

# Pass Guaranteed 2025 Newest Oracle 1Z0-184-25 Study Guide



P.S. Free & New 1Z0-184-25 dumps are available on Google Drive shared by DumpTorrent: [https://drive.google.com/open?id=1K2\\_vHwKbmRpafK1nwsXP0OffW51Vda](https://drive.google.com/open?id=1K2_vHwKbmRpafK1nwsXP0OffW51Vda)

The goal of 1Z0-184-25 exam torrent is to help users pass the exam with the shortest possible time and effort. With 1Z0-184-25 exam torrent, you neither need to keep yourself locked up in the library for a long time nor give up a rare vacation to review. You will never be frustrated by the fact that you can't solve a problem. With 1Z0-184-25 question torrent, you will suddenly find the joy of learning and you will pass the professional qualification exam very easily.

## 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>Using Vector Indexes: This section evaluates the expertise of AI Database Specialists in optimizing vector searches using indexing techniques. It covers the creation of vector indexes to enhance search speed, including the use of HNSW and IVF vector indexes for performing efficient search queries in AI-driven applications.</li> </ul>
Topic 3	<ul style="list-style-type: none"> <li>Understand Vector Fundamentals: This section of the exam measures the skills of Data Engineers in working with vector data types for storing embeddings and enabling semantic queries. It covers vector distance functions and metrics used in AI vector search. Candidates must demonstrate proficiency in performing DML and DDL operations on vectors to manage data efficiently.</li> </ul>
Topic 4	<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>

>> 1Z0-184-25 Study Guide <<

## Exam 1Z0-184-25 Vce Format - Instant 1Z0-184-25 Download

Dear customers, we would like to make it clear that learning knowledge and striving for certificates of exam is a self-improvement process, and you will realize yourself rather than offering benefits for anyone. So our 1Z0-184-25 practice materials are once a lifetime opportunity you cannot miss. With all advantageous features introduced as follow, please read them carefully.

## Oracle AI Vector Search Professional Sample Questions (Q16-Q21):

### NEW QUESTION # 16

What is the significance of splitting text into chunks in the process of loading data into Oracle AI Vector Search?

- A. To reduce the computational burden on the embedding model
- B. To facilitate parallel processing of the data during vectorization
- C. To minimize token truncation as each vector embedding model has its own maximum token limit

**Answer: C**

Explanation:

Splitting text into chunks (C) in Oracle AI Vector Search (e.g., via DBMS\_VECTOR\_CHAIN.UTL\_TO\_CHUNKS) ensures that each segment fits within the token limit of embedding models (e.g., 512 tokens for BERT), preventing truncation that loses semantic content. This improves vector quality for similarity search. Reducing computational burden (A) is a secondary effect, not the primary goal. Parallel processing (B) may occur but isn't the main purpose; chunking is about model compatibility. Oracle's documentation emphasizes chunking to align with embedding model constraints.

### NEW QUESTION # 17

What is the primary purpose of a similarity search in Oracle Database 23ai?

- A. Optimize relational database operations to compute distances between all data points in a database
- B. To group vectors by their exact scores
- C. To find exact matches in BLOB data
- D. To retrieve the most semantically similar entries using distance metrics between different vectors

**Answer: D**

Explanation:

Similarity search in Oracle 23ai (C) uses vector embeddings in VECTOR columns to retrieve entries semantically similar to a query vector, based on distance metrics (e.g., cosine, Euclidean) via functions like VECTOR\_DISTANCE. This is key for AI applications like RAG, finding "close" rather than exact matches. Optimizing relational operations (A) is unrelated; similarity search is vector-specific. Exact matches in BLOBs (B) don't leverage vector semantics. Grouping by scores (D) is a post-processing step, not the primary purpose. Oracle's documentation defines similarity search as retrieving semantically proximate vectors.

## NEW QUESTION # 18

What is the default distance metric used by the VECTOR\_DISTANCE function if none is specified?

- A. Cosine
- B. Euclidean
- C. Hamming
- D. Manhattan

**Answer: A**

Explanation:

The VECTOR\_DISTANCE function in Oracle 23ai computes vector distances, and if no metric is specified (e.g., VECTOR\_DISTANCE(v1, v2)), it defaults to Cosine (C). Cosine distance (1 - cosine similarity) is widely used for text embeddings due to its focus on angular separation, ignoring magnitude-fitting for normalized vectors from models like BERT. Euclidean (A) measures straight-line distance, not default. Hamming (B) is for binary vectors, rare in 23ai's FLOAT32 context. Manhattan (D) sums absolute differences, less common for embeddings. Oracle's choice of Cosine reflects its AI focus, as documentation confirms, aligning with industry norms for semantic similarity-vital for users assuming defaults in queries.

## NEW QUESTION # 19

Which parameter is used to define the number of closest vector candidates considered during HNSW index creation?

- A. VECTOR\_MEMORY\_SIZE
- B. NEIGHBOURS
- C. EFCONSTRUCTION
- D. TARGET\_ACCURACY

**Answer: C**

Explanation:

In Oracle 23ai, EFCONSTRUCTION (A) controls the number of closest vector candidates (edges) considered during HNSW index construction, affecting the graph's connectivity and search quality. Higher values improve accuracy but increase build time. VECTOR\_MEMORY\_SIZE (B) sets memory allocation, not candidate count. NEIGHBOURS (C) isn't a parameter; it might confuse with NEIGHBOR\_PARTITION (IVF). TARGET\_ACCURACY (D) adjusts query-time accuracy, not index creation. Oracle's HNSW documentation specifies EFCONSTRUCTION for this purpose.

## NEW QUESTION # 20

In the following Python code, what is the significance of prepending the source filename to each text chunk before storing it in the vector database?

bash

CollapseWrapCopy

```
docs = [{"text": filename + "|" + section, "path": filename} for filename, sections in faqs.items() for section in sections]
# Sample the resulting data
docs[2]
```

- A. It helps differentiate between chunks from different files but has no impact on vectorization
- B. It improves the accuracy of the LLM by providing additional training data
- C. It preserves context and aids in the retrieval process by associating each vectorized chunk with its original source file
- D. It speeds up the vectorization process by providing a unique identifier for each chunk

**Answer: C**

Explanation:

Prepending the filename to each text chunk (e.g., filename + "|" + section) in the Python code (A) preserves contextual metadata, linking each chunk-and its resulting vector-to its source file. This aids retrieval in RAG applications by allowing the application to trace back to the original document, enhancing response context (e.g., "from Book1"). While it differentiates chunks (B), its impact goes beyond identification, affecting retrieval usability. It doesn't speed up vectorization (C); embedding models process text regardless of prefixes. It also doesn't train the LLM (D); it's metadata for retrieval, not training data. Oracle's RAG examples emphasize metadata preservation for context-aware responses.

## NEW QUESTION # 21

Our 1Z0-184-25 study guide and training materials of DumpTorrent are summarized by experienced IT experts, who combine the 1Z0-184-25 original questions and real answers. Due to our professional team, the passing rate of 1Z0-184-25 test of our DumpTorrent is the highest in the 1Z0-184-25 exam training. So, choosing DumpTorrent, choosing success.

Exam 1Z0-184-25 Vce Format: <https://www.dumptorrent.com/1Z0-184-25-braindumps-torrent.html>

What's more, part of that DumpTorrent 1Z0-184-25 dumps now are free: [https://drive.google.com/open?id=1K2\\_vHwKbmRpafK1nwsXP0OfWS51Vda](https://drive.google.com/open?id=1K2_vHwKbmRpafK1nwsXP0OfWS51Vda)