

# 2026 1Z0-184-25 Valid Test Blueprint Free PDF | High Pass-Rate 1Z0-184-25 Reliable Real Exam: Oracle AI Vector Search Professional



P.S. Free 2025 Oracle 1Z0-184-25 dumps are available on Google Drive shared by DumpsTests: [https://drive.google.com/open?id=1n3MEICDN6y7yLSLW\\_FaWwU1Ktw6xN8z](https://drive.google.com/open?id=1n3MEICDN6y7yLSLW_FaWwU1Ktw6xN8z)

Our 1Z0-184-25 practice prep provides you with a brand-new learning method that lets you get rid of heavy schoolbags, lose boring textbooks, and let you master all the important knowledge in the process of making a question. Please believe that with 1Z0-184-25 Real Exam, you will fall in love with learning. Our 1Z0-184-25 exam questions are contained in three versions: the PDF, Software and APP online which can cater to different needs of our customers.

## Oracle 1Z0-184-25 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"><li>Using Vector Embeddings: This section measures the abilities of AI Developers in generating and storing vector embeddings for AI applications. It covers generating embeddings both inside and outside the Oracle database and effectively storing them within the database for efficient retrieval and processing.</li></ul>
Topic 2	<ul style="list-style-type: none"><li>Building a RAG Application: This section assesses the knowledge of AI Solutions Architects in implementing retrieval-augmented generation (RAG) applications. Candidates will learn to build RAG applications using PL</li><li>SQL and Python to integrate AI models with retrieval techniques for enhanced AI-driven decision-making.</li></ul>
Topic 3	<ul style="list-style-type: none"><li>Performing Similarity Search: This section tests the skills of Machine Learning Engineers in conducting similarity searches to find relevant data points. It includes performing exact and approximate similarity searches using vector indexes. Candidates will also work with multi-vector similarity search to handle searches across multiple documents for improved retrieval accuracy.</li></ul>

Topic 4	<ul style="list-style-type: none"> <li>• <b>Leveraging Related AI Capabilities:</b> This section evaluates the skills of Cloud AI Engineers in utilizing Oracle's AI-enhanced capabilities. It covers the use of Exadata AI Storage for faster vector search, Select AI with Autonomous for querying data using natural language, and data loading techniques using SQL Loader and Oracle Data Pump to streamline AI-driven workflows.</li> </ul>
---------	---

## >> 1Z0-184-25 Valid Test Blueprint <<

### 1Z0-184-25 Reliable Real Exam | New 1Z0-184-25 Test Cram

DumpsTests is responsible for our 1Z0-184-25 study materials. Every exam product of DumpsTests have sold to customer will enjoy considerate after-sales service. If you have problems about our 1Z0-184-25 study materials such as installation, operation and so on, we will quickly reply to you after our online workers have received your emails. We are not afraid of troubles. We warmly welcome to your questions and suggestions on the 1Z0-184-25 Exam Questions. We sincerely hope we can help you solve your problem and help you pass the 1Z0-184-25 exam.

### Oracle AI Vector Search Professional Sample Questions (Q45-Q50):

#### NEW QUESTION # 45

Which PL/SQL package is primarily used for interacting with Generative AI services in Oracle Database 23ai?

- A. DBMS\_VECTOR\_CHAIN
- B. DBMS\_GENAI
- C. DBMS\_ML
- **D. DBMS\_AI**

**Answer: D**

Explanation:

Oracle Database 23ai introduces DBMS\_AI as the primary PL/SQL package for interacting with Generative AI services, such as OCI Generative AI, enabling features like natural language query processing (e.g., Select AI) and AI-driven insights. DBMS\_ML (B) focuses on machine learning model training and management, not generative AI. DBMS\_VECTOR\_CHAIN (C) supports vector processing workflows (e.g., document chunking, embedding), but it's not the main interface for generative AI services. DBMS\_GENAI (D) is not a recognized package in 23ai documentation. DBMS\_AI's role is highlighted in Oracle's AI integration features for 23ai.

#### NEW QUESTION # 46

Which is a characteristic of an approximate similarity search in Oracle Database 23ai?

- **A. It trades off accuracy for faster performance**
- B. It compares every vector in the dataset
- C. It always guarantees 100% accuracy
- D. It is slower than exact similarity search

**Answer: A**

Explanation:

Approximate similarity search (ANN) in Oracle 23ai (B) uses indexes (e.g., HNSW, IVF) to trade accuracy for speed, returning near-matches faster by not comparing all vectors. Exact search compares every vector (A), not ANN. It doesn't guarantee 100% accuracy (C); that's exact search. It's faster, not slower (D), than exact search due to indexing. Oracle's documentation defines ANN's speed-accuracy trade-off as its hallmark.

### NEW QUESTION # 47

What is the primary purpose of the DBMS\_VECTOR\_CHAIN.UTL\_TO\_CHUNKS package in a RAG application?

- A. To convert a document into a single, large text string
- **B. To split a large document into smaller chunks to improve vector quality by minimizing token truncation**
- C. To load a document into the database
- D. To generate vector embeddings from a text document

**Answer: B**

Explanation:

In Oracle Database 23ai, the DBMS\_VECTOR\_CHAIN package supports Retrieval Augmented Generation (RAG) workflows by providing utilities for vector processing. The UTL\_TO\_CHUNKS function specifically splits large documents into smaller, manageable text chunks. This is critical in RAG applications because embedding models (e.g., BERT, ONNX models) have token limits (e.g., 512 tokens). Splitting text minimizes token truncation, ensuring that each chunk retains full semantic meaning, which improves the quality of subsequent vector embeddings and search accuracy. Generating embeddings (A) is handled by functions like VECTOR\_EMBEDDING, not UTL\_TO\_CHUNKS. Loading documents (B) is a separate process (e.g., via SQL\*Loader). Converting to a single text string (D) contradicts the chunking purpose and risks truncation. Oracle's documentation on DBMS\_VECTOR\_CHAIN emphasizes chunking for optimizing vector quality in RAG.

### NEW QUESTION # 48

A machine learning team is using IVF indexes in Oracle Database 23ai to find similar images in a large dataset. During testing, they observe that the search results are often incomplete, missing relevant images. They suspect the issue lies in the number of partitions probed. How should they improve the search accuracy?

- **A. Add the TARGET\_ACCURACY clause to the query with a higher value for the accuracy**
- B. Change the index type to HNSW for better accuracy
- C. Increase the VECTOR\_MEMORY\_SIZE initialization parameter
- D. Re-create the index with a higher EFCONSTRUCTION value

**Answer: A**

Explanation:

IVF (Inverted File) indexes in Oracle 23ai partition vectors into clusters, probing a subset during queries for efficiency. Incomplete results suggest insufficient partitions are probed, reducing recall. The TARGET\_ACCURACY clause (A) allows users to specify a desired accuracy percentage (e.g., 90%), dynamically increasing the number of probed partitions to meet this target, thus improving accuracy at the cost of latency. Switching to HNSW (B) offers higher accuracy but requires re-indexing and may not be necessary if IVF tuning suffices. Increasing VECTOR\_MEMORY\_SIZE (C) allocates more memory for vector operations but doesn't directly affect probe count. EFCONSTRUCTION (D) is an HNSW parameter, irrelevant to IVF. Oracle's IVF documentation highlights TARGET\_ACCURACY as the recommended tuning mechanism.

### NEW QUESTION # 49

What is the primary function of an embedding model in the context of vector search?

- A. To store vectors in a structured format for efficient retrieval
- B. To execute similarity search operations within a database
- **C. To transform text or data into numerical vector representations**
- D. To define the schema for a vector database

**Answer: C**

Explanation:

An embedding model in the context of vector search, such as those used in Oracle Database 23ai, is fundamentally a machine learning construct (e.g., BERT, SentenceTransformer, or an ONNX model) designed to transform raw data—typically text, but also images or other modalities—into numerical vector representations (C). These vectors, stored in the VECTOR data type, encapsulate semantic meaning in a high-dimensional space where proximity reflects similarity. For instance, the word "cat" might be mapped to a 512-dimensional vector like [0.12, -0.34, ...], where its position relative to "dog" indicates relatedness. This transformation is the linchpin of vector search, enabling mathematical operations like cosine distance to find similar items. Option A (defining schema) misattributes a database design role to the model; schema is set by DDL (e.g., CREATE TABLE with

What's more, part of that DumpsTests 1Z0-184-25 dumps now are free: [https://drive.google.com/open?id=1n3MEICDN6y7yLSLW\\_FaWwU1Ktw6xN8z](https://drive.google.com/open?id=1n3MEICDN6y7yLSLW_FaWwU1Ktw6xN8z)