

Oracle 1Z0-184-25 Dumps - Try Free 1Z0-184-25 Exam Questions and Answer



2025 Latest TestPassed 1Z0-184-25 PDF Dumps and 1Z0-184-25 Exam Engine Free Share: https://drive.google.com/open?id=1D39eomlyi7-aWJ4BqtSjvwBCP_D_njv4

The learning material is open in three excellent formats; Oracle 1Z0-184-25 dumps PDF, a desktop Oracle 1Z0-184-25 dumps practice test, and a web-based Oracle 1Z0-184-25 dumps practice test. Oracle 1Z0-184-25 dumps is organized by experts while saving the furthest down-the-line plan to them for the Oracle 1Z0-184-25 Exam. The sans bug plans have been given to you all to drift through the Oracle certificate exam.

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">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 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">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 PLSQL and Python to integrate AI models with retrieval techniques for enhanced AI-driven decision-making.

Topic 5	<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.
---------	--

>> 1Z0-184-25 Advanced Testing Engine <<

VCE Oracle 1Z0-184-25 Dumps | Exam 1Z0-184-25 Labs

Our PDF version of our 1Z0-184-25 exam practice guide is convenient for the clients to read and supports the printing. If the clients use our PDF version they can read the PDF form conveniently and take notes. The 1Z0-184-25 quiz prep can be printed onto the papers. If the clients need to take note of the important information they need they can write them on the papers to be convenient for reading or print them on the papers. The clients can read our 1Z0-184-25 Study Materials in the form of PDF or on the printed papers. Thus the clients learn at any time and in any place and practice the 1Z0-184-25 exam practice guide repeatedly.

Oracle AI Vector Search Professional Sample Questions (Q32-Q37):

NEW QUESTION # 32

Which SQL function is used to create a vector embedding for a given text string in Oracle Database 23ai?

- A. `EMBED_TEXT`
- B. `CREATE_VECTOR_EMBEDDING`
- C. `GENERATE_EMBEDDING`
- D. `VECTOR_EMBEDDING`

Answer: D

Explanation:

The `VECTOR_EMBEDDING` function in Oracle Database 23ai generates a vector embedding from input data (e.g., a text string) using a specified model, such as an ONNX model loaded into the database. It's designed for in-database embedding creation, supporting vector search and AI applications. Options A, B, and C (`GENERATE_EMBEDDING`, `CREATE_VECTOR_EMBEDDING`, `EMBED_TEXT`) are not valid SQL functions in 23ai. `VECTOR_EMBEDDING` integrates seamlessly with the `VECTOR` data type and is documented as the standard method for embedding generation in SQL queries.

NEW QUESTION # 33

What is the purpose of the Vector Pool in Oracle Database 23ai?

- A. To manage database partitioning
- B. To enable longer SQL execution
- C. To store HNSW vector indexes and IVF index metadata
- D. To store non-vector data types

Answer: C

Explanation:

The Vector Pool in Oracle 23ai is a dedicated SGA memory region (controlled by `VECTOR_MEMORY_SIZE`) for vector operations, specifically storing HNSW indexes (graph structures) and IVF index metadata (e.g., centroids) (B). This optimizes memory usage for vector search, keeping critical index data accessible for fast queries. Partitioning (A) is unrelated; that's a tablespace feature. Longer SQL execution (C) might benefit indirectly from memory efficiency, but it's not the purpose. Non-vector data (D) resides elsewhere (e.g., PGA, buffer cache). Oracle allocates the Vector Pool to enhance AI workloads, ensuring indexes don't compete with other memory, a design choice reflecting vector search's growing importance.

NEW QUESTION # 34

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

- A. To find exact matches in BLOB data
- B. To group vectors by their exact scores
- C. Optimize relational database operations to compute distances between all data points in a database
- **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 # 35

How is the security interaction between Autonomous Database and OCI Generative AI managed in the context of Select AI?

- A. By encrypting all communication between the Autonomous Database and OCI Generative AI using TLS/SSL protocols
- B. By requiring users to manually enter their OCI API keys each time they execute a natural language query
- **C. By utilizing Resource Principals, which grant the Autonomous Database instance access to OCI Generative AI without exposing sensitive credentials**
- D. By establishing a secure VPN tunnel between the Autonomous Database and OCI Generative AI service

Answer: C

Explanation:

In Oracle Database 23ai's Select AI, security between the Autonomous Database and OCI Generative AI is managed using Resource Principals (B). This mechanism allows the database instance to authenticate itself to OCI services without hardcoding credentials, enhancing security by avoiding exposure of sensitive keys. TLS/SSL encryption (A) is used for data-in-transit security, but it's a complementary layer, not the primary management method. A VPN tunnel (C) is unnecessary within OCI's secure infrastructure and not specified for Select AI. Manual API key entry (D) is impractical and insecure for automated database interactions. Oracle's documentation on Select AI highlights Resource Principals as the secure, scalable authentication method.

NEW QUESTION # 36

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

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

Answer: D

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

NEW QUESTION # 37

VCE 1Z0-184-25 Dumps: <https://www.testpassed.com/1Z0-184-25-still-valid-exam.html>

- P.S. Free & New 1Z0-184-25 dumps are available on Google Drive shared by TestPassed: https://drive.google.com/open?id=1D39eomlyi7-aWJ4BqtSjvwBCP_D_mjv4