

Oracle.1z0-1195-26.v2026-10-19.q23

Exam Code:	1z0-1195-26
Exam Name:	Oracle AI Database Foundations Associate
Certification Provider:	Oracle
Free Question Number:	23
Version:	v2026-10-19
# of views:	138
# of Questions views:	230
https://www.freepdfdumps.com/Oracle.1z0-1195-26.v2026-10-19.q23.html	

NEW QUESTION: 1

What does Oracle mean by a converged database strategy?

- A. One reporting tool replaces SQL, JSON APIs, and graph queries.
- B. One application server manages all databases without schema changes.
- C. One storage tier is reserved only for AI workloads and vector indexes.
- D. One database architecture supports multiple data types and workloads together.

Answer: D (LEAVE A REPLY)

Comprehensive and Detailed 100 to 150 words of Explanation From Oracle AI Database topics:

A converged database strategy means using one database engine to support multiple modern data models and workload types rather than deploying a separate specialized database for each requirement. Oracle AI Database provides native support for relational, JSON/document, vector, graph, spatial, text, and other data, while supporting transactional, analytic, AI Vector Search, and mixed workloads. This reduces data movement, synchronization, security fragmentation, and operational complexity. A converged database does not mean one reporting tool replaces every access language, one application server manages unrelated databases, or one storage tier is reserved only for AI. Those choices confuse application tooling or storage design with the database architecture itself. Option D accurately expresses Oracle's converged strategy: multiple data types and workloads supported together in a unified database platform. Oracle

NEW QUESTION: 2

Which pair correctly matches an AI domain to an example?

- A. Vision - image classification
- B. Language - object detection in photos
- C. Forecasting - entity extraction from text

D. Speech - product demand prediction

Answer: (SHOW ANSWER)

Vision - image classification is the correctly matched AI domain and use case. The uploaded source explicitly identifies option A as correct. Oracle Cloud Infrastructure Vision documentation confirms that Vision performs image analysis and includes image-classification capabilities for identifying objects and scene-based characteristics in images. The distinction among the answer choices is based on the type of input being analyzed and the objective of the AI model. Computer vision works with images and visual content; classification assigns labels or categories based on visual characteristics. Language capabilities operate primarily on natural-language text and support functions such as entity recognition, sentiment analysis, text classification, and key-phrases extraction. Therefore, object detection in photographs belongs to vision rather than language.

Likewise, forecasting predicts future numerical or temporal outcomes from historical patterns; product-demand prediction is a typical forecasting scenario. Speech focuses on spoken audio, such as transcription or speech recognition, rather than business-demand prediction.

Option A is therefore the only domain/example relationship that is semantically and technically aligned.

Oracle Vision explicitly supports image classification, making the mapping unambiguous. Study Guide reference: Identifying Oracle AI Database 26ai Strategy and Autonomous AI Database Basics - AI domains, vision, language, speech, forecasting, and practical AI use cases.

NEW QUESTION: 3

Which set of characteristics is commonly used to describe Autonomous AI Database?

- A.** Self-training, self-upgrading, and self-documenting
- B.** Self-hosted, self-audited, and self-architected
- C.** Self-coding, self-licensing, and self-networking
- D.** Self-driving, self-securing, and self-repairing

Answer: (SHOW ANSWER)

The uploaded question set identifies self-driving, self-securing, and self-repairing as the defining Autonomous AI Database characteristics. Oracle documentation uses the same terminology when describing Autonomous Database services.

Self-driving refers to automated database-management activities that traditionally require significant DBA intervention, including provisioning, tuning, optimization, backups, patching, and scaling. Self-securing encompasses automated security practices designed to protect database infrastructure and data, including encryption and security maintenance. Self-repairing describes automated availability and fault-management capabilities intended to reduce downtime and recover from infrastructure or database failures with minimal manual involvement.

These attributes are central to Oracle's Autonomous Database strategy because the service shifts routine infrastructure and database operations from customer-managed procedures toward automated cloud-service capabilities. This allows development and data teams to focus on application logic and business workloads rather than routine database administration.

The alternative combinations are not Oracle's established characterization. Autonomous AI Database does not promise autonomous application coding, licensing, documentation creation, or architectural design. Those terms incorrectly broaden the scope of database automation beyond the capabilities Oracle associates with the Autonomous platform.

Study Guide reference: Identifying Oracle AI Database 26ai Strategy and Autonomous AI Database Basics - Autonomous characteristics and automated database operations.

NEW QUESTION: 4

What is the main difference between Autonomous AI Database Serverless and Dedicated deployment choices?

- A.** Serverless requires customers to manage isolated Exadata infrastructure, while Dedicated is fully shared by default
- B.** Serverless is used only for JSON workloads, while Dedicated is used only for relational workloads
- C.** Serverless is simple and elastic, while Dedicated provides isolated private resources and more customization.
- D.** Serverless and Dedicated differ only in the SQL tools available to developers

Answer: (SHOW ANSWER)

The fundamental distinction is that Serverless emphasizes simplicity and elasticity , whereas Dedicated provides isolated infrastructure and greater operational customization . This is the answer explicitly identified in the uploaded question set. Oracle documentation describes the Serverless model as ultra-simple and elastic: customers manage the Autonomous AI Database while Oracle manages the underlying Exadata infrastructure. Dedicated, by contrast, provides exclusive compute, storage, network, and database resources.

Oracle also characterizes Dedicated as a private-cloud-in-public-cloud deployment model with high levels of security isolation and governance. Dedicated environments can support customizable operational policies involving workload placement, update scheduling, availability, capacity usage, and other infrastructure-level concerns. Serverless removes much of that infrastructure planning and is therefore well suited to organizations prioritizing rapid provisioning and elastic consumption.

Neither deployment is restricted exclusively to JSON or relational workloads, and the distinction is not primarily about available developer SQL tools. Option A reverses the infrastructure characteristics: it is Dedicated-not Serverless-that supplies the isolated dedicated resource model.

Study Guide reference: Identifying Oracle AI Database 26ai Strategy and Autonomous AI Database Basics - Serverless versus Dedicated deployment architecture.

NEW QUESTION: 5

What is the operational benefit of using a converged database instead of several point-solution databases?

- A. It automatically removes the need to review security policies for database workloads
- B. It guarantees that all workloads use the same physical schema and access method
- C. It can reduce operational complexity by using a more consistent security and patching approach
- D. It requires every application team to maintain a separate synchronization process

Answer: C (LEAVE A REPLY)

A converged database can reduce operational complexity because multiple data models and workload types can use a common database platform with a more consistent approach to security, upgrades, patching, and maintenance. The uploaded question source explicitly identifies option C. Oracle documentation states that Oracle AI Database is a converged, multimodel database and specifically highlights a common approach for security, upgrades, patching, and maintenance.

With separate point-solution databases, organizations may need distinct administrator skill sets, identity configurations, encryption mechanisms, backup procedures, monitoring systems, patch schedules, and replication pipelines. Consolidating appropriate workloads onto a converged database can remove portions of that duplicated operational footprint while allowing relational, JSON, graph, spatial, vector, and other capabilities to work together.

However, convergence does not mean security policies can stop being reviewed.

Governance remains necessary. It also does not require every workload to use an identical physical schema or application-access mechanism; Oracle supports multiple data models and APIs precisely because applications have different access requirements.

Option D describes a disadvantage of fragmented point solutions: separate databases frequently introduce synchronization processes. The converged architecture is intended to reduce, rather than require, such duplication.

Study Guide reference: Identifying Oracle AI Database 26ai Strategy and Autonomous AI Database Basics - converged architecture, operational simplification, security, patching, and maintenance.

NEW QUESTION: 6

Which set is included in the Oracle AI Database 26ai converged approach?

- A. Relational, JSON, vector, object-only files, and separate document stores
- B. JSON, document APIs, storage buckets, monitoring, and billing
- C. Relational, XML, spatial, external-only graph stores, and separate vector databases
- D. Relational, JSON, vector, graph, and spatial

Answer: D ([LEAVE A REPLY](#))

Oracle AI Database 26ai follows a converged database strategy in which multiple data models and workload capabilities are supported within a unified database platform. Oracle describes 26ai as supporting AI, microservices, graph, document, spatial, and relational applications within one converged database. AI Vector Search extends that approach by adding native vector storage and similarity search alongside established relational, JSON, graph, text, and spatial functionality.

The architectural objective is to avoid creating separate point-solution databases whenever an application needs a different representation of data. For example, structured customer information can remain relational, document-oriented information can use JSON, relationships can be analyzed as graphs, geospatial information can use Spatial capabilities, and semantic representations can reside as vectors. These capabilities can then participate in integrated queries instead of forcing applications to continuously replicate and synchronize data among independent database products.

Options A and C contradict this strategy by introducing separate document, graph, or vector stores. Option B mixes workload capabilities with infrastructure and operational concepts such as storage buckets and billing rather than identifying the database's core converged data models. Option D correctly enumerates relational, JSON, vector, graph, and spatial capabilities and matches the source question's designated answer.

Study Guide reference: Identifying Oracle AI Database 26ai Strategy and Autonomous AI Database Basics - converged database architecture and multi-model data support.

NEW QUESTION: 7

When comparing a query vector with stored vectors, what does a smaller vector distance indicate?

- A.** The items require a larger maintenance window
- B.** The items are more semantically similar
- C.** The stored vectors were created with different models
- D.** The query must switch from SQL to GraphQL

Answer: ([SHOW ANSWER](#))

Vector distance quantifies how far apart two vector representations are according to a specified mathematical distance metric. In semantic search, embeddings place semantically related content near one another in a multidimensional vector space.

Consequently, when a distance-oriented metric such as cosine distance or Euclidean distance produces a smaller value, the vectors are considered closer and therefore generally more similar according to the embedding model. Oracle's cosine-distance documentation explains the inverse relationship between cosine similarity and cosine distance: increasingly similar vector directions produce smaller cosine-distance values.

For example, cosine distance is calculated as $1 - \text{cosine similarity}$. Identical or highly aligned semantic representations therefore approach a distance of zero, while increasingly different vectors produce larger distances. This principle is why Oracle SQL similarity

searches commonly order rows by VECTOR_DISTANCE(...) in ascending order and fetch the first N rows: the rows with the smallest distances are the nearest semantic matches. Distance has no relationship to database maintenance windows, and a smaller value does not prove that different embedding models were used. Vector search also remains available through SQL and does not require GraphQL.

Study Guide reference: Working with AI and Vector Foundations - vector embeddings, distance metrics, semantic similarity, and top-K ranking.

NEW QUESTION: 8

A developer wants one place to work with SQL, REST endpoints, JSON features, APEX, and machine learning tools. Which entry point should the developer use?

- A. Customer-managed key rotation
- B. Database Actions
- C. Vector target accuracy settings
- D. Property Graph View definition only

Answer: B (LEAVE A REPLY)

Database Actions is the correct centralized entry point because it is the browser-based development and administration environment bundled with Autonomous AI Database. Oracle documents Database Actions as providing development, data, administration, monitoring, and download capabilities. Its Development area includes SQL, Data Modeler, REST, JSON, Oracle Machine Learning, Graph Studio, and Oracle APEX, which directly matches the developer's requirement for a single integrated workspace. The SQL worksheet supports SQL and PL/SQL development, while REST and JSON tools expose database data through modern application interfaces. APEX provides low-code application development, and Oracle Machine Learning integrates analytical and machine-learning workflows with database-resident data.

The other choices are individual configuration concepts rather than comprehensive developer entry points.

Customer-managed key rotation relates to encryption-key lifecycle management; vector target accuracy controls approximate vector-search behavior; and a Property Graph View is a graph modeling construct.

Therefore, none provides the breadth of tooling required. The uploaded question set likewise identifies Database Actions as the correct response.

Study Guide reference: Using Oracle Database Actions and Data Studio Tools - Database Actions development tools and integrated workspaces.

NEW QUESTION: 9

A data analyst wants to discover data, load it into the platform, and prepare it for later use without leaving the built-in tools. Which workspace fits this task?

- A. Data Studio
- B. OCI Vault

C. VECTOR_DISTANCE function only

D. Property Graph View definition

Answer: (SHOW ANSWER)

Data Studio is the correct built-in workspace because its purpose directly covers data discovery, loading, preparation, transformation, analysis, cataloging, sharing, and related self-service workflows. The uploaded source designates Data Studio as the correct answer. Oracle's current Autonomous AI Database documentation states that Data Studio can be used to load, discover, catalog, transform, analyze, share, enrich, and automate data workflows through a web-based interface.

Data Studio is available through Database Actions and contains specialized tools including Data Load , Catalog , Data Transforms , Data Analysis , and Data Insights . Data Load can ingest or link data from local files, remote databases, cloud storage, and live feeds, while Catalog supports searching, discovering, inspecting, and managing available data assets. OCI Vault addresses secrets and encryption-key management rather than data preparation.

VECTOR_DISTANCE is an individual SQL function used for vector comparisons, not an integrated analyst workspace. A Property Graph View defines a graph representation over data but likewise does not provide general discovery and loading functionality.

Study Guide reference: Using Oracle Database Actions and Data Studio Tools - Data Studio, Data Load, Catalog, Data Transforms, and self-service data preparation.

NEW QUESTION: 10

A retailer wants to find products semantically similar to a shopper's description, but only from items that are in stock and sold in the shopper's region. Which design meets this requirement?

A. Use JSON Duality Views because similarity search requires document-only access

B. Use a single SQL statement that combines stock and region filters with vector ranking

C. Use a property graph because vector search cannot work with business predicates

D. Use vector search only, then manually remove unavailable products outside the database

Answer: B (LEAVE A REPLY)

Oracle AI Vector Search is designed to combine semantic similarity with conventional business predicates inside the same SQL statement. Oracle's native VECTOR data type and vector-distance operators allow embeddings to coexist with relational attributes such as inventory status, region, category, price, or security classification. The application can therefore restrict rows using ordinary SQL predicates-for example, in_stock = 'Y' and region = :region-while ranking qualifying products using vector similarity or distance.

Oracle explicitly positions the converged database architecture as enabling vector similarity searches together with relational, JSON, graph, text, and spatial criteria in a single database query.

JSON Duality Views are not a prerequisite for semantic search, and a property graph does not replace the vector engine. Performing vector retrieval first and filtering unavailable inventory in application code is also inferior because it wastes retrieval capacity and can distort the top-K result set. Applying business filters and semantic ranking together keeps processing close to the data and produces the appropriate eligible top matches. The uploaded source identifies this integrated SQL design as the correct choice.

Study Guide reference: Implementing Select AI and AI Vector Search in Autonomous AI Database - vector similarity search combined with relational filtering.

NEW QUESTION: 11

What is an embedding in the context of AI Vector Search?

- A. A text-only backup of a JSON collection
- B. A maintenance schedule attached to a database instance
- C. A graph edge label stored in a property graph
- D. A sequence of numbers that represents semantic content

Answer: D (LEAVE A REPLY)

An embedding is a numerical vector representation of the meaning or characteristics of some source content.

Oracle describes vector embeddings as mathematical vector representations that capture semantic meaning for content such as words, documents, images, audio, or other data objects. The embedding model transforms the source input into an ordered sequence of numerical values, placing semantically related items close together within a multidimensional vector space.

This numerical representation enables AI Vector Search. Instead of comparing documents solely by literal keyword equality, Oracle can calculate mathematical distances between embeddings. Smaller distances generally identify objects that are semantically more closely related according to the embedding model. This supports capabilities such as semantic document retrieval, recommendation systems, image similarity, RAG, and natural-language search.

An embedding is therefore not a backup mechanism, database maintenance configuration, or graph label.

Those concepts belong to unrelated database subsystems. The crucial certification distinction is that the original content and its embedding are different representations: the content may be text, image, audio, or another object, while its embedding is a numeric vector representing learned semantic characteristics. The uploaded assessment identifies "a sequence of numbers that represents semantic content" as the correct description.

Study Guide reference: Working with AI and Vector Foundations - embeddings, semantic representations, vector spaces, and similarity search.

NEW QUESTION: 12

Which statement describes Artificial Intelligence?

- A.** Systems use data and algorithms to perform tasks that normally require human intelligence
- B.** Systems eliminate the need for training data in every use case
- C.** Systems are identical to AGI in all current enterprise products
- D.** Systems replace databases by storing all data as graphs

Answer: A (LEAVE A REPLY)

Artificial Intelligence describes computing systems designed to perform functions associated with human intelligence, including learning, reasoning, language understanding, recognition, prediction, and decision support. Oracle characterizes AI as computing systems trained to simulate human intelligence and notes that AI systems can learn from data, solve problems, process diverse inputs, and pursue defined objectives. Oracle also describes AI technologies as systems or machines that mimic human intelligence when performing tasks.

Data and algorithms are central to modern AI implementations. Machine-learning models learn patterns from training data, while generative models can synthesize new outputs based on learned representations and prompts. However, AI does not imply that training or grounding data is unnecessary. Nor is current enterprise AI equivalent to Artificial General Intelligence (AGI); most deployed systems remain specialized around particular tasks and workloads.

AI also does not replace database technology. Oracle AI Database instead integrates AI capabilities with relational, JSON, graph, spatial, vector, and other database models so that intelligent applications can operate directly on governed enterprise data. Oracle AI Database 26ai specifically emphasizes AI capabilities while preserving converged data-management functionality. The uploaded question source marks option A as correct. Study Guide reference: Identifying Oracle AI Database 26ai Strategy and Autonomous AI Database Basics - AI fundamentals and Oracle's AI-enabled database strategy.

NEW QUESTION: 13

In a property graph, what are vertices and edges?

- A.** Vertices are path lengths and edges are SQL functions
- B.** Vertices represent entities, and edges represent relationships between them
- C.** Vertices are storage partitions, and edges are encryption keys
- D.** Vertices are JSON collections, and edges are maintenance policies

Answer: B (LEAVE A REPLY)

A property graph models information as vertices representing entities and edges representing relationships among those entities. Oracle's Property Graph documentation defines a graph as a collection of objects or vertices connected by arrows or edges. Both vertices and edges can contain properties expressed as key-value pairs. This model is particularly effective when understanding relationships is as important as examining individual records.

For example, in a banking model, customer accounts can be vertices while transfers between accounts become edges. In a social network, people are vertices and relationships such as "knows," "follows," or "works with" become edges. Each edge identifies source and destination vertices and can carry a label and additional properties. Oracle AI Database can then perform graph pattern matching and analytics over these relationships.

Path lengths are results or attributes of graph traversal rather than vertices themselves. SQL functions are query constructs, not graph edges. Likewise, database partitions, encryption keys, JSON collections, and maintenance policies do not define the fundamental property-graph data model. The uploaded assessment explicitly identifies entities and relationships as the proper interpretation.

Study Guide reference: Working with JSON and Graph in Oracle AI Database - property graphs, vertices, edges, properties, and relationship analysis.

NEW QUESTION: 14

What does a JSON Duality View enable?

- A.** It replaces SQL with a document-only query language
- B.** It exposes relational data as JSON documents while preserving relational storage
- C.** It converts a graph model into a vector index automatically
- D.** It requires data to be copied into a separate JSON database

Answer: B (LEAVE A REPLY)

A JSON-Relational Duality View allows applications to work with relationally stored information as JSON documents without maintaining a separate document-store copy. The uploaded source identifies this exact capability as the correct answer. Oracle AI Database documentation confirms that a duality view maps relational table data to hierarchical JSON documents that are materialized on demand rather than separately stored. Applications can therefore access and, when permitted, modify the same underlying information either through relational tables or through its document representation.

This architecture preserves relational advantages such as normalization, integrity constraints, SQL processing, and transactional consistency while giving document-oriented applications a natural JSON interface. A change made through an updatable JSON document is reflected in the underlying relational data, and relational changes are correspondingly visible through the duality view.

The feature does not eliminate SQL, transform property graphs into vector indexes, or require synchronization with an independent document database. Oracle specifically positions JSON-Relational Duality as a mechanism for combining relational and document development models around one authoritative data representation.

Study Guide reference: Working with JSON and Graph in Oracle AI Database - JSON-Relational Duality Views, relational storage, and document-oriented access.

NEW QUESTION: 15

A team already uses Database Actions and now needs the part of the environment focused on discovery, load, integration, and cataloging. Which workspace should they open?

- A. Database Users
- B. Data Studio
- C. Graph Studio
- D. SQL Worksheet

Answer: B (LEAVE A REPLY)

Data Studio is the Database Actions workspace specifically designed for data-oriented activities such as loading, discovering, cataloging, transforming, integrating, analyzing, and preparing data. The uploaded source marks Data Studio as the correct option. Oracle's current Autonomous AI Database documentation states that Data Studio enables users to load, discover, catalog, transform, analyze, share, enrich, and automate data workflows through a web-based interface.

Data Studio contains purpose-built tools including Data Load , Catalog , Data Transforms , Data Analysis

, and Data Insights . Catalog provides a central mechanism for browsing, searching, discovering, inspecting, and acting on local or connected data assets, while Data Transforms supports graphical data integration and transformation workflows.

The alternatives perform different functions. Database Users handles user administration. Graph Studio focuses on graph modeling and analytics. SQL Worksheet provides interactive execution of SQL and PL/SQL rather than an integrated discovery, cataloging, and loading environment.

Therefore, when the question combines the terms discovery, load, integration, and cataloging , Data Studio is the definitive Oracle workspace.

Study Guide reference: Using Oracle Database Actions and Data Studio Tools - Data Studio, Data Load, Catalog, and Data Transforms.

NEW QUESTION: 16

What does the Oracle VECTOR data type enable?

- A. Browser-based low-code page design
- B. Automatic JSON schema validation for all APIs
- C. Automatic conversion of every table into a graph model
- D. Native storage of vector values in the database

Answer: D (LEAVE A REPLY)

The Oracle VECTOR data type provides native database storage for vector values used by AI and machine- learning workloads. Oracle AI Database 26ai represents vectors as ordered numerical values with defined dimensionality and element formats. This allows vector embeddings representing text, images, audio, documents, or other content to reside directly alongside conventional business data rather than requiring a separate specialized vector database.

Native vector storage is foundational to Oracle AI Vector Search. Once embeddings are stored in VECTOR columns, SQL can apply vector-distance functions, perform exact or approximate similarity searches, create vector indexes, and combine semantic rankings with relational, JSON, text, spatial, or graph predicates.

Oracle emphasizes that keeping vectors with business data reduces data movement, lowers architecture complexity, and permits similarity searches against current transactional information.

The VECTOR type does not provide APEX page design—that is an Oracle APEX function. It does not universally validate JSON schemas, nor does it automatically convert relational tables into graph structures.

Those are separate Oracle Database capabilities. Consequently, native storage of vector values precisely describes its core function, consistent with the uploaded question source. Study Guide reference: Working with AI and Vector Foundations - VECTOR data type, vector embeddings, vector columns, and AI Vector Search.

Valid 1z0-1195-26 Dumps shared by Actual4test.com for Helping Passing 1z0-1195-26 Exam! Actual4test.com now offer the **newest 1z0-1195-26 exam dumps**, the Actual4test.com 1z0-1195-26 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com 1z0-1195-26 dumps with Test Engine here: https://www.actual4test.com/1z0-1195-26_examcollection.html (54 Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

NEW QUESTION: 17

A recruiting application must find the top job postings that semantically match a candidate's resume, but only in the candidate's city. How should the application meet this requirement?

- A. Export all job postings to a separate vector platform and remove city filtering
- B. Replace job descriptions with graph edge labels only
- C. Use one SQL query that combines relational filters with vector distance ordering
- D. Use Early maintenance so results are ranked sooner

Answer: C (LEAVE A REPLY)

The correct Oracle AI Vector Search design is a single SQL query combining the relational city restriction with vector-distance ranking. The uploaded source explicitly identifies option C. Oracle AI Database's native VECTOR type enables vector similarity searches without moving business information to a separate vector database. Oracle specifically states that vector searches can be combined with sophisticated business- data searches using SQL and the capabilities of its converged database.

The recruiting system can store an embedding of each job description alongside ordinary relational attributes such as city, employer, salary, employment type, and status. The

candidate's resume or search request is converted into a query vector. SQL can then apply a predicate such as `city = :candidate_city`, calculate vector distance against qualifying job embeddings, order the result by that distance, and return the top matches.

Exporting postings to another vector platform unnecessarily introduces data movement and loses the direct integration with current relational attributes. Graph edge labels are intended for relationship modeling, not semantic document representation. "Early maintenance" is unrelated to vector ranking.

The key Oracle AI Database principle being tested is semantic vector retrieval and conventional relational business filtering within one SQL operation .

Study Guide reference: Implementing Select AI and AI Vector Search in Autonomous AI Database - vector-distance ordering, relational predicates, top-K retrieval, and converged SQL.

NEW QUESTION: 18

A support portal must search product manuals semantically while also filtering results by product line and support level stored in relational columns. How does Oracle AI Vector Search support this requirement?

- A. It requires semantic search results to be filtered outside the database after retrieval
- B. It uses document APIs for search while excluding SQL-based business predicates
- C. It supports semantic search with relational business filters in an integrated database approach
- D. It stores structured filter values as graph relationships before search can run

Answer: C (LEAVE A REPLY)

Oracle AI Vector Search supports semantic similarity search together with conventional relational predicates in SQL , making option C correct. The uploaded source explicitly identifies this integrated approach. Oracle AI Database 26ai provides the native VECTOR data type so embeddings can reside directly alongside relational business attributes.

Oracle states that AI-powered vector similarity searches can be combined with business-data searches using SQL and the full capabilities of the converged database.

For the support portal, each manual or document chunk can have an embedding while relational columns identify its product line, entitlement level, version, or support tier. A SQL statement can restrict rows using predicates such as product line and support level while ordering eligible records by vector distance from the user's query embedding.

This approach is superior to retrieving a broad semantic result set and filtering it later in application code.

External post-filtering can waste processing and may remove highly ranked records without correctly replacing them with the next eligible matches.

Document APIs and property graphs are also not prerequisites. Oracle's converged architecture permits relational and vector criteria to operate together directly.

Study Guide reference: Implementing Select AI and AI Vector Search in Autonomous AI Database - hybrid business filtering, vector similarity ranking, and integrated SQL.

NEW QUESTION: 19

Which example is a common AI use case?

- A. Listing database tables alphabetically in a schema browser
- B. Detecting fraudulent transactions from behavior patterns
- C. Rotating an encryption key on a fixed compliance schedule
- D. Manually assigning a subnet CIDR during database provisioning

Answer: B (LEAVE A REPLY)

Comprehensive and Detailed 100 to 150 words of Explanation From Oracle AI Database topics:

Fraud detection is a standard AI use case because machine-learning and anomaly-detection models can identify unusual behavioral patterns across transaction attributes such as amount, account, merchant, location, and historical activity. Oracle documents AI-driven fraud-prevention architectures that analyze transaction behavior and flag suspicious or anomalous activity for investigation. The other options are administrative database or network tasks rather than AI inference problems. Alphabetically listing tables is ordinary metadata browsing; rotating an encryption key is deterministic security administration; and assigning a subnet CIDR is infrastructure configuration. None requires a model to learn patterns from data. Detecting fraudulent transactions from behavior patterns therefore best represents an AI workload and aligns with the Oracle AI Database foundations coverage of practical AI domains and use cases. Oracle Docs

NEW QUESTION: 20

How can developers create vectors for data objects in Oracle AI Database 26ai?

- A. By using a built-in vectorization function
- B. By loading the data into a separate vector service before querying it
- C. By defining every vector as a JSON document manually
- D. By converting all source objects into Property Graph Views first

Answer: A (LEAVE A REPLY)

Oracle AI Database 26ai provides native SQL and PL/SQL facilities for generating vector embeddings from source data. A principal example is the VECTOR_EMBEDDING SQL function, which generates an embedding by applying an embedding or feature-extraction model to an input value. Oracle also provides vector utilities such as UTL_TO_EMBEDDING, DBMS_VECTOR, and DBMS_VECTOR_CHAIN for vectorization, chunking, embedding generation, similarity-search pipelines, and integration with supported embedding providers.

When an embedding model is imported into Oracle AI Database-for example, in supported ONNX form- the database can perform text-to-vector transformation internally. Oracle also supports accessing external embedding providers through REST where appropriate, but sending all data to a separate vector database or service is not a prerequisite. Keeping

vectorization and vector storage within Oracle AI Database can reduce data movement and enables embeddings to remain integrated with the underlying business objects. Vectors also do not need to be manually represented as JSON documents, and Property Graph Views serve a different purpose: modeling entities and relationships. Consequently, the built-in vectorization capability is the direct Oracle-native mechanism described by the question. The uploaded source explicitly identifies option A as correct.

Study Guide reference: Working with AI and Vector Foundations - vector generation, VECTOR_EMBEDDING, vector utilities, embedding models, and native AI Vector Search.

NEW QUESTION: 21

Modern applications often need to work with relational data, JSON documents, graph relationships, and vector embeddings. What challenge does using a different specialized database for each need create?

- A. They can reduce administration because every platform shares one security model
- B. They can avoid all data transformation and synchronization work
- C. They can query all data types together with one built-in model automatically
- D. They can create data silos that make integrated access and management harder

Answer: (SHOW ANSWER)

Using a different point-solution database for each data model can create data silos , increasing integration, synchronization, governance, and operational complexity. The uploaded assessment identifies option D as correct. Oracle AI Database 26ai is explicitly positioned as a converged database platform supporting AI, graph, document, spatial, relational, and other application models within one database architecture.

The problem with separate specialized stores is that an application may need to duplicate relational records into a document database, copy embeddings into a vector database, and maintain relationships in a graph database. These copies must remain synchronized as source data changes. Security controls, backups, monitoring, patching, access policies, and application integrations may also differ across platforms.

Oracle's converged model instead allows different representations and workloads to operate on centrally governed data. Oracle specifically notes that the converged platform provides synergy among multiple data models and enables different types of information to be joined and manipulated together.

Therefore, separate databases do not automatically share security or eliminate transformation. Those are precisely the architectural burdens that convergence is intended to reduce.

Study Guide reference: Identifying Oracle AI Database 26ai Strategy and Autonomous AI Database Basics - converged database strategy and elimination of data silos.

NEW QUESTION: 22

A business wants one data platform where semantic similarity, relational consistency, and SQL-based filtering all work together.

Which statement aligns with this design goal?

- A.** Vector search requires data to be modeled as graph paths before SQL predicates are applied.
- B.** A converged database can keep data consistent while combining vector search with SQL.
- C.** SQL filtering is handled in a separate application step after vector search returns results.
- D.** Semantic search is usually handled after relational filters are removed from the query flow.

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 23

A team needs faster similarity search at scale and accepts approximate top-K results. Which feature should they use?

- A.** A standard B-tree index optimized for exact match and range queries
- B.** A graph view with ONE ROW PER STEP
- C.** A JSON duality view with update restrictions
- D.** A vector index with approximate search

Answer: ([SHOW ANSWER](#))

A vector index with approximate similarity search is designed specifically for high-performance top-K retrieval over large vector collections. Exact vector search calculates distances against all candidate vectors that satisfy the query predicates, which can become computationally expensive at scale. Approximate nearest-neighbor search uses vector indexing structures to reduce the number of candidate vectors evaluated, significantly improving search latency while accepting a controlled trade-off between performance and recall or accuracy. Oracle AI Database supports vector indexes with organizations such as INMEMORY NEIGHBOR GRAPH and NEIGHBOR PARTITIONS and allows administrators to configure target accuracy.

This requirement explicitly states that approximate top-K results are acceptable, making an approximate vector index the intended architecture. A conventional B-tree index is appropriate for scalar equality, ordering, or range-access patterns, not high-dimensional semantic similarity. JSON Duality Views provide document-relational mapping rather than nearest-neighbor acceleration. Property graph views model entities and relationships and likewise do not serve as vector similarity indexes. The uploaded assessment identifies "a vector index with approximate search" as the correct option.

Study Guide reference: Working with AI and Vector Foundations - vector indexes, approximate nearest-neighbor search, top-K retrieval, and target accuracy.

Valid 1z0-1195-26 Dumps shared by Actual4test.com for Helping Passing 1z0-1195-26 Exam! Actual4test.com now offer the **newest 1z0-1195-26 exam dumps**, the Actual4test.com 1z0-1195-26 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com 1z0-1195-26 dumps with Test Engine here: https://www.actual4test.com/1z0-1195-26_examcollection.html (**54** Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)