

Oracle AI Vector Search Professional actual exam questions & 1Z0-184-25 valid study dumps & Oracle AI Vector Search Professional test torrent

astramieri/oracle-ai-vector-search



Note-taking repository for Oracle AI Vector Search Professional exam (1Z0-184-25)

1 Contributor 0 Issues 9 Stars 8 Forks

BONUS!!! Download part of Lead1Pass 1Z0-184-25 dumps for free: <https://drive.google.com/open?id=1d0wuMHof4pvXBsHY3UpVqWWbIGnPsQ0T>

To want to pass Oracle 1Z0-184-25 certification test can't be done just depend on the exam related books. Instead of blindly studying relevant knowledge the exam demands, you can do some valuable questions. The efficient exam dumps is essential tool to prepare for 1Z0-184-25 test. Come on and purchase Lead1Pass Oracle 1Z0-184-25 Practice Test dumps. This braindump's hit accuracy is high and it works best the other way around. Lead1Pass Oracle 1Z0-184-25 questions and answers are a rare material which can help you pass you exam first time.

Oracle 1Z0-184-25 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none">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 2	<ul style="list-style-type: none">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 3	<ul style="list-style-type: none">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.
Topic 4	<ul style="list-style-type: none">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.

>> Pass 1Z0-184-25 Test Guide <<

Latest Oracle 1Z0-184-25 Learning Materials & New 1Z0-184-25 Exam Format

It is universally acknowledged that Oracle certification can help present you as a good master of some knowledge in certain areas, and it also serves as an embodiment in showcasing one's personal skills. However, it is easier to say so than to actually get the Oracle certification. We have to understand that not everyone is good at self learning and self-discipline, and thus many people need outside help to cultivate good study habits, especially those who have trouble in following a timetable. To handle this, our 1Z0-184-25 test training will provide you with a well-rounded service so that you will not lag behind and finish your daily task step by step. At the same time, our 1Z0-184-25 study torrent will also save your time and energy in well-targeted learning as we are going to make everything done in order that you can stay focused in learning our 1Z0-184-25 study materials without worries behind. We are so honored and pleased to be able to read our detailed introduction and we will try our best to enable you a better understanding of our 1Z0-184-25 test training better.

Oracle AI Vector Search Professional Sample Questions (Q15-Q20):

NEW QUESTION # 15

Which SQL statement correctly adds a VECTOR column named "v" with 4 dimensions and FLOAT32 format to an existing table named "my_table"?

- A. UPDATE my_table SET v = VECTOR(4, FLOAT32)
- B. **ALTER TABLE my_table ADD (v VECTOR(4, FLOAT32))**
- C. ALTER TABLE my_table MODIFY (v VECTOR(4, FLOAT32))
- D. ALTER TABLE my_table ADD v VECTOR(4, FLOAT32)

Answer: B

Explanation:

To add a new column to an existing table, Oracle uses the ALTER TABLE statement with the ADD clause. Option B, ALTER TABLE my_table ADD (v VECTOR(4, FLOAT32)), correctly specifies the column name "v", the VECTOR type, and its attributes (4 dimensions, FLOAT32 precision) within parentheses, aligning with Oracle's DDL syntax for VECTOR columns. Option A uses MODIFY, which alters existing columns, not adds new ones, making it incorrect here. Option C uses UPDATE, a DML statement for updating data, not a DDL operation for schema changes. Option D omits parentheses around the VECTOR specification, which is syntactically invalid as Oracle requires dimensions and format to be enclosed. The SQL Language Reference confirms this syntax for adding VECTOR columns.

NEW QUESTION # 16

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?

- A. **Using the same embedding model for both vector creation and similarity search**
- B. The physical storage location of the vector data
- C. Regularly updating vector embeddings to reflect changes in the source data
- D. The specific distance algorithm employed for vector comparisons

Answer: A

Explanation:

In Oracle Database 23ai, vector search accuracy hinges on the consistency of the embedding model. The VECTOR data type stores embeddings as fixed-dimensional arrays, and similarity searches (e.g., using VECTOR_DISTANCE) assume that all vectors stored and query-are generated by the same model. This ensures they occupy the same semantic space, making distance calculations meaningful. Regular updates (B) maintain data freshness, but if the model changes, integrity is compromised unless all embeddings are regenerated consistently. The distance algorithm (C) (e.g., cosine, Euclidean) defines how similarity is measured but relies on consistent embeddings; an incorrect model mismatch undermines any algorithm. Physical storage location (D) affects performance, not integrity. Oracle's documentation stresses model consistency as a prerequisite for reliable vector search within its native capabilities.

NEW QUESTION # 17

What are the key advantages and considerations of using Retrieval Augmented Generation (RAG) in the context of Oracle AI Vector Search?

- A. It prioritizes real-time data extraction and summarization from various sources to ensure the LLM always has the most up-to-date information

- B. It excels at optimizing the performance and efficiency of LLM inference through advanced caching and precomputation techniques, leading to faster response times but potentially increasing storage requirements
- C. It leverages existing database security and access controls, thereby enabling secure and controlled access to both the database content and the LLM
- D. It focuses on training specialized LLMs within the database environment for specific tasks, offering greater control over model behavior and data privacy but potentially requiring more development effort

Answer: C

Explanation:

RAG in Oracle AI Vector Search integrates vector search with LLMs, leveraging database-stored data. A key advantage is its use of existing database security and access controls (D), ensuring that sensitive enterprise data remains secure while being accessible to LLMs, aligning with Oracle's security model (e.g., roles, privileges). Performance optimization (A) occurs but isn't the primary focus; storage increases are minimal compared to security benefits. Real-time extraction (B) is possible but not RAG's core strength, which lies in static data augmentation. Training LLMs (C) is unrelated to RAG, which uses pre-trained models. Oracle emphasizes security integration as a standout RAG feature.

NEW QUESTION # 18

What is the function of the COSINE parameter in the SQL query used to retrieve similar vectors?

topk = 3

sql = f'''select payload, vector_distance(vector, :vector, COSINE) as score from {table_name} order by score fetch approximate {topk} rows only'''

- A. It converts the vectors to a format compatible with the SQL database
- B. It indicates that the cosine distance metric should be used to measure similarity between vectors
- C. It filters out vectors with a cosine similarity below a certain threshold
- D. It specifies the type of vector encoding used in the database

Answer: B

Explanation:

In Oracle Database 23ai, the VECTOR_DISTANCE function calculates the distance between two vectors using a specified metric. The COSINE parameter in the query (vector_distance(vector, :vector, COSINE)) instructs the database to use the cosine distance metric (C) to measure similarity. Cosine distance, defined as 1 - cosine similarity, is ideal for high-dimensional vectors (e.g., text embeddings) as it focuses on angular separation rather than magnitude. It doesn't filter vectors (A); filtering requires additional conditions (e.g., WHERE clause). It doesn't convert vector formats (B); vectors are already in the VECTOR type. It also doesn't specify encoding (D), which is defined during vector creation (e.g., FLOAT32). Oracle's documentation confirms COSINE as one of the supported metrics for similarity search.

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. TARGET_ACCURACY
- C. NEIGHBOURS
- D. EFCONSTRUCTION

Answer: D

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

.....

For candidates who are going to attend the exam, some practice is necessary, for the practice can build up the confidence. 1Z0-184-25 exam torrent of us can help you pass the exam successfully. 1Z0-184-25 exam braindumps are edited by professional experts, and the quality can be guaranteed. In addition, 1Z0-184-25 exam materials cover most knowledge points for the exam, and you can master the major knowledge points for the exam, therefore your confidence for the exam will be strengthened. We provide you with free demo for you to have a try before buying 1Z0-184-25 Exam Braindumps, so that you can know what the complete version is like.

Latest 1Z0-184-25 Learning Materials: <https://www.lead1pass.com/Oracle/1Z0-184-25-practice-exam-dumps.html>

BONUS!!! Download part of Lead1Pass 1Z0-184-25 dumps for free: <https://drive.google.com/open?id=1d0wuMHof4pvXBsHY3UpVqWWbIGnPsQ0T>