



Oracle 1Z0-1122-26

Oracle Cloud Infrastructure 2026 AI Foundations Associate

**Thank You for Downloading 1Z0-1122-26 Updated Exam
Questions**

<https://certs4sale.com/oracle/1z0-1122-26-pdf-exam-dumps>

Oracle 1Z0-1122-26

Oracle Cloud Infrastructure AI Foundations Associate

ORACLE CLOUD
PRACTICE ASSESSMENT

TARGET CERTIFICATION
OCI AI Foundations Associate

EXAM CODE
1Z0-1122-26

DOCUMENT PURPOSE
Practice Questions & Detailed Explanations

CORE FOCUS AREAS
Generative AI, ML, OCI AI Services & Prompting

Question 1

SINGLE CHOICE

Your organization is building a customer service chatbot that needs to understand and respond to customer inquiries in natural language. You're evaluating different AI approaches for this task. Which of the following best describes why a Large Language Model (LLM) with Transformer architecture is more suitable than a traditional supervised learning classifier for this application?

- A. LLMs use Transformer architecture which enables parallel processing of input sequences and attention mechanisms that can capture long-range dependencies in text, making them superior for understanding context in open-ended customer inquiries
- B. LLMs require less training data than supervised classifiers and automatically scale to handle any input without fine-tuning
- C. Transformers are recurrent neural networks that process text sequentially, which is more efficient for real-time customer service applications
- D. LLMs use convolutional layers that excel at extracting local patterns from text, similar to how CNNs work on images

Correct Answer: A

Explanation:

The correct answer explains the core advantage of Transformers: they use self-attention mechanisms that allow the model to weigh the importance of different parts of the input sequence, enabling understanding of context and long-range dependencies without the sequential processing limitations of RNNs. This is critical for understanding nuanced customer inquiries.

Why other options are incorrect:

- *Option B* is misleading—LLMs typically require substantial training data, though they can be fine-tuned with less data than training from scratch.

Question 2

SINGLE CHOICE

A financial services company wants to detect fraudulent transactions in real-time. They have historical transaction data labeled as fraudulent or legitimate. Which machine learning approach is most appropriate for this scenario, and what key consideration should guide their implementation?

- A. Supervised learning classification is appropriate because labeled historical data exists, but they must be careful to handle class imbalance since fraudulent transactions typically represent a small fraction of total transactions
- B. Unsupervised learning clustering is the best approach because it doesn't require labeled data and can discover fraud patterns automatically
- C. Reinforcement learning is necessary because the model must learn to optimize the decision to approve or deny transactions through trial and error
- D. Deep learning with LLMs is required because transaction fraud detection requires natural language understanding of customer descriptions

Correct Answer: A

Explanation:

The correct answer identifies this as a supervised learning classification problem (labeled data exists) and highlights the critical real-world challenge: fraudulent transactions are typically rare, creating a class imbalance problem that requires special handling (e.g., resampling, weighted loss functions, or adjusted thresholds).

Why other options are incorrect:

- *Option B:* While unsupervised learning can discover patterns, having labeled historical data makes supervised learning more effective and directly aligned with the business objective.
- *Option C:* Reinforcement learning is designed for sequential decision-making with rewards/penalties over time, not for classification of static transactions.
- *Option D:* LLMs are designed for language tasks; transaction fraud detection uses structured numerical and categorical data, not natural language text.

Question 3

SINGLE CHOICE

Your team is deploying an OCI AI Vision service to automatically categorize incoming medical images (X-rays, CT scans) and flag potential abnormalities for physician review. Which of the following statements about this implementation are correct?

- A. OCI Vision uses pre-trained deep learning models including convolutional neural networks optimized for image classification, and the service handles the underlying infrastructure, allowing your team to focus on integration and data preparation
- B. OCI Vision requires you to build and train your own CNN models from scratch, providing raw APIs to upload your custom models
- C. Since this is a medical application, OCI Vision automatically ensures regulatory compliance (HIPAA, FDA approval) without additional governance measures
- D. OCI Vision can only process images in real-time streaming mode and cannot handle batch processing of historical image archives

Correct Answer: A

Explanation:

The correct answer accurately describes OCI Vision as a managed service with pre-trained models (including CNNs for image classification) that abstracts away infrastructure management. The service provides ready-to-use APIs for image analysis without requiring customers to build models from scratch.

Why other options are incorrect:

- *Option B:* Mischaracterizes OCI Vision as a raw model development platform; it is a managed service with pre-trained models.
- *Option C:* While OCI has responsible AI features, automated compliance is not a feature; organizations must still implement proper governance, data handling, and regulatory controls.
- *Option D:* OCI Vision supports both real-time and batch processing modes, giving flexibility for different use cases.

Question 4

SINGLE CHOICE

An e-commerce company is implementing prompt engineering strategies for their chatbot powered by OCI Generative AI. They want to improve response quality and consistency. Which of the following best practices should they prioritize?

- A. Providing clear, specific instructions with examples (few-shot prompting), maintaining consistent formatting, and iterating on prompts based on output quality to refine behavior
- B. Making prompts as short as possible to minimize computational cost, and avoiding any examples or context as they slow down processing
- C. Instructing the LLM to always provide responses in a single paragraph regardless of topic complexity, as this standardizes output format
- D. Using abstract, philosophical language in prompts because LLMs perform better with creative rather than direct instruction

Correct Answer: A

Explanation:

The correct answer reflects established prompt engineering best practices: clarity, specificity, in-context examples, and iterative refinement. These techniques significantly improve LLM output quality and consistency for business applications.

Why other options are incorrect:

- *Option B:* While brevity matters, omitting context and examples actually reduces quality; modern LLMs handle more detailed prompts efficiently and benefit from examples.
- *Option C:* Forcing single-paragraph format regardless of content complexity reduces usability and clarity.
- *Option D:* LLMs respond better to direct, clear instructions than abstract language; clarity beats creativity for predictable business outcomes.

Question 5

MULTIPLE CHOICE

A retail company wants to use OCI AI Language service to analyze customer reviews at scale and extract insights about product sentiment, key topics, and named entities (product names, competitors). Which capabilities are provided by OCI Language service for this use case? (Select all that apply)

- A. Sentiment analysis to classify reviews as positive, negative, or neutral with confidence scores
- B. Named entity recognition (NER) to extract entities such as product names, company names, and competitor mentions from text
- C. Text classification and key phrase extraction to identify topics, themes, and important concepts mentioned in reviews
- D. Automatic generation of customer refund policies based on detected negative sentiment in reviews
- E. Conversion of customer review text into images or video format for marketing purposes

Correct Answer: A, B, C

Explanation:

The correct answers identify the core capabilities of OCI Language service: sentiment analysis, named entity recognition, and text classification/key phrase extraction. These are all standard NLP capabilities provided by the service for analyzing unstructured text.

Why other options are incorrect:

- *Option D:* OCI Language is a text analysis service; it does not generate business policies or refund decisions. Such decisions require business logic and human judgment.
- *Option E:* Text-to-media conversion is not a function of the Language service; that would fall under different OCI services or require separate tools.

End of Assessment Module

You have completed the preview module for **Oracle 1Z0-1122-26: OCI AI Foundations Associate**.

THANK YOU FOR DOWNLOADING 1Z0-1122-26 UPDATED EXAM QUESTIONS

Note: Thanks For Trying The Demo Of Our 1Z0-1122-26 Exam Product

Visit Our Site to Purchase the Full Set of Actual 1Z0-1122-26 Exam Questions With Answers.

Money Back Guarantee



Click The Link Below

<https://certs4sale.com/oracle/1z0-1122-26-pdf-exam-dumps>