

Prepare Smart for Success Free Oracle 1Z0-184-25 Exam Questions and Answers

Ready to pass faster? Grab free and updated Oracle Database AI Vector Search Professional exam PDF questions now. Get authentic 1Z0-184-25 dumps packed with verified answers and secure your certification success with [PrepBolt](https://prepbolt.com/1Z0-184-25.html) 1Z0-184-25 exam pdf questions and answers.

Thank you for Downloading 1Z0-184-25 exam PDF Demo

<https://prepbolt.com/1Z0-184-25.html>

Question 1

Question Type: MultipleChoice

You are tasked with creating a table to store vector embeddings with the following characteristics: Each vector must have exactly 512 dimensions, and the dimensions should be stored as 32-bit floating point numbers. Which SQL statement should you use?

Options:

- A- CREATE TABLE vectors (id NUMBER, embedding VECTOR(512))
- B- CREATE TABLE vectors (id NUMBER, embedding VECTOR)
- C- CREATE TABLE vectors (id NUMBER, embedding VECTOR(*, INT8))
- D- CREATE TABLE vectors (id NUMBER, embedding VECTOR(512, FLOAT32))

Answer:

D

Explanation:

In Oracle 23ai, the VECTOR data type can specify dimensions and precision. CREATE TABLE vectors (id NUMBER, embedding VECTOR(512, FLOAT32)) (D) defines a column with exactly 512 dimensions and FLOAT32 (32-bit float) format, meeting both requirements. Option A omits the format (defaults vary), risking mismatch. Option B is unspecified, allowing variable dimensions---not "exactly 512." Option C uses INT8, not FLOAT32, and '*' denotes undefined dimensions. Oracle's SQL reference confirms this syntax for precise VECTOR definitions.

Question 2

Question Type: MultipleChoice

In Oracle Database 23ai, which SQL function calculates the distance between two vectors using the Euclidean metric?

Options:

- A- L1_DISTANCE

- B- L2_DISTANCE
- C- HAMMING_DISTANCE
- D- COSINE_DISTANCE

Answer:

B

Explanation:

In Oracle Database 23ai, vector distance calculations are primarily handled by the VECTOR_DISTANCE function, which supports multiple metrics (e.g., COSINE, EUCLIDEAN) specified as parameters (e.g., VECTOR_DISTANCE(v1, v2, EUCLIDEAN)). However, the question implies distinct functions, a common convention in some databases or libraries, and Oracle's documentation aligns L2_DISTANCE (B) with the Euclidean metric. L2 (Euclidean) distance is the straight-line distance between two points in vector space, computed as $\sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$, where x_i and y_i are vector components. For example, for vectors [1, 2] and [4, 6], L2 distance is $\sqrt{(1-4)^2 + (2-6)^2} = \sqrt{9 + 16} = 5$.

Option A, L1_DISTANCE, represents Manhattan distance ($|x_1 - x_2| + |y_1 - y_2|$), summing absolute differences---not Euclidean. Option C, HAMMING_DISTANCE, counts differing bits, suited for binary vectors (e.g., INT8), not continuous Euclidean spaces typically used with FLOAT32 embeddings. Option D, COSINE_DISTANCE (1 - cosine similarity), measures angular separation, distinct from Euclidean's magnitude-inclusive approach. While VECTOR_DISTANCE is the general function in 23ai, L2_DISTANCE may be an alias or a contextual shorthand in some Oracle AI examples, reflecting Euclidean's prominence in geometric similarity tasks. Misinterpreting this could lead to choosing COSINE for spatial tasks where magnitude matters, skewing results. Oracle's vector search framework supports Euclidean via VECTOR_DISTANCE, but B aligns with the question's phrasing.

Question 3

Question Type: MultipleChoice

What is a key advantage of using GoldenGate 23ai for managing and distributing vector data for AI applications?

Options:

- A- Real-time vector data updates across locations
- B- Automatic translation of vector embeddings between formats
- C- Specialized vector embedding compression
- D- Built-in version control for vector data

Answer:

A

Explanation:

Oracle GoldenGate 23ai is a real-time data replication and integration tool, extended in 23ai to handle the VECTOR data type for AI applications. Its key advantage (A) is enabling real-time updates of vector data across distributed locations---e.g., replicating VECTOR columns from a primary database in New York to a secondary in London with sub-second latency. This ensures AI models (e.g., for similarity search or RAG) access the latest embeddings as source data (e.g., documents) changes, critical for dynamic environments like customer support systems where new queries demand current context. Imagine a VECTOR column storing embeddings of support tickets; GoldenGate keeps these synchronized across regions, minimizing staleness that could degrade AI responses.

Option B (automatic translation) is fictional; GoldenGate doesn't convert vector formats (e.g., FLOAT32 to INT8)---that's a model or application task. Option C (compression) isn't a GoldenGate feature; compression might occur at the storage layer, but GoldenGate focuses on replication fidelity, not size reduction. Option D (version control) misaligns with GoldenGate's purpose; it ensures data consistency, not historical versioning like Git. Real-time replication (A) stands out, as Oracle's documentation emphasizes GoldenGate's role in keeping vector-driven AI applications globally consistent, a game-changer for distributed AI deployments where latency or inconsistency could disrupt user trust. Without this, static exports (e.g., Data Pump) would lag, undermining real-time AI use cases.

Question 4

Question Type: MultipleChoice

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?

Options:

- 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.

Question 5

Question Type: MultipleChoice

What happens when you attempt to insert a vector with an incorrect number of dimensions into a VECTOR column with a defined number of dimensions?

Options:

- A- The database truncates the vector to fit the defined dimensions
- B- The database pads the vector with zeros to match the defined dimensions
- C- The database ignores the defined dimensions and inserts the vector as is
- D- The insert operation fails, and an error message is thrown

Answer:

D

Explanation:

In Oracle Database 23ai, a VECTOR column with a defined dimension count (e.g., VECTOR(4, FLOAT32)) enforces strict dimensional integrity to ensure consistency for similarity search and indexing. Attempting to insert a vector with a mismatched number of dimensions---say, TO_VECTOR('[1.2, 3.4, 5.6]') (3D) into a VECTOR(4)---results in the insert operation failing with an error (D), such as ORA-13199: 'vector dimension mismatch.' This rigidity protects downstream AI operations; a 3D vector in a 4D column would misalign with indexed data (e.g., HNSW graphs),

breaking similarity calculations like cosine distance, which require uniform dimensionality.

Option A (truncation) is tempting but incorrect; Oracle doesn't silently truncate [1.2, 3.4, 5.6] to [1.2, 3.4]---this would discard data arbitrarily, risking semantic loss (e.g., a truncated sentence embedding losing meaning). Option B (padding with zeros) seems plausible---e.g., [1.2, 3.4, 5.6] becoming [1.2, 3.4, 5.6, 0]---but Oracle avoids implicit padding to prevent unintended semantic shifts (zero-padding could alter distances). Option C (ignoring dimensions) only applies to undefined VECTOR columns (e.g., VECTOR without size), not fixed ones; here, the constraint is enforced. The failure (D) forces developers to align data explicitly (e.g., regenerate embeddings), ensuring reliability---a strict but necessary design choice in Oracle's AI framework. In practice, this error prompts debugging upstream data pipelines, avoiding silent failures that could plague production AI systems.

Question 6

Question Type: MultipleChoice

You need to generate a vector from the string '[1.2, 3.4]' in FLOAT32 format with 2 dimensions. Which function will you use?

Options:

- A- TO_VECTOR
- B- VECTOR_DISTANCE
- C- FROM_VECTOR
- D- VECTOR_SERIALIZE

Answer:

A

Explanation:

In Oracle Database 23ai, the TO_VECTOR function (A) converts a string representation of a vector (e.g., '[1.2, 3.4]') into a VECTOR data type with specified format (e.g., FLOAT32) and dimensions (here, 2). It's designed for creating vectors from text input, matching the requirement. VECTOR_DISTANCE (B) calculates distances between vectors, not generates them. FROM_VECTOR (C) isn't a documented function; it might be confused with serialization or extraction, but it's not standard. VECTOR_SERIALIZE (D) converts a vector to a string, the opposite of what's needed. Oracle's SQL reference confirms TO_VECTOR for this purpose, parsing the string into a 2D FLOAT32 vector.

Question 7

Question Type: MultipleChoice

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

Options:

- A- EFCONSTRUCTION
- B- VECTOR_MEMORY_SIZE
- C- NEIGHBOURS
- D- TARGET_ACCURACY

Answer:

A

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_PARTITIONS (IVF). TARGET_ACCURACY (D) adjusts query-time accuracy, not index creation. Oracle's HNSW documentation specifies EFCONSTRUCTION for this purpose.

Thank You for trying 1Z0-184-25 PDF Demo

To try our 1Z0-184-25 practice exam software
visit link below

<https://prepbolt.com/1Z0-184-25.html>

Start Your 1Z0-184-25 Preparation

Use Coupon "SAVE50" for extra 50% discount on the purchase of Practice Test Software. Test your 1Z0-184-25 preparation with actual exam questions.