve modification, but PEFT is more efficient. This distinction is critical for practical LLM customization.

OCI 2025 Generative AI documentation likely compares Fine-tuning and PEFT under customization techniques.

Here is the next batch of 10 questions (31-40) 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 # 76

.....

Our 1Z0-1127-25 study materials are very popular in the international market and enjoy wide praise by the people in and outside the circle. We have shaped our 1Z0-1127-25 exam braindumps into a famous and top-ranking brand and we enjoy well-deserved reputation among the clients. Our 1Z0-1127-25 Training Questions boost many outstanding and superior advantages which other same kinds of products don't have. You won't regret if you buy them!

1Z0-1127-25 Valid Exam Notes: <https://www.prep4sureview.com/1Z0-1127-25-latest-braindumps.html>

- 1Z0-1127-25 Exam Dumps Collection New Exam 1Z0-1127-25 Materials 1Z0-1127-25 Exam Overview The page for free download of ⇒ 1Z0-1127-25 ⇐ on www.verifiedumps.com will open immediately Exam 1Z0-1127-25 Simulator Fee
- 1Z0-1127-25 Exam Overview 1Z0-1127-25 Latest Exam Online 1Z0-1127-25 Latest Real Exam Search for 1Z0-1127-25 and obtain a free download on ➡ www.pdfvce.com Valid 1Z0-1127-25 Test Labs
- Training 1Z0-1127-25 Online Reliable 1Z0-1127-25 Test Online 1Z0-1127-25 Exam Overview Open website ⇒ www.vceengine.com ⇐ and search for [1Z0-1127-25] for free download 1Z0-1127-25 Reasonable Exam Price
- Free PDF Oracle - 1Z0-1127-25 - Trustable New Oracle Cloud Infrastructure 2025 Generative AI Professional Test Camp The page for free download of [1Z0-1127-25] on ▶ www.pdfvce.com ◀ will open immediately 1Z0-1127-25 Exam Overview
- Best Oracle New 1Z0-1127-25 Test Camp Professionally Researched by Oracle Certified Trainers Download 1Z0-1127-25 for free by simply entering ➡ www.vce4dumps.com website New 1Z0-1127-25 Study Plan
- 1Z0-1127-25 Exam Dumps Collection 1Z0-1127-25 Latest Real Exam New Exam 1Z0-1127-25 Materials Open [www.pdfvce.com] and search for ➡ 1Z0-1127-25 to download exam materials for free 1Z0-1127-25 Latest Exam Online
- 1Z0-1127-25 Test Guide Online New Exam 1Z0-1127-25 Materials Valid Exam 1Z0-1127-25 Practice Open www.prepawayete.com and search for [1Z0-1127-25] to download exam materials for free 1Z0-1127-25 Valid Dumps Demo
- 1Z0-1127-25 Reasonable Exam Price 1Z0-1127-25 Exam Overview ➡ 1Z0-1127-25 Latest Exam Online Download ➡ 1Z0-1127-25 for free by simply entering ➡ www.pdfvce.com website Valid 1Z0-1127-25 Exam Duration
- Here's the Quick Way to Crack Oracle 1Z0-1127-25 Certification Exam www.prepawaypdf.com is best website to obtain ⇒ 1Z0-1127-25 ⇐ for free download Training 1Z0-1127-25 Online

