

1Z0-184-25 exam materials & 1Z0-184-25 practice questions & 1Z0-184-25 study guide

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

P.S. Free & New 1Z0-184-25 dumps are available on Google Drive shared by It-Tests: https://drive.google.com/open?id=1yX3QYcQCpRwsA_FC-hh14E808zAxqK0J

Our 1Z0-184-25 Study Materials are compiled by domestic first-rate experts and senior lecturer and the contents of them contain all the important information about the test and all the possible answers of the questions which maybe appear in the test. You can use the practice test software to check your learning outcomes. Our 1Z0-184-25 study materials' self-learning and self-evaluation functions, the statistics report function, the timing function and the function of stimulating the test could assist you to find your weak links, check your level, adjust the speed and have a warming up for the real exam. You will feel your choice to buy Oracle Database 23ai study materials are too right.

Oracle 1Z0-184-25 Exam Syllabus Topics:

Topic	Details
Topic 1	<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.

Topic 2	<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 3	<ul style="list-style-type: none"> 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 4	<ul style="list-style-type: none"> 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 SQL and Python to integrate AI models with retrieval techniques for enhanced AI-driven decision-making

>> 1Z0-184-25 Exam Testking <<

New 1Z0-184-25 Exam Practice & Valid 1Z0-184-25 Study Notes

When prepare a exam, we may face the situation like this: there are so many books in front of me, which one should I choose for preparing for the exam? If you are ready to attend the 1Z0-184-25 exam, then just choose us, our product is the one you can trust, with the experienced professionals to expect and update, the quality of the product is quite high. Furthermore, our company respect the privacy of the customers, with our product, there is no need for you to worry about the problem. Except for this, if you buy product for the 1Z0-184-25 Exam, you will get the free update for one year, and money back guarantee within 60 days after you buy it, so don't hesitate, just do it.

Oracle AI Vector Search Professional Sample Questions (Q43-Q48):

NEW QUESTION # 43

What security enhancement is introduced in Exadata System Software 24ai?

- A. Integration with third-party security tools
- B. SNMP security (Security Network Management Protocol)
- 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 # 44

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

- A. Creating a new table using CTAS (CREATE TABLE AS SELECT) that includes the VECTOR column from the original table
- B. Modifying the data type of an existing VECTOR column to a non-VECTOR type
- C. Adding a new VECTOR column to the table
- D. Dropping an existing VECTOR column from the table

Answer: B

Explanation:

Oracle Database 23ai imposes restrictions on DDL operations for tables with VECTOR columns to preserve data integrity. CTAS (A) is permitted, as it copies the VECTOR column intact into a new table, maintaining its structure. Dropping a VECTOR column

(B) is allowed via ALTER TABLE DROP COLUMN, as it simply removes the column without altering its type. Adding a new VECTOR column (D) is supported with ALTER TABLE ADD, enabling schema evolution. However, modifying an existing VECTOR column's data type to a non-VECTOR type (C) (e.g., VARCHAR2, NUMBER) is not permitted because VECTOR is a specialized type with dimensional and format constraints (e.g., FLOAT32), and Oracle does not support direct type conversion due to potential loss of semantic meaning and structure. This restriction is documented in Oracle's SQL reference.

NEW QUESTION # 45

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

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

Answer: D

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 VECTOR). Option B (executing searches) confuses the model with database functions like VECTOR_DISTANCE, which use the embeddings, not create them. Option D (storing vectors) pertains to the database's storage engine, not the model's function-storage is handled by Oracle's VECTOR type and indexes (e.g., HNSW). The embedding model's role is purely generative, not operational or structural. In practice, Oracle 23ai integrates this via VECTOR_EMBEDDING, which calls the model to produce vectors, underscoring its transformative purpose. Misunderstanding this could lead to conflating data preparation with query execution, a common pitfall for beginners.

NEW QUESTION # 46

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

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

Answer: A

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 the probability 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 # 47

Which is NOT a feature or capability related to AI and Vector Search in Exadata?

- A. Loading Vector Data using SQL*Loader
- B. Vector Replication with GoldenGate
- C. AI Smart Scan
- **D. Native Support for Vector Search Only within the Database Server**

Answer: D

