Instant 1Z0-184-25 Access & 1Z0-184-25 Latest Dumps Ppt

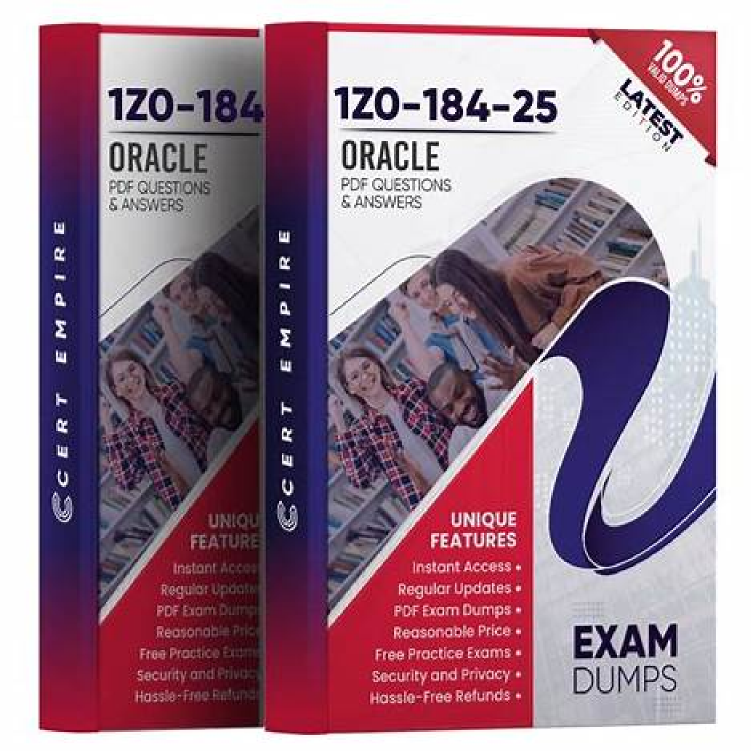

BTW, DOWNLOAD part of GuideTorrent 1Z0-184-25 dumps from Cloud Storage: https://drive.google.com/open?id=1GdGzTBm_gY61lDOdXjE46ZzYQDHf-6UF

All 1Z0-184-25 practice questions you should know are written in them with three versions to choose from: the PDF, the Software and the APP online. At the same time, the experts who compiled the 1Z0-184-25 learning engine are assiduously over so many years in this filed. I can say that our experts have became the authority in this career. And they are good at simplifying the content of the 1Z0-184-25 Exam Braindumps to be understood by our customers all over the world.

Oracle 1Z0-184-25 Exam Syllabus Topics:

Topic	Details
Topic 1	• 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 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	• 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	• 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.

Topic 5	• 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.

>> Instant 1Z0-184-25 Access <<

1Z0-184-25 Latest Dumps Ppt | 1Z0-184-25 Latest Material

We can ensure you a pass rate as high as 99% of our 1Z0-184-25 exam questions. So with our 1Z0-184-25 study guide, you will pass the 1Z0-184-25 exam. And this is the right thing you can imagine. You surely desire the 1Z0-184-25 certification. So with a tool as good as our 1Z0-184-25 Exam Material, why not study and practice for just 20 to 30 hours and then pass the examination? It is more convenient for you to study and practice anytime, anywhere with our varied versions of 1Z0-184-25 exam braindumps.

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

NEW QUESTION # 32
What is created to facilitate the use of OCI Generative AI with Autonomous Database?

- A. A secure VPN tunnel
- B. An AI profile for OCI Generative AI
- C. A new user account with elevated privileges
- D. A dedicated OCI compartment

Answer: B

Explanation:
To integrate OCI Generative AI with Autonomous Database in Oracle 23ai (e.g., for Select AI), an AI profile (A) is created within the database using DBMS_AI. This profile configures the connection to OCI Generative AI, specifying the LLM and authentication (e.g., Resource Principals). A compartment (B) organizes OCI resources but isn't "created" specifically for this integration; it's a prerequisite. A new user account (C) or VPN tunnel (D) isn't required; security leverages existing mechanisms. Oracle's Select AI setup documentation highlights the AI profile as the key facilitator.

NEW QUESTION # 33
When generating vector embeddings outside the database, what is the most suitable option for storing the embeddings for later use?

- A. In a dedicated vector database
- B. In a CSV file
- C. In the database as BLOB (Binary Large Object) data
- D. In a binary FVEC file with the relational data in a CSV file

Answer: A

Explanation:
When vector embeddings are generated outside the database, the storage choice must balance efficiency, scalability, and usability for similarity search. A CSV file (A) is simple and human-readable but inefficient for large-scale vector operations due to text parsing overhead and lack of indexing support. A binary FVEC file (B) offers a compact format for vectors, reducing storage size and improving read performance, but separating relational data into a CSV complicates integration and querying, making it suboptimal for unified workflows. Storing embeddings as BLOBs in a relational database (C) integrates well with structured data and supports SQL access, but it lacks the specialized indexing (e.g., HNSW, IVF) and query optimizations that dedicated vector databases provide. A dedicated vector database (D), such as Milvus or Pinecone (or Oracle 23ai's vector capabilities if internal), is purpose-built for high-dimensional vectors, offering efficient storage, advanced indexing, and fast approximate nearest neighbor (ANN)

searches. For external generation scenarios, where embeddings are not immediately integrated into Oracle 23ai, a dedicated vector database is the most suitable due to its performance and scalability advantages. Oracle's AI Vector Search documentation indirectly supports this by emphasizing optimized vector storage for search efficiency, though it focuses on in-database solutions.

NEW QUESTION # 34

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

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

Answer: D

Explanation:
Exadata in Oracle Database 23ai enhances AI and vector search capabilities. Vector Replication with GoldenGate (B) supports real-time vector data distribution. SQL*Loader (C) loads vector data into VECTOR columns. AI Smart Scan (D) accelerates AI workloads using Exadata's storage optimizations. However, "Native Support for Vector Search Only within the Database Server" (A) is not a feature; vector search is natively supported across Exadata's architecture, leveraging both database and storage layers (e.g., via Smart Scan), not restricted to the server alone. This option misrepresents Exadata's distributed capabilities, making it the correct "NOT" answer.

NEW QUESTION # 35

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

- A. HAMMING_DISTANCE
- B. COSINE_DISTANCE
- C. L1_DISTANCE
- D. L2_DISTANCE

Answer: D

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 $\sqrt{\sum(xi - yi)^2}$, where xi and yi 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 ($\sum|xi - yi|$), 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.

NEW QUESTION # 36

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?

- A. Change the index type to HNSW for better accuracy
- B. Increase the VECTOR_MEMORY_SIZE initialization parameter
- C. Add the TARGET_ACCURACY clause to the query with a higher value for the accuracy
- D. Re-create the index with a higher EFCONSTRUCTION value

Answer: C

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.

NEW QUESTION # 37
......