

# Reliable 1Z0-1127-25 Exam Syllabus - New Study 1Z0-1127-25 Questions



P.S. Free & New 1Z0-1127-25 dumps are available on Google Drive shared by Exam4Docs: <https://drive.google.com/open?id=1Er5DsXBruf4nk07-w8mvZYoYbP3LPGtX>

We will give you full refund if you fail to pass the exam after purchasing 1Z0-1127-25 learning materials from us. We are pass guarantee and money back guarantee, and money will be returned to your payment account. We have a professional team to collect and research the latest information for 1Z0-1127-25 Exam Dumps, we can ensure you that the exam dumps you receive are the latest one we have. In order to let you know the latest information for the 1Z0-1127-25 learning materials, we offer you free update for one year, and the update version will be sent to your email automatically.

## Oracle 1Z0-1127-25 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"><li>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.</li></ul>
Topic 2	<ul style="list-style-type: none"><li>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.</li></ul>
Topic 3	<ul style="list-style-type: none"><li>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.</li></ul>

Topic 4	<ul style="list-style-type: none"> <li>• 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.</li> </ul>
---------	---

>> Reliable 1Z0-1127-25 Exam Syllabus <<

## Quiz Oracle - 1Z0-1127-25 –Newest Reliable Exam Syllabus

If you find someone around has a nice life go wild, it is because that they may have favored the use of study & work method different from normal people. 1Z0-1127-25 dumps torrent files may be the best method for candidates who are preparing for their IT exam and eager to clear exam as soon as possible. People's success lies in their good use of every change to self-improve. Our 1Z0-1127-25 Dumps Torrent files will be the best resources for your real test. If you choose our products, we will choose efficient & high-passing preparation materials.

### Oracle Cloud Infrastructure 2025 Generative AI Professional Sample Questions (Q87-Q92):

#### NEW QUESTION # 87

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

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

**Answer: B**

Explanation:

Comprehensive and Detailed In-Depth Explanation=

A "model endpoint" in OCI's inference workflow is an API or interface where users send requests and receive responses from a deployed model. Option B is correct. Option A (weight updates) occurs during fine-tuning, not inference. Option C (metrics) is for evaluation, not endpoints. Option D (training data) relates to storage, not inference. Endpoints enable real-time interaction. OCI 2025 Generative AI documentation likely describes endpoints under inference deployment.

#### NEW QUESTION # 88

When is fine-tuning an appropriate method for customizing a Large Language Model (LLM)?

- A. When the LLM does not perform well on a task and the data for prompt engineering is too large
- B. When the LLM already understands the topics necessary for text generation
- C. When the LLM requires access to the latest data for generating outputs
- D. When you want to optimize the model without any instructions

**Answer: A**

Explanation:

Comprehensive and Detailed In-Depth Explanation=

Fine-tuning is suitable when an LLM underperforms on a specific task and prompt engineering alone isn't feasible due to large, task-specific data that can't be efficiently included in prompts. This adjusts the model's weights, making Option B correct. Option A suggests no customization is needed. Option C favors RAG for latest data, not fine-tuning. Option D is vague-fine-tuning requires data and goals, not just optimization without direction. Fine-tuning excels with substantial task-specific data.

OCI 2025 Generative AI documentation likely outlines fine-tuning use cases under customization strategies.

#### NEW QUESTION # 89

What differentiates Semantic search from traditional keyword search?

- A. It relies solely on matching exact keywords in the content.
- B. It is based on the date and author of the content.
- **C. It involves understanding the intent and context of the search.**
- D. It depends on the number of times keywords appear in the content.

**Answer: C**

Explanation:

Comprehensive and Detailed In-Depth Explanation=

Semantic search uses embeddings and NLP to understand the meaning, intent, and context behind a query, rather than just matching exact keywords (as in traditional search). This enables more relevant results, even if exact terms aren't present, making Option C correct. Options A and B describe traditional keyword search mechanics. Option D is unrelated, as metadata like date or author isn't the primary focus of semantic search. Semantic search leverages vector representations for deeper understanding.

OCI 2025 Generative AI documentation likely contrasts semantic and keyword search under search or retrieval sections.

#### NEW QUESTION # 90

What does a higher number assigned to a token signify in the "Show Likelihoods" feature of the language model token generation?

- **A. The token is more likely to follow the current token.**
- B. The token is less likely to follow the current token.
- C. The token is unrelated to the current token and will not be used.
- D. The token will be the only one considered in the next generation step.

**Answer: A**

Explanation:

Comprehensive and Detailed In-Depth Explanation=

In "Show Likelihoods," a higher number (probability score) indicates a token's greater likelihood of following the current token, reflecting the model's prediction confidence-Option B is correct. Option A (less likely) is the opposite. Option C (unrelated) misinterprets likelihood ties tokens contextually. Option D (only one) assumes greedy decoding, not the feature's purpose. This helps users understand model preferences.

OCI 2025 Generative AI documentation likely explains "Show Likelihoods" under token generation insights.

#### NEW QUESTION # 91

You create a fine-tuning dedicated AI cluster to customize a foundational model with your custom training data. How many unit hours are required for fine-tuning if the cluster is active for 10 hours?

- A. 40 unit hours
- B. 25 unit hours
- C. 30 unit hours
- **D. 20 unit hours**

**Answer: D**

Explanation:

Comprehensive and Detailed In-Depth Explanation=

In OCI, unit hours typically equal actual hours of cluster activity unless specified otherwise (e.g., per GPU scaling). For 10 hours of activity, it's  $10 \text{ hours} \times 1 \text{ unit/hour} = 10 \text{ unit hours}$ , but options suggest a multiplier (common in cloud pricing). Assuming a standard 2-unit/hour rate (e.g., for GPU clusters), it's  $10 \times 2 = 20 \text{ unit hours}$ -Option C fits best. Options A, B, and D imply inconsistent rates (2.5, 4, 3).

OCI 2025 Generative AI documentation likely specifies unit hour rates under DedicatedAI Cluster pricing.

#### NEW QUESTION # 92

.....

Just like the old saying goes, there is no royal road to success, and only those who do not dread the fatiguing climb of gaining its

numinous summits. In a similar way, there is no smoothly paved road to the 1Z0-1127-25 certification. You have to work on it and get started from now. If you want to gain the related certification, it is very necessary that you are bound to spend some time on carefully preparing for the 1Z0-1127-25 Exam, including choosing the convenient and practical study materials, sticking to study and keep an optimistic attitude and so on.

New Study 1Z0-1127-25 Questions: <https://www.exam4docs.com/1Z0-1127-25-study-questions.html>

P.S. Free 2026 Oracle 1Z0-1127-25 dumps are available on Google Drive shared by Exam4Docs: <https://drive.google.com/open?id=1Er5DsXBruf4nk07-w8mvZYoYbP3LPGtX>