1Z0-184-25 Vce Download & Latest 1Z0-184-25 Braindumps

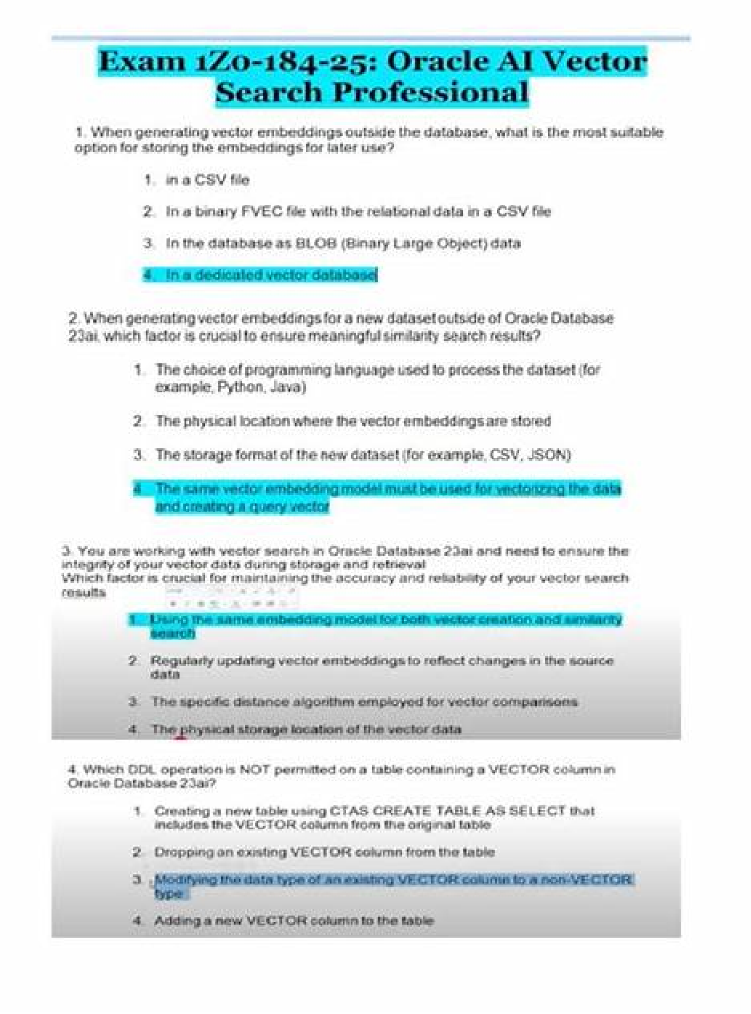

Exam 1Z0-184-25: Oracle AI Vector Search Professional

1. When generating vector embeddings outside the database, what is the most suitable option for storing the embeddings for later use?

    1. in a CSV file
    2. In a binary FVEC file with the relational data in a CSV file
    3. In the database as BLOB (Binary Large Object) data
    4. In a dedicated vector database

2. When generating vector embeddings for a new dataset outside of Oracle Database 23ai, which factor is crucial to ensure meaningful similarity search results?

    1. The choice of programming language used to process the dataset (for example, Python, Java)
    2. The physical location where the vector embeddings are stored
    3. The storage format of the new dataset (for example, CSV, JSON)
    4. The same vector embedding model must be used for vectorizing the data and creating a query vector

3. You are working with vector search in Oracle Database 23ai and need to ensure the integrity of your vector data during storage and retrieval
Which factor is crucial for maintaining the accuracy and reliability of your vector search results

    1. Using the same embedding model for both vector creation and similarity search
    2. Regularly updating vector embeddings to reflect changes in the source data
    3. The specific distance algorithm employed for vector comparisons
    4. The physical storage location of the vector data

4. Which DDL operation is NOT permitted on a table containing a VECTOR column in Oracle Database 23ai?

    1. Creating a new table using CTAS CREATE TABLE AS SELECT that includes the VECTOR column from the original table
    2. Dropping an existing VECTOR column from the table
    3. Modifying the data type of an existing VECTOR column to a non-VECTOR type
    4. Adding a new VECTOR column to the table

BTW, DOWNLOAD part of DumpsQuestion 1Z0-184-25 dumps from Cloud Storage: https://drive.google.com/open?id=1FUtGryyO4wdVf01ELoxHNP_XnAgElydg

The Oracle AI Vector Search Professional (1Z0-184-25) PDF dumps are suitable for smartphones, tablets, and laptops as well. So you can study actual Oracle AI Vector Search Professional (1Z0-184-25) questions in PDF easily anywhere. DumpsQuestion updates Oracle AI Vector Search Professional (1Z0-184-25) PDF dumps timely as per adjustments in the content of the actual Oracle 1Z0-184-25 exam.

Oracle 1Z0-184-25 Exam Syllabus Topics:

Topic	Details
Topic 1	• 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.
Topic 2	• 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.

Topic 3	• 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.
Topic 4	• Leveraging Related AI Capabilities: 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.
Topic 5	• 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.

Latest Oracle 1Z0-184-25 Braindumps, Valid 1Z0-184-25 Exam Notes

Under the support of our study materials, passing the exam won't be an unreachable mission. More detailed information is under below. We are pleased that you can spare some time to have a look for your reference about our 1Z0-184-25 test prep. As long as you spare one or two hours a day to study with our laTest 1Z0-184-25 Quiz prep, we assure that you will have a good command of the relevant knowledge before taking the exam. What you need to do is to follow the 1Z0-184-25 exam guide system at the pace you prefer as well as keep learning step by step.

Oracle AI Vector Search Professional Sample Questions (Q54-Q59):

NEW QUESTION # 54
What security enhancement is introduced in Exadata System Software 24ai?

- A. SNMP security (Security Network Management Protocol)
- B. Integration with third-party security tools
- C. Enhanced encryption algorithm for data at rest

Answer: C

Explanation:
Exadata System Software 24ai (noted in context beyond 23ai) introduces an enhanced encryption algorithm for data at rest (B), strengthening security for stored data, including vectors. Third-party integration (A) isn't highlighted as a 24ai feature. SNMP security (C) relates to network monitoring, not a primary Exadata enhancement. Oracle's Exadata documentation for 24ai emphasizes advanced encryption as a key security upgrade.

NEW QUESTION # 55
What is the advantage of using Euclidean Squared Distance rather than Euclidean Distance in similarity search queries?

- A. It guarantees higher accuracy than Euclidean Distance
- B. It supports hierarchical partitioning of vectors
- C. It is the default distance metric for Oracle AI Vector Search
- D. It is simpler and faster because it avoids square-root calculations

Answer: D

Explanation:
Euclidean Squared Distance (L2-squared) skips the square-root step of Euclidean Distance (L2), i.e., $\sum(xi - yi)^2$ vs. $\sqrt{\sum(xi - yi)^2}$. Since the square root is monotonic, ranking order remains identical, but avoiding it (C) reduces computational cost, making queries faster-crucial for large-scale vector search. It's not the default metric (A); cosine is often default in Oracle 23ai. It doesn't relate to partitioning (B), an indexing feature. Accuracy (D) is equivalent, as rankings are preserved. Oracle's documentation notes L2-

squared as an optimization for performance.

NEW QUESTION # 56
Which of the following actions will result in an error when using VECTOR_DIMENSION_COUNT() in Oracle Database 23ai?

- A. Using a vector with a data type that is not supported by the function
- B. Calling the function on a vector that has been created with TO_VECTOR()
- C. Providing a vector with duplicate values for its components
- D. Providing a vector with a dimensionality that exceeds the specified dimension count

Answer: A

Explanation:
The VECTOR_DIMENSION_COUNT() function in Oracle 23ai returns the number of dimensions in a VECTOR-type value (e.g., 512 for VECTOR(512, FLOAT32)). It's a metadata utility, not a validator of content or structure beyond type compatibility. Option B-using a vector with an unsupported data type-causes an error because the function expects a VECTOR argument; passing, say, a VARCHAR2 or NUMBER instead (e.g., '1,2,3' or 42) triggers an ORA-error (e.g., ORA-00932: inconsistent datatypes). Oracle enforces strict typing for vector functions.
Option A (exceeding specified dimensions) is a red herring; the function reports the actual dimension count of the vector, not the column's defined limit-e.g., VECTOR_DIMENSION_COUNT(TO_VECTOR('[1,2,3]')) returns 3, even if the column is VECTOR(2), as the error occurs at insertion, not here. Option C (duplicate values, like [1,1,2]) is valid; the function counts dimensions (3), ignoring content. Option D (using TO_VECTOR()) is explicitly supported; VECTOR_DIMENSION_COUNT(TO_VECTOR('[1.2, 3.4]')) returns 2 without issue. Misinterpreting this could lead developers to over-constrain data prematurely-B's type mismatch is the clear error case, rooted in Oracle's vector type system.

NEW QUESTION # 57
Which vector index available in Oracle Database 23ai is known for its speed and accuracy, making it a preferred choice for vector search?

- A. Inverted File System (IFS) index
- B. Inverted File (IVF) index
- C. Binary Tree (BT) index
- D. Hierarchical Navigable Small World (HNSW) index

Answer: D

Explanation:
Oracle 23ai supports two main vector indexes: IVF and HNSW. HNSW (D) is renowned for its speed and accuracy, using a hierarchical graph to connect vectors, enabling fast ANN searches with high recall-ideal for latency-sensitive applications like real-time RAG. IVF (C) partitions vectors for scalability but often requires tuning (e.g., NEIGHBOR_PARTITIONS) to match HNSW's accuracy, trading off recall for memory efficiency. BT (A) isn't a 23ai vector index; it's a generic term unrelated here. IFS (B) seems a typo for IVF; no such index exists. HNSW's graph structure outperforms IVF in small-to-medium datasets or where precision matters, as Oracle's documentation and benchmarks highlight, making it a go-to for balanced performance.

NEW QUESTION # 58
What happens when querying with an IVF index if you increase the value of the NEIGHBOR_PARTITIONS probes parameter?

- A. The number of centroids decreases
- B. Index creation time is reduced
- C. More partitions are probed, improving accuracy, but also increasing query latency
- D. Accuracy decreases

Answer: C

Explanation:
The NEIGHBOR_PARTITIONS parameter in Oracle 23ai's IVF index controls how many partitions are probed during a query. Increasing this value examines more clusters, raising theprobability of finding relevant vectors, thus improving accuracy (recall). However, this increases computational effort, leading to higher query latency-a classic ANN trade-off. The number of centroids (A)

is fixed during index creation and unaffected by query parameters. Accuracy does not decrease (B); it improves. Index creation time (C) is unrelated to query-time settings. Oracle's documentation on IVF confirms that NEIGHBOR_PARTITIONS directly governs this accuracy-latency balance.

NEW QUESTION # 59

......

We have to admit that the exam of gaining the 1Z0-184-25 certification is not easy for a lot of people, especial these people who have no enough time. If you also look forward to change your present boring life, maybe trying your best to have the 1Z0-184-25 latest questions are a good choice for you. Now it is time for you to take an exam for getting the certification. If you have any worry about the 1Z0-184-25 Exam, do not worry, we are glad to help you. Because the 1Z0-184-25 cram simulator from our company are very useful for you to pass the 1Z0-184-25 exam and get the certification.

Latest 1Z0-184-25 Braindumps: https://www.dumpsquestion.com/1Z0-184-25-exam-dumps-collection.html

BONUS!!! Download part of DumpsQuestion 1Z0-184-25 dumps for free: https://drive.google.com/open?id=1FUtGryyO4wdVf01ELoxHNP_XnAgElydg