

Snowflake.SOL-C01.v2026-03-05.q89

Exam Code:	SOL-C01
Exam Name:	Snowflake Certified SnowPro Associate - Platform Certification
Certification Provider:	Snowflake
Free Question Number:	89
Version:	v2026-03-05
# of views:	107
# of Questions views:	969
https://www.freecram.net/torrent/Snowflake.SOL-C01.v2026-03-05.q89.html	

NEW QUESTION: 1

What is the purpose of the Public Schema in Snowflake? (Choose any 2 options)

- A. Storing user-specific data
- B. Providing a default schema for new objects
- C. Schema to store Account level meta data
- D. Storing shared resources accessible by all users

Answer: (SHOW ANSWER)

The Public schema is created by default within every Snowflake database and serves primarily as a shared workspace that is accessible to all users who have the necessary privileges. One of its main roles is to act as a default schema for objects created without explicitly specifying a schema name. This ensures usability and prevents object creation errors when users forget to define a schema. Additionally, the Public schema is commonly used to store shared or common resources that multiple users or teams need to access. However, it should not be used for sensitive or user-specific data because access can be broad. It does not store account-level metadata; that information resides in INFORMATION_SCHEMA views or account-level metadata functions. Overall, the Public schema acts as both a convenient default location and a shared-access area within a Snowflake database.

NEW QUESTION: 2

Which of the following data format supported to unload semi-structured data from Snowflake? (Choose any 2 options)

- A. JSON
- B. CSV
- C. Parquet
- D. XML

Answer: (SHOW ANSWER)

Snowflake supports unloading semi-structured data specifically into JSON and Parquet formats. JSON preserves semi-structured hierarchy and nested fields, making it ideal for interchange, backup, and integration with downstream systems. Parquet provides columnar, compressed, and metadata-rich storage well-suited for analytical reprocessing in data lakes or other systems. CSV is not semi-structured and cannot represent nested data accurately. XML, although a semi-structured format, is not supported by Snowflake as an unload format.

When performing COPY INTO <location>, users can specify FILE_FORMAT with TYPE=JSON or TYPE=PARQUET to export semi-structured content efficiently.

NEW QUESTION: 3

What is the purpose of the COPY INTO SQL command?

- A. To copy data into a virtual warehouse
- B. To copy data from one table into another
- C. To load data from a stage into a table
- D. To duplicate all files on a stage

Answer: ([SHOW ANSWER](#))

COPY INTO <table> loads data from staged files (internal or external stages) into a Snowflake table.

It supports pattern matching, file formats, error-handling modes, and parallel loading.

It does not move files, does not load to a warehouse (warehouses only supply compute), and table-to-table copying uses INSERT INTO...SELECT.

NEW QUESTION: 4

What is the PRIMARY purpose of the use of the PARSE_DOCUMENT function in Snowflake?

- A. To identify any Personally Identifiable Information (PII) in text
- B. To identify data that will benefit from the use of a directory table
- C. To extract text from PDF files
- D. To parse JSON data

Answer: ([SHOW ANSWER](#))

The PARSE_DOCUMENT function is part of Snowflake Cortex AI and is designed specifically to extract text, layout information, and structured elements from unstructured documents, especially PDFs. It supports OCR-based extraction for scanned files and layout-aware extraction to preserve tables, headings, and format structure.

Its purpose is not PII detection; Snowflake does not provide built-in automatic PII identification via PARSE_DOCUMENT. It does not identify candidate data for directory tables and is unrelated to JSON parsing-Snowflake uses PARSE_JSON for JSON data.

PARSE_DOCUMENT is primarily used for workflows such as contract analysis, invoice extraction, document classification, compliance automation, and downstream AI enrichment.

NEW QUESTION: 5

What are the key benefits of using Snowflake? (Select TWO).

- A. Independent scaling of compute and storage
- B. Vendor lock-in
- C. Manual infrastructure management
- D. Near-unlimited concurrency
- E. On-premises deployment

Answer: ([SHOW ANSWER](#))

Snowflake's multi-cluster shared data architecture enables independent scaling of compute and storage, meaning workloads can be sized differently without affecting stored data.

It also provides near-unlimited concurrency, because Snowflake can automatically add compute clusters during peak loads.

Snowflake does not require manual infrastructure management and runs fully in the cloud-not on-premises.

Multi-cloud support reduces vendor lock-in, not increases it.

NEW QUESTION: 6

How can Snowsight be used to monitor the performance of a virtual warehouse?

- A. Use the SHOW WAREHOUSES command.
- B. Use the LOAD_HISTORY view.
- C. Use the DESCRIBE WAREHOUSE command.
- D. Use the QUERY_HISTORY view.

Answer: ([SHOW ANSWER](#))

To monitor virtual warehouse performance in Snowsight, the correct approach is using the QUERY_HISTORY view or its Snowsight equivalent under the Activity section.

QUERY_HISTORY provides detailed metrics on execution time, queuing, credit consumption, compilation time, and warehouse utilization for every executed query. By analyzing this information, users can diagnose issues such as concurrency bottlenecks, under-sized warehouses, long-running queries, or excessive I/O.

SHOW WAREHOUSES returns metadata about warehouse configurations (size, state, scaling policy) but does not provide performance data. LOAD_HISTORY tracks COPY INTO operations only and is not representative of general warehouse workload. DESCRIBE WAREHOUSE shows configuration details but no runtime or historical performance metrics.

Thus, the most accurate and complete way to assess warehouse performance-particularly execution delays, queuing, or excessive credit usage-is by using QUERY_HISTORY.

NEW QUESTION: 7

Which of the following are benefits of using a multi-clustered warehouse in Snowflake? (Choose any 3 options)

- A. Improved query performance during peak load
- B. Automatic scaling of clusters based on query demand
- C. Reduced storage costs
- D. Increased compute capacity

Answer: (SHOW ANSWER)

A multi-cluster warehouse raises Snowflake's capability to handle concurrent workloads efficiently. When large numbers of users or queries hit the system, Snowflake can automatically start additional clusters (scaling out) to distribute workloads, reducing queuing and improving performance. As query demand drops, Snowflake scales in by shutting down clusters to conserve credits. This dynamic auto-scaling provides significantly increased compute capacity during peak usage periods and enhances responsiveness.

Importantly, multi-cluster warehouses affect only compute, not storage, so they do not reduce storage costs.

Their primary benefits revolve around performance stability, concurrency handling, workload isolation, and compute elasticity.

NEW QUESTION: 8

When unloading data from a Snowflake table, what is the target destination for the COPY INTO @stage command?

- A. An internal or external stage
- B. Another Snowflake table
- C. A local file system
- D. A data stream

Answer: (SHOW ANSWER)

COPY INTO @stage writes query results or table data into files stored in either internal or external stages.

Unloading cannot target a local filesystem directly, cannot load into another table (use INSERT instead), and does not write to streams.

NEW QUESTION: 9

Which of the following is a potential use case for the CLASSIFY_TEXT function?

- A. Generating SQL code from natural language description
- B. Categorizing customer support tickets based on topic
- C. Translating a user manual into multiple languages
- D. Extracting data from scanned invoices

Answer: (SHOW ANSWER)

CLASSIFY_TEXT is a Snowflake Cortex task-specific function designed to categorize free-form text into predefined buckets. It operates by applying LLM-driven classification logic based on category labels or descriptions provided by the user. A primary and practical use case is automatically categorizing customer support tickets based on content, such as identifying billing issues, login problems, or product defects. This enables automated routing, prioritization, and operational analytics.

The function outputs structured results with the predicted label and confidence score. It is optimized for domain classification rather than open-ended text generation.

Incorrect options:

- * Generating SQL code is handled by the COMPLETE function.
 - * Translation is handled by TRANSLATE.
 - * Extracting structured fields from documents is the role of PARSE_DOCUMENT.
- Thus, classification of incoming text streams is the natural use case of CLASSIFY_TEXT.

NEW QUESTION: 10

What file extension is commonly used for Snowflake notebooks?

- A. .ipnb
- B. .sql
- C. .ipynb
- D. .txt

Answer: (SHOW ANSWER)

Snowflake notebooks use the .ipynb file extension, the standard format for Jupyter notebooks. This format stores executable code, markdown, metadata, and cell outputs in a structured JSON layout. Snowflake adopts this format to ensure compatibility with the broader Python ecosystem, thereby enabling seamless migration between Snowflake and external notebook environments. The .ipynb structure allows mixed SQL and Python cells, visualizations, Streamlit components, documentation, and stepwise development within Snowsight. It supports reproducibility, collaboration, and integration with Snowpark and Cortex.

Incorrect formats:

- * .ipnb is a misspelling and invalid.
- * .sql is used for SQL scripts only.
- * .txt cannot represent notebook metadata or cell structure.

Thus, .ipynb is the correct and only supported notebook format.

NEW QUESTION: 11

What is the purpose of auto-resume on a virtual warehouse?

- A. To restart the warehouse manually after suspension.
- B. To automatically start the warehouse when a query is submitted to it.
- C. To automatically scale up the warehouse size.
- D. To prevent the warehouse from suspending.

Answer: (SHOW ANSWER)

When auto-resume is enabled, Snowflake automatically starts a suspended warehouse when a query or task is submitted. This ensures smooth operations without manual intervention. It does not scale the warehouse, prevent suspension, or require manual restarting.

NEW QUESTION: 12

What are Snowflake notebooks primarily used for?

- A. Storing raw data
- B. User authentication

C. Writing and executing SQL queries, Python commands, building visualizations, and documenting analysis

D. Managing cloud storage

Answer: (SHOW ANSWER)

Snowflake Notebooks provide an interactive environment where developers and analysts can combine SQL, Python, visualization libraries, and documentation. Use cases include exploratory analysis, ETL prototyping, data engineering workflows, machine learning model development, and dashboard creation using Streamlit components.

Notebooks allow mixing of SQL and Python cells with shared session state, enabling smooth transitions between Snowpark DataFrames, SQL queries, visual charts, and markdown explanations. They enhance reproducibility and collaboration, supporting versioning, parameterization, and seamless Snowflake compute integration.

Incorrect options:

* Raw data is stored in Snowflake tables, not in notebooks.

* User authentication is handled by the Cloud Services Layer.

* Managing cloud storage is automatic and not a notebook responsibility.

Thus, Snowflake notebooks are an end-to-end development and analytics interface.

NEW QUESTION: 13

What is the SHOW FILE FORMATS SQL command used for in Snowflake?

A. Create a list of all of the file formats currently supported by Snowflake.

B. Get information about how different file formats should be used.

C. List the file formats in the currently-selected schema.

D. List all of the files in a database that have data in a given file format.

Answer: C (LEAVE A REPLY)

SHOW FILE FORMATS lists file format objects that exist in the currently selected schema (or optionally, a specific database or schema). These objects define how Snowflake interprets staged files during data loading.

It does not list supported built-in formats (CSV, Parquet, etc.).

It does not show example usage information and does not list files-LIST @stage handles that.

NEW QUESTION: 14

What are the steps to create a new schema in Snowsight?

A. Data >> Add Data >> Create Schema

B. Data >> Create Schema

C. Data >> Database >> Create Schema

D. Data Product >> Create Schema

Answer: (SHOW ANSWER)

To create a new schema in Snowsight, users must follow the hierarchical navigation of Snowflake's object structure. The correct sequence is:

Data # Database # Create Schema

Steps:

- * Open the `Datab` in the left navigation.
- * Select a `Database` where the schema will be created.
- * Choose `Create # Schema`.
- * Enter the schema name and (optionally) description.
- * Click `Create`.

Incorrect options:

- * "Add Data" is used for loading files, not creating schema objects.
- * "Data # Create Schema" is incomplete because Snowflake must know which database the schema belongs to.
- * "Data Product" is unrelated to object creation.

Thus, Snowflake requires explicit database context before schema creation.

NEW QUESTION: 15

Which layer of Snowflake's architecture provides security features like encryption and authentication?

- A. Compute Layer
- B. Data Storage Layer
- C. Cloud Services Layer
- D. Query Processing Layer

Answer: (SHOW ANSWER)

The Cloud Services Layer is responsible for Snowflake's global security, coordination, and metadata management. It orchestrates all authentication and authorization processes, including username/password login, MFA, OAuth, SSO, key-pair authentication, and integration with enterprise identity providers. It also manages role-based access control (RBAC), privilege enforcement, session management, and secure query orchestration.

This layer centralizes encryption policy enforcement, including the management of Snowflake-managed and customer-managed encryption keys (Tri-Secret Secure). While all data stored and transmitted in Snowflake is encrypted by default, it is the Cloud Services Layer that handles the key lifecycle, rotation, and hierarchy.

The Compute Layer only runs queries.

The Storage Layer keeps compressed, encrypted columnar data but does not manage security controls.

The Query Processing Layer executes SQL but does not manage authentication or RBAC.

Thus, the Cloud Services Layer is the authoritative control plane for Snowflake security.

NEW QUESTION: 16

What are compute resources called in Snowflake?

- A. Data Nodes
- B. Virtual Warehouses
- C. Compute Clusters

D. Virtual Machines

Answer: (SHOW ANSWER)

Snowflake compute resources are referred to as Virtual Warehouses. A virtual warehouse is a cluster of compute nodes that executes SQL queries, performs DML operations (INSERT/UPDATE/DELETE), and runs data loading or transformation tasks.

Virtual Warehouses provide:

- * Dedicated compute isolation
- * Independent scaling (resize at any time)
- * Concurrency support through multi-cluster mode
- * Auto-suspend and auto-resume for cost efficiency

While Virtual Warehouses consist of compute clusters under the hood, Snowflake abstracts the underlying VM and node architecture, exposing only the warehouse construct to users. This ensures simplicity and avoids operational burdens such as node management.

Incorrect terms like Data Nodes or Virtual Machines represent underlying infrastructure concepts not exposed to end users.

Valid SOL-C01 Dumps shared by ExamDiscuss.com for Helping Passing SOL-C01 Exam! ExamDiscuss.com now offer the **newest SOL-C01 exam dumps**, the ExamDiscuss.com SOL-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com SOL-C01 dumps with Test Engine here: <https://www.examdiscuss.com/Snowflake/exam/SOL-C01/premium/> (218 Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 17

What SQL command will return a list of all files in a stage?

- A. PUT
- B. LIST
- C. GET
- D. COPY INTO

Answer: (SHOW ANSWER)

LIST @stage_name displays all files present in a stage (internal or external).

PUT uploads local files to an internal stage.

GET downloads files from an internal stage.

COPY INTO loads or unloads data but does not list files.

NEW QUESTION: 18

What is the primary purpose of the LIMIT clause in a Snowflake query?

- A. To sort the result set
- B. To filter the result set

- C. To restrict the number of rows returned
- D. To group the result set

Answer: ([SHOW ANSWER](#))

The LIMIT clause limits the number of rows Snowflake returns from a query. It is commonly used to preview datasets, improve performance when exploring large tables, and paginate results. LIMIT may be combined with OFFSET to skip a number of rows. Sorting is performed with ORDER BY, filtering with WHERE, and grouping with GROUP BY-not with LIMIT. Therefore, restricting returned rows is the primary purpose of LIMIT.

NEW QUESTION: 19

How does Snowflake's compute layer handle query execution?

- A. By optimizing data in cloud storage
- B. With shared-disk architecture
- C. Using MPP (massively parallel processing) compute
- D. Using single-threaded processing

Answer: ([SHOW ANSWER](#))

Snowflake's compute layer uses Massively Parallel Processing (MPP), meaning queries are divided into smaller tasks distributed across multiple compute nodes in the Virtual Warehouse. Each node processes a portion of the data simultaneously, maximizing parallelism and drastically reducing query times.

Although Snowflake uses a central storage layer (shared-disk model), the compute engine behaves like a shared-nothing MPP system, where each node handles local processing independently, minimizing contention.

Incorrect options:

- * Snowflake does not rely on single-thread execution.
- * Storage optimization occurs at the Storage Layer, not compute.
- * Snowflake does not use traditional shared-disk execution; compute nodes work in parallel independently.

This architecture enables high performance for large analytical workloads.

NEW QUESTION: 20

Which function would you use to determine the sentiment of a customer review in Snowflake Cortex?

- A. TRANSLATE
- B. COMPLETE
- C. SENTIMENT
- D. PARSE_DOCUMENT

Answer: ([SHOW ANSWER](#))

The SENTIMENT function in Snowflake Cortex analyzes text and returns a numerical sentiment score between -1 and 1. This value indicates whether the text expresses negative, positive, or neutral sentiment. It accepts English-language text and uses Snowflake's managed AI models to infer polarity based on contextual understanding. A negative score represents negative emotion,

a positive score represents positive emotion, and values near zero indicate neutrality. It is ideal for customer reviews, feedback analysis, and text mining.

TRANSLATE is used for language translation, COMPLETE for text generation, and PARSE_DOCUMENT for extracting text from documents-not sentiment analysis. SENTIMENT provides fast, in-database AI inference without external model hosting or integration.

NEW QUESTION: 21

What parameter is used to define how long Time Travel can be used to access a table?

- A. DATE_OUTPUT_FORMAT
- B. DATA_RETENTION_TIME_IN_DAYS
- C. TIMEZONE
- D. USE_CACHED_RESULT

Answer: (SHOW ANSWER)

The DATA_RETENTION_TIME_IN_DAYS parameter controls the Time Travel retention period for Snowflake objects such as tables, schemas, and databases. It specifies the number of days that historical data (prior versions of rows or dropped objects) is retained and accessible via Time Travel. Within this retention window, users can query data "as of" a previous time, restore dropped objects, or clone objects at a historical point.

DATE_OUTPUT_FORMAT determines the display format for date values and is unrelated to historical retention. TIMEZONE affects how timestamps are interpreted and displayed, not how long data history is preserved. USE_CACHED_RESULT governs whether Snowflake may return cached query results, not Time Travel behavior.

NEW QUESTION: 22

Changing the size of a virtual warehouse from a Medium to a Large refers to what type of warehouse tuning?

- A. Scaling down
- B. Scaling in
- C. Scaling out
- D. Scaling up

Answer: (SHOW ANSWER)

Changing a Snowflake virtual warehouse from Medium to Large increases the compute resources allocated to that warehouse-such as CPU cores, memory, and I/O bandwidth. This process is called scaling up, or vertical scaling. It is typically performed to improve performance for complex queries, large joins, or computationally heavy workloads.

Scaling down is the reverse-reducing warehouse size from Large to Medium or similar.

Scaling out refers to adding more clusters to a multi-cluster warehouse, enabling increased concurrency rather than increasing power per cluster.

Scaling in is the reverse of scaling out-reducing the number of active clusters.

Thus, increasing the size from Medium to Large is vertical scaling, also known as scaling up.

NEW QUESTION: 23

How are privileges inherited in the Snowflake role hierarchy?

- A. Privileges are only inherited by direct child roles in the hierarchy.
- B. Privileges are inherited by any roles at the same level in the hierarchy.
- C. Privileges are inherited by any roles above a given role in the hierarchy.
- D. Privileges are only inherited by the direct parent role in the hierarchy.

Answer: ([SHOW ANSWER](#))

Snowflake uses a hierarchical Role-Based Access Control (RBAC) model. When Role A is granted to Role B, the higher-level role (Role B) inherits all privileges of the lower-level role (Role A). This upward inheritance allows broad administrative roles (such as SYSADMIN) to automatically inherit permissions granted to subordinate roles.

Key principle: privilege inheritance flows upward, not downward.

* Parent roles inherit from child roles.

* Child roles do not inherit from parent roles.

Sibling roles do not share privileges unless explicitly granted.

Thus, roles above a given role in the hierarchy inherit privileges, making option C correct.

Options A and D incorrectly describe downward inheritance, which Snowflake does not perform.

Option B (same-level inheritance) also does not occur unless privileges are explicitly granted.

This hierarchical model simplifies governance by allowing higher-level administrative roles to centrally manage access across multiple subordinate roles.

NEW QUESTION: 24

What is a key characteristic of the Snowflake architecture's Cloud Services Layer?

- A. It stores all customer data.
- B. It manages virtual warehouses.
- C. It handles security and metadata management.
- D. It provides the user interface for Snowsight.

Answer: ([SHOW ANSWER](#))

The Cloud Services Layer is the coordination and control layer of Snowflake's architecture. One of its primary responsibilities is managing security, metadata, authentication, and system-wide services. This layer handles user authentication, role-based access control, metadata services (such as table structures, micro-partition metadata, statistics), query parsing, optimization, execution coordination, and transaction management.

It does not store customer data; storage is handled by the Database Storage Layer using micro-partitions. It does not manage virtual warehouses directly; warehouses are part of the Compute Layer. While Snowsight is a UI that interacts with the Cloud Services Layer, the interface itself is not part of the architectural layer.

The Cloud Services Layer essentially acts as the "brain" of Snowflake, ensuring the platform is consistent, secure, optimized, and able to scale operations intelligently across compute clusters and cloud-native storage environments.

NEW QUESTION: 25

Where do tables reside in Snowflake?

- A. Stages
- B. Virtual warehouses
- C. Schemas
- D. File formats

Answer: (SHOW ANSWER)

In Snowflake's logical hierarchy, tables reside within schemas, and schemas reside within databases. This structured arrangement allows clear organization, access management, and separation of workloads or business domains. A schema functions as a container that holds tables, views, functions, sequences, stages, and file formats. When users reference a table, they typically qualify it using the full three-part name:

database.schema.table.

Stages are used to store or reference files, not to hold structured database tables. Virtual warehouses provide compute power for executing queries and performing DML operations but do not store data. File formats define how staged files should be interpreted (CSV, JSON, Parquet, etc.) and are unrelated to where tables physically reside.

Snowflake separates compute and storage, so while compute happens in virtual warehouses, all persistent data-including tables-is managed in Snowflake's centralized storage layer, organized through schemas and databases. This organization ensures consistency, ease of management, and predictable access control.

NEW QUESTION: 26

What is the default Time Travel retention period for standard tables in Snowflake?

- A. 0 days
- B. 1 day
- C. 7 days
- D. 14 days

Answer: (SHOW ANSWER)

Standard permanent tables in Snowflake have a default 1-day Time Travel retention period.

During this period, users may query historical data, restore dropped tables, or clone past states.

Higher editions allow extending this period up to 90 days. Temporary and transient tables have 0-day retention by default.

NEW QUESTION: 27

What information can users view in the Query History option in Snowsight?

- A. Only successful queries
- B. Data storage configuration
- C. Network settings
- D. Details of executed queries, including status, duration, and resource usage

Answer: (SHOW ANSWER)

Snowsight's Query History interface provides comprehensive visibility into executed queries. Users can view:

- * Query text
- * Execution status (running, succeeded, failed)
- * Duration and start/end timestamps
- * Warehouse used and credits consumed
- * Partition pruning details
- * Error messages
- * Query profile for performance diagnostics

This allows developers, analysts, and administrators to troubleshoot slow queries, identify bottlenecks, and monitor compute usage.

Incorrect options:

- * Query History shows failed and running queries, not just successful ones.
- * Storage configuration is unrelated to query operations.
- * Network settings are managed by the Cloud Services Layer, not Query History.

Thus, Query History is a key operational tool in Snowsight.

NEW QUESTION: 28

What is the CREATE FILE FORMAT command used for in Snowflake?

- A.** To delete a file format
- B.** To define the format of data files for processing
- C.** To modify an existing file format in a table
- D.** To create a new table

Answer: (SHOW ANSWER)

The CREATE FILE FORMAT command creates a file format object in Snowflake that defines how data files should be interpreted when loading or unloading. Parameters in the definition include file type (for example, CSV, JSON, PARQUET), field delimiters, compression, quoting rules, and other parsing options. Once created, this file format object can be referenced by COPY INTO commands or external table definitions to ensure consistent handling of files.

Deleting a file format is performed using DROP FILE FORMAT. Modifying an existing file format is done with ALTER FILE FORMAT. Creating a new table is achieved with CREATE TABLE.

Therefore, the purpose of CREATE FILE FORMAT is specifically to define how Snowflake should process data files.

NEW QUESTION: 29

What is the primary purpose of Snowflake's Cloning feature?

- A.** To back up data in real-time
- B.** To create an exact copy of a database, schema, or table without duplicating the storage
- C.** To manage user permissions
- D.** To create a backup of the entire Snowflake account

Answer: (SHOW ANSWER)

Snowflake's Cloning feature provides the ability to create a zero-copy clone of tables, schemas, or entire databases. This means that the cloned object references the same underlying micro-partition data as the original, so no duplicate storage costs are incurred upon creation. Instead, Snowflake uses metadata pointers to reference the existing data. If changes occur in the original or cloned objects, Snowflake performs copy-on-write operations to maintain isolation while still optimizing storage. This cloning mechanism is extremely valuable for development, testing environments, analytics validation, and troubleshooting historical states. It is not a real-time backup tool; backup and restore are handled through Time Travel and Fail-safe. Cloning does not grant or manage permissions nor copy the entire Snowflake account; it works solely at the database, schema, and table object levels.

NEW QUESTION: 30

Which of the following are key components of Snowflake's RBAC system? (Choose any 3 options)

- A. Permissions
- B. Roles
- C. Privileges
- D. Users

Answer: (SHOW ANSWER)

Snowflake's RBAC security model revolves around Users, Roles, and Privileges. Users represent individual identities or service accounts. Roles group privileges and are assigned to users. Privileges define specific actions that can be performed on objects (e.g., SELECT, INSERT, OWNERSHIP). "Permissions" is not a formal RBAC component—permissions are effectively the result of privileges assigned to roles. Therefore, the three core RBAC components are Users, Roles, and Privileges.

NEW QUESTION: 31

Which Snowflake object is a logical container for schemas and other database objects?

- A. Schema
- B. View
- C. Table
- D. Database

Answer: (SHOW ANSWER)

A Database is the highest logical container for schemas and their objects in Snowflake.

Hierarchy:

Account # Database # Schema # Tables, Views, Stages, etc.

Schemas organize objects within a database; tables and views do not contain objects.

C01 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com SOL-C01 dumps with Test Engine here:
<https://www.examdiscuss.com/Snowflake/exam/SOL-C01/premium/> (218 Q&As Dumps,
35%OFF Special Discount Code: freecram)

NEW QUESTION: 32

What is the purpose of Time Travel?

- A. To automatically manage timestamp data types
- B. To ensure that users' data can be recovered at any time
- C. To facilitate the loading of historical data into Snowflake
- D. To allow users to access historical data

Answer: (SHOW ANSWER)

Time Travel enables Snowflake users to query, clone, or restore historical versions of data. This includes retrieving previous states of tables, schemas, or databases—even after updates, deletes, or drops. Time Travel operates within a retention period (default: 1 day, up to 90 days on higher editions).

Users can query historical data using the AS OF or BEFORE clause, restore dropped objects, and clone databases at specific points in time for backup or analysis.

Time Travel does not automatically manage timestamp data types. It does not guarantee indefinite recovery—after the retention window expires, data moves into Fail-safe. It also is not primarily designed for loading historical datasets; its purpose is to access past states of Snowflake-managed data.

Thus, the correct purpose is to enable access to historical data inside Snowflake.

NEW QUESTION: 33

What is the purpose of the "auto-suspend" feature in Snowflake virtual warehouses?

- A. To automatically increase the size of the warehouse during peak load
- B. To automatically shut down the warehouse after a period of inactivity
- C. To automatically add more nodes to the warehouse
- D. To automatically reduce the number of clusters

Answer: (SHOW ANSWER)

The auto-suspend feature saves compute costs by automatically suspending a warehouse when it becomes idle for a configured period. Since Snowflake charges for compute time while a warehouse is running, auto-suspend prevents unnecessary credit consumption by stopping the warehouse when no queries are executing.

The warehouse can resume automatically when a new query is submitted, provided auto-resume is enabled.

Auto-suspend does not change warehouse size (scaling up/down) nor adjust clusters in a multi-cluster warehouse (scaling in/out). It strictly controls when compute resources turn off due to inactivity, making it an essential cost-optimization feature.

NEW QUESTION: 34

Which command is used to view the details of a file format object in Snowflake?

- A. LIST FILE FORMAT
- B. ALTER FILE FORMAT
- C. DESCRIBE FILE FORMAT
- D. SHOW FILE FORMATS

Answer: (SHOW ANSWER)

DESCRIBE FILE FORMAT <name> returns configuration settings such as FIELD_DELIMITER, TYPE, SKIP_HEADER, and COMPRESSION for the specified format.

SHOW FILE FORMATS only lists existing formats, not their internals.

LIST FILE FORMAT is invalid. ALTER modifies but does not display details.

NEW QUESTION: 35

What is the LIST command used for in Snowflake?

- A. Lists all the users assigned to a given role.
- B. Lists all the files in a given stage.
- C. Lists all the tables in a specific schema.
- D. Lists all the privileges granted on a specific object.

Answer: (SHOW ANSWER)

TheLISTcommand displays all files available in a stage-internal, external, user, or table stages.

It is typically used prior to COPY INTO to confirm which files are available for loading.

SHOW TABLES lists tables; SHOW GRANTS shows privileges; GRANTS queries show user-role assignments-not LIST.

NEW QUESTION: 36

What is the key difference between a "private listing" and a "public listing" in a Snowflake Data Exchange context?

- A. Public listings require payment, while private listings are free.
- B. Public listings are only accessible to Snowflake employees, while private listings are for external users.
- C. Private listings are shared with specific Snowflake accounts, while public listings are available to all Snowflake users.
- D. Private listings are for sharing code, while public listings are for sharing data.

Answer: (SHOW ANSWER)

Aprivate listingis shared only with explicitly selected Snowflake accounts, giving providers fine-grained control over who can access their data. This is ideal for partner organizations, internal business units, or customers requiring restricted access. Apublic listing, on the other hand, is accessible to all Snowflake users via the Marketplace, making it available for broad distribution. Pricing models for both types can be free or paid; neither category inherently requires payment. Both listing types share data, not code. Public listings are not restricted to Snowflake employees-

they are open to the entire Snowflake ecosystem. The distinction lies entirely in access scope: targeted (private) vs. universal (public).

NEW QUESTION: 37

What is a key benefit of using Streamlit for data visualization in Snowflake notebooks?

- A. Manual data storage
- B. Simplified code execution
- C. Fixed-size compute clusters
- D. Real-time interactivity and dynamic visualizations

Answer: D (LEAVE A REPLY)

Streamlit inside Snowflake Notebooks enables users to build interactive, real-time visualizations and data applications directly within Snowsight. Because Streamlit is event-driven, users can create UI components such as sliders, dropdowns, date pickers, buttons, and filters that dynamically update charts, tables, and computed outputs.

Data remains inside Snowflake, processed through Python with Snowpark, eliminating data movement to external systems. This allows rapid prototyping of dashboards, analytical workflows, and AI-driven apps.

Streamlit also supports layout customization, reactive components, and seamless integration with Snowflake compute for scalable analytics.

Incorrect options:

- * Streamlit does not manage data storage.
- * Streamlit does not simplify code execution; it enhances visualization and interactivity.
- * Warehouse clusters are not fixed-size-Snowflake warehouses can scale.

Thus, its primary benefit is enabling rich interactive analytics directly within Snowflake.

NEW QUESTION: 38

What is a key feature of custom roles in Snowflake? (Choose any 2 options)

- A. They cannot be granted to users
- B. They can be tailored to specific organizational requirements
- C. They have predefined privileges
- D. They are created by Snowflake administrators

Answer: (SHOW ANSWER)

Custom roles in Snowflake allow administrators to design access patterns aligned with organizational security requirements. They are created by Snowflake administrators and do not have predefined privileges; instead, privileges must be granted explicitly. Custom roles can be assigned to users or other roles, just like system-defined roles. Thus, the key attributes are that they are created by administrators and tailored to meet specific organizational needs.

NEW QUESTION: 39

Which command is synonymous with the LIST command in Snowflake?

- A. LS

- B. SHOW
- C. DESC
- D. LIS

Answer: ([SHOW ANSWER](#))

Snowflake allows users to list files in a stage using the LIST command, and it provides a shorthand alias: LS.

Both commands return identical results, including file name, size, last modified timestamp, and other metadata stored in the stage. This is relevant for internal stages (user, table, named stages) and external stages backed by cloud storage. SHOW and DESC are unrelated commands- SHOW lists Snowflake metadata objects such as tables, schemas, or roles, while DESC describes object structures. "LIS" is not a Snowflake command. Therefore, LS is the exact shorthand synonym for LIST and is used identically in queries.

NEW QUESTION: 40

Which command is used to grant privileges to a role in Snowflake?

- A. ASSIGN
- B. CREATE
- C. REVOKE
- D. GRANT

Answer: ([SHOW ANSWER](#))

The GRANT command is used in Snowflake to assign privileges to a role. GRANT allows administrators to authorize roles to perform actions on objects such as tables, warehouses, schemas, or databases. The REVOKE command removes privileges, CREATE creates objects, and ASSIGN is not a valid Snowflake command. Therefore, GRANT is the correct command for privilege assignment.

NEW QUESTION: 41

Which of the following are examples of unstructured data? (Choose any 3 options)

- A. JSON
- B. Image
- C. Video
- D. Pdf

Answer: ([SHOW ANSWER](#))

Unstructured data refers to content that lacks a predefined schema or consistent structural pattern. It cannot be represented naturally in tabular or hierarchical formats without extraction or transformation. Snowflake supports unstructured data through external tables and native features allowing storage and machine learning analyses. Images, videos, and PDFs fall into this category because they contain raw multimedia or document content without intrinsic row/column structure, often requiring OCR, video/audio processing, or content extraction. In contrast, JSON is semi-structured, containing keys, hierarchical structures, and metadata that Snowflake's VARIANT type can interpret directly. Unstructured data is becoming increasingly important for AI-driven

workloads, enabling use cases such as OCR with PARSE_DOCUMENT, deep-learning analysis of images and video, and document intelligence tasks.

NEW QUESTION: 42

What does Snowflake UI called?

- A. Snowsite
- B. SnowGUI
- C. SnowView
- D. Snowsight

Answer: ([SHOW ANSWER](#))

The official Snowflake web interface is called Snowsight. Snowsight is the modern, browser-based UI that enables SQL development, Python worksheets, dashboarding, visual exploration, and account monitoring. It replaces the classic UI and is now the default interface. Snowsight provides integrated capabilities such as notebooks, data loading interfaces, dashboards, role and warehouse management, and visual query history.

Other options like Snowsite, SnowGUI, and SnowView are not valid Snowflake product names.

NEW QUESTION: 43

What does "warehouse scaling in/out" refer to in Snowflake?

- A. Changing the size of the warehouse (e.g., from Small to Medium or vice versa).
- B. Moving data between different storage locations.
- C. Changing the region of the warehouse.
- D. Adjusting the number of clusters in a multi-cluster warehouse.

Answer: ([SHOW ANSWER](#))

Scaling in/out in Snowflake refers to modifying the number of compute clusters associated with a multi-cluster virtual warehouse. Scaling out increases cluster count to accommodate higher concurrency or workload spikes, allowing more queries to run simultaneously without queuing. Scaling in reduces cluster count during periods of lower demand, optimizing compute usage and costs. This is distinct from scaling up/down, which refers to changing warehouse size (e.g., Small, Medium). Scaling does not involve data movement or region changes; warehouse compute is stateless and operates independently of storage. Multi-cluster warehouses allow Snowflake to automatically add or remove clusters based on demand when auto-scale policies are configured.

NEW QUESTION: 44

Which of the following are types of tables in Snowflake? (Choose any 3 options)

- A. Permanent
- B. Temporary
- C. Transient
- D. Virtual

Answer: ([SHOW ANSWER](#))

Snowflake provides three main table types: Permanent, Transient, and Temporary. Permanent tables support full Time Travel and Fail-safe, and are designed for long-term storage. Transient tables do not include Fail-safe and are best suited for intermediate or less durable datasets. Temporary tables exist only for the duration of the user session that created them, making them ideal for staging or ephemeral workloads. Snowflake does not have a table type called "Virtual"; therefore, the correct options are Permanent, Temporary, and Transient.

NEW QUESTION: 45

When loading data into a table using the COPY INTO command, what is the default error-handling setting?

- A. ABORT_STATEMENT
- B. SKIP_FILE
- C. VALIDATE
- D. CONTINUE

Answer: (SHOW ANSWER)

By default, Snowflake uses `ON_ERROR = ABORT_STATEMENT` for COPY INTO. This means that if a single row fails to load due to issues such as type mismatches or malformed data, Snowflake immediately stops the entire load operation. This behavior ensures data integrity by preventing partial or incomplete loads when errors occur.

Other error-handling modes must be explicitly specified: CONTINUE skips erroneous rows, SKIP_FILE skips whole files with errors, and VALIDATE is part of the validation mode feature, not a default COPY behavior.

Thus, unless explicitly changed, COPY INTO always uses ABORT_STATEMENT.

NEW QUESTION: 46

How can you search for listings on Snowflake Marketplace? (Choose any 2 options)

- A. Browsing categories and tags
- B. Contact Snowflake Support
- C. User Universal Search
- D. Using the search bar at the top of the marketplace

Answer: (SHOW ANSWER)

Snowflake Marketplace provides multiple intuitive discovery mechanisms to help users explore available listings efficiently. The search bar at the top of the Marketplace is a primary entry point, enabling keyword-based discovery across providers, industries, and dataset types. This facilitates targeted searching for specific topics or data categories.

Universal Search, available in Snowsight, extends search beyond databases and worksheets to include Marketplace listings. This unified search experience allows users to locate datasets without navigating through menus, improving productivity and discoverability.

Browsing categories and tags supports exploratory search by grouping listings under predefined themes such as financial data, geospatial data, demographics, and more. These tags promote

structured navigation, especially useful when users are exploring rather than seeking something specific.

Contacting Snowflake Support is not required for accessing or discovering listings and is irrelevant to Marketplace search workflows.

Valid SOL-C01 Dumps shared by ExamDiscuss.com for Helping Passing SOL-C01 Exam! ExamDiscuss.com now offer the **newest SOL-C01 exam dumps**, the ExamDiscuss.com SOL-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com SOL-C01 dumps with Test Engine here:
<https://www.examdiscuss.com/Snowflake/exam/SOL-C01/premium/> (218 Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 47

What is the primary purpose of the COPY INTO command in Snowflake for data loading?

- A. To replicate databases across regions.
- B. To create new virtual warehouses.
- C. To load data from staged files into a table.
- D. To manage user roles and permissions.

Answer: (SHOW ANSWER)

COPY INTO <table> bulk-loads data from internal or external stages into Snowflake tables. It supports transformations, validation, and multiple file formats.

It does not replicate databases, manage roles, or create warehouses.

NEW QUESTION: 48

What filter can be used when searching for data products on the Snowflake Marketplace?

- A. Size
- B. Compatibility
- C. Price
- D. Performance

Answer: C (LEAVE A REPLY)

On the Snowflake Marketplace, users can apply a Price filter when searching for data products. This filter helps distinguish between free offerings and paid listings, making it easier to align data acquisition with budget and commercial terms. Marketplace metadata includes pricing information that this filter uses to refine search results.

Dataset "Size" is not exposed as a primary Marketplace search filter. "Compatibility" is implied, as all Marketplace products are consumable within Snowflake. "Performance" is not a filterable attribute, because performance is influenced mainly by the consumer's compute configuration and query design, rather than a fixed property of the Marketplace listing.

NEW QUESTION: 49

Which SQL command is used to view the definition of a Table?

- A. SHOW TABLE
- B. SELECT GET_DDL('TABLE','<table_name>');
- C. SELECT * FROM INFORMATION_SCHEMA.ABLES
- D. DESCRIBE TABLE

Answer: B (LEAVE A REPLY)

The GET_DDL function retrieves the Data Definition Language (DDL) script that was used to create a given object. Executing:

```
SELECT GET_DDL('TABLE', 'MY_TABLE');
```

returns the full DDL including column definitions, comments, clustering keys, constraints, masking policies, and other metadata. This makes it essential for schema migration, replication, compliance audits, and environment synchronization.

DESCRIBE TABLE shows only column-level metadata-names, types, nullability-but does not return the full DDL.

SHOW TABLE is not a valid Snowflake command; tables are listed using SHOW TABLES.

INFORMATION_SCHEMA.TABLES lists table metadata, but the provided option "ABLES" is invalid and still would not provide DDL even if corrected.

Thus, GET_DDL() is the only method that returns the complete table definition directly.

NEW QUESTION: 50

Which command can be used to query the Information Schema in Snowflake to retrieve a list of tables?

- A. SELECT * FROM PUBLIC.INFORMATION_SCHEMA;
- B. SELECT * FROM INFORMATION_SCHEMA.COLUMNS;
- C. SELECT * FROM INFORMATION_SCHEMA.TABLES;
- D. SELECT * FROM PUBLIC.TABLES;

Answer: (SHOW ANSWER)

Snowflake's INFORMATION_SCHEMA contains metadata views about databases, schemas, tables, columns, constraints, and other objects. To retrieve a list of tables, the correct view is INFORMATION_SCHEMA.

TABLES, which returns:

- * Table name
- * Schema name
- * Table type (BASE TABLE, VIEW, MATERIALIZED VIEW)
- * Creation timestamp
- * Owner and retention properties

The query:

```
SELECT * FROM INFORMATION_SCHEMA.TABLES;
```

returns metadata for all tables in the current database context.

INFORMATION_SCHEMA.COLUMNS returns column-level metadata, not table listings.

PUBLIC.INFORMATION_SCHEMA is invalid-INFORMATION_SCHEMA is not inside PUBLIC schema.

PUBLIC.TABLES is not a system object unless created manually.

NEW QUESTION: 51

Which of the following file formats can be processed by the PARSE_DOCUMENT function?

- A. PDF
- B. PPTX
- C. XLSX
- D. DOCX

Answer: ([SHOW ANSWER](#))

PARSE_DOCUMENT extracts text and, depending on mode, layout information from supported document formats. The function currently supports PDF, PPTX, and DOCX files. For PDFs, it can perform OCR to extract scanned text or use LAYOUT mode to detect tables and structured regions. For PPTX files, it extracts slide text and content. For DOCX files, the function captures text and applies layout analysis. XLSX files are not supported-Excel content requires different extraction methods. PARSE_DOCUMENT enables automated document intelligence workflows entirely inside Snowflake, eliminating the need for external OCR or parsing services.

NEW QUESTION: 52

What is the typical syntax for using the LIMIT clause to retrieve the first 10 rows of a table named employees?

- A. SELECT * FROM employees LIMIT 10;
- B. SELECT LIMIT 10 * FROM employees;
- C. LIMIT 10 SELECT * FROM employees;
- D. SELECT * FROM employees WHERE LIMIT 10;

Answer: ([SHOW ANSWER](#))

The LIMIT clause is used to restrict the number of rows returned by a query. The correct syntax in Snowflake is:

```
SELECT * FROM employees LIMIT 10;
```

This retrieves the first 10 rows of the result set. LIMIT appears at the end of the SELECT statement.

Snowflake follows a SQL syntax similar to MySQL/PostgreSQL regarding LIMIT usage.

Incorrect forms include:

- * Using LIMIT before SELECT
- * Embedding LIMIT inside WHERE
- * Reversing the order of SELECT and LIMIT

LIMIT is frequently used for sampling, record previewing, debugging, and exploratory analysis.

NEW QUESTION: 53

What are the key benefits of the Snowflake multi-cluster shared data architecture? (Select TWO).

- A. It provides enhanced data security features.
- B. It optimizes data loading for unstructured data.
- C. It allows for independent scaling of compute and storage.
- D. It stores data in a columnar format to improve performance.
- E. It offers near-unlimited concurrency and elasticity.

Answer: (SHOW ANSWER)

The Snowflake multi-cluster shared data architecture separates compute and storage, enabling independent scaling of each. Compute resources (virtual warehouses) operate independently from the centralized storage layer, allowing users to increase compute power without impacting storage costs-or vice versa. This flexibility provides major advantages for cost optimization and workload performance tuning.

The architecture also supports near-unlimited concurrency and elasticity. When many users or workloads run simultaneously, Snowflake can automatically add additional clusters to a multi-cluster warehouse, ensuring that no queries experience queuing or performance degradation. This capability is crucial for BI dashboards, ETL pipelines, and large organizations with varied usage patterns.

Columnar storage (option D) is indeed a Snowflake feature but is not unique to the multi-cluster architecture.

Enhanced security arrives via Snowflake's platform-wide mechanisms, not specifically because of multi-cluster architecture. Unstructured data loading (option B) is also not governed by multi-cluster behavior.

NEW QUESTION: 54

Which of the following are valid components of Snowflake's core architecture? (Select THREE)

- A. On-premises Data Centers
- B. User Interface Layer
- C. Third-Party ETL Tools
- D. Database Storage Layer
- E. Query Processing Layer
- F. Cloud Services Layer

Answer: (SHOW ANSWER)

Snowflake's architecture consists of:

- * Database Storage Layer - Stores optimized columnar data.
- * Query Processing Layer - Executes workloads via virtual warehouses.
- * Cloud Services Layer - Handles metadata, optimization, access control.

ETL tools and UI are external integrations. Snowflake is cloud-native and has no on-premises components.

NEW QUESTION: 55

Which of the following can be cloned using Snowflake's Cloning feature? (Choose any 3 options)

- A. Databases

- B. Schemas
- C. Roles
- D. Tables

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

Snowflake supports zero-copy cloning at the database, schema, and table levels. Cloning creates a new metadata reference to existing data without physically duplicating micro-partitions, significantly reducing storage cost and clone creation time. Cloning a database replicates all schemas and tables within it at that moment in time.

Cloning a schema replicates all objects inside that schema, including tables and views. Cloning a table creates an isolated version of that table for development, testing, or analysis.

Importantly, roles cannot be cloned because they are account-level security objects, not data objects. Privileges must be manually granted or recreated for cloned objects. Cloning enables rapid, efficient creation of consistent environments for workload isolation and testing.

NEW QUESTION: 56

How is the initial activation of the Snowflake Data Exchange feature, which is necessary for establishing private exchanges, typically accomplished?

- A. Any Snowflake account administrator can enable Data Exchange with a click.
- B. Data Exchange is automatically enabled for all accounts.
- C. Activation requires a SQL command executed by ACCOUNTADMIN.
- D. Activation requires a request to Snowflake Support.

Answer: ([SHOW ANSWER](#))

Snowflake Data Exchange is not enabled by default. Accounts must open a support request to Snowflake Support, which then activates the feature. It cannot be enabled via UI or SQL. After activation, administrators may create private exchanges and listings.

NEW QUESTION: 57

Which SQL keyword is used to constrain the number of rows returned by a query in Snowflake?
(Select TWO)

- A. TOP
- B. SELECT
- C. GET
- D. LIMIT
- E. ROW_NUMBER

Answer: ([SHOW ANSWER](#))

Snowflake supports both TOP and LIMIT as valid SQL mechanisms for restricting the number of rows returned by a query. They function similarly but are used in different syntactic positions.

LIMIT appears at the end of the query:

```
SELECT * FROM my_table LIMIT 10;
```

This tells Snowflake to return only the specified number of rows.

TOP is used in the SELECT clause:

```
SELECT TOP 10 * FROM my_table;
```

Both approaches are equivalent in functionality and are supported to maintain compatibility with ANSI SQL and other SQL dialects.

The other options are not correct row-limiting mechanisms:

*SELECT is a required keyword but does not limit rows on its own.

*GET is not a SQL keyword for row restriction; GET is used for downloading files from internal stages.

*ROW_NUMBER is a window function used for ranking rows, not constraining row count.

Therefore, TOP and LIMIT are the only correct SQL keywords in Snowflake for controlling how many rows are returned by a query.

NEW QUESTION: 58

When should AUTO_SUSPEND be enabled for a virtual warehouse?

- A. To ensure the warehouse is always active for immediate queries.
- B. To save credits during periods of inactivity.
- C. When the warehouse is expected to run continuously.
- D. To increase the warehouse's performance.

Answer: (SHOW ANSWER)

AUTO_SUSPEND minimizes cost by suspending compute after a specified idle period.

It prevents unnecessary credit consumption while ensuring compute resumes automatically when needed.

Performance improvements depend on warehouse size, not suspension settings.

NEW QUESTION: 59

What are the steps to create a new stage in Snowsight?

- A. Data >> Create Schema
- B. Data >> add data >> Create Schema
- C. Data Product >> Create Schema
- D. Data >> Database >> Schema >> Create Stage

Answer: D (LEAVE A REPLY)

To create a new stage in Snowsight, users navigate through: Data # Database # Schema # Create Stage. Within the schema, selecting the "+" icon opens the option to create a stage. Users may configure internal or external stages, encryption settings, and directory table options. Other paths listed reference creating schemas rather than creating stages and are therefore incorrect.

NEW QUESTION: 60

Which of the following are examples of system-defined roles in Snowflake? (Choose any 3 options)

- A. USERADMIN
- B. CUSTOM_ROLE
- C. ACCOUNTADMIN
- D. SECURITYADMIN

Answer: (SHOW ANSWER)

Snowflake includes several system-defined RBAC roles to provide structured access control across the platform.

ACCOUNTADMIN is the highest-privileged role, governing the entire account, including billing, replication, failover groups, and object ownership inheritance.

SECURITYADMIN is responsible for user, role, and privilege management, including granting and revoking roles, password policies, and MFA enforcement. It is essential for identity lifecycle governance.

USERADMIN manages user creation, role creation, and user authentication administration. It is commonly assigned to identity administrators.

CUSTOM_ROLE is not a Snowflake system role; it is a user-created role, allowing organizations to design fine-grained, domain-specific access policies.

These predefined roles form the foundation of Snowflake's security model and support least-privilege implementations.

NEW QUESTION: 61

How can roles access objects in Snowflake? (Choose any 2 options)

- A. By being directly assigned privileges on the objects
- B. By inheriting privileges from parent roles
- C. By configuring network settings
- D. By storing raw data

Answer: A,B (LEAVE A REPLY)

Roles gain access to objects in Snowflake through direct privilege grants or through inheritance within a role hierarchy. A role may be explicitly granted privileges such as SELECT, OWNERSHIP, or USAGE.

Alternatively, if a role is assigned to another role, it inherits all privileges of its parent role. Neither network settings nor data storage processes determine role access.

Valid SOL-C01 Dumps shared by ExamDiscuss.com for Helping Passing SOL-C01 Exam! ExamDiscuss.com now offer the **newest SOL-C01 exam dumps**, the ExamDiscuss.com SOL-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com SOL-C01 dumps with Test Engine here:

<https://www.examdiscuss.com/Snowflake/exam/SOL-C01/premium/> (218 Q&As Dumps,

35%OFF Special Discount Code: freecram)

NEW QUESTION: 62

Which types of stages are available in Snowflake? (Choose any 2 options)

- A. Internal
- B. External

- C. Public
- D. Private

Answer: (SHOW ANSWER)

Snowflake provides two main types of stages: Internal stages and External stages. Internal stages reside inside Snowflake and include user stages, table stages, and named internal stages. These offer convenient, secure storage within Snowflake's managed environment. External stages reference cloud storage locations such as AWS S3, Azure Blob Storage, or Google Cloud Storage, enabling integration with external data repositories.

Snowflake does not define a stage type called "Public" or "Private," although stages can be granted secure access. Thus, the correct types are Internal and External stages.

NEW QUESTION: 63

What is a primary function of a view in Snowflake?

- A. Providing a virtual table based on the result set of a query
- B. Executing network configurations
- C. Managing user authentication
- D. Storing raw data

Answer: (SHOW ANSWER)

A view in Snowflake represents a virtual table whose contents are defined by a stored SQL query. Views allow users to encapsulate transformation logic, simplify complex joins, enforce column-level security, and provide curated datasets for downstream consumers.

Views do not store data themselves; the underlying tables store the data, and Snowflake executes the view's query each time the view is referenced. This makes views ideal for abstraction layers, semantic modeling, and separating compute costs across user groups.

Incorrect options:

- * Snowflake views do not store raw data.
- * Views have no role in authentication or network configuration.

Views help streamline business logic and control data access efficiently.

NEW QUESTION: 64

What Snowflake object is used to organize database objects into logical groupings?

- A. Schemas
- B. Roles
- C. Tables
- D. Stages

Answer: (SHOW ANSWER)

Schemas serve as logical containers within a database, grouping tables, views, file formats, functions, and other objects.

Roles control access, not organization.

Tables store data but do not group objects.

Stages store files or point to external storage; they are not organizational containers.

NEW QUESTION: 65

What can a schema in Snowflake contain? (Choose any 3 options)

- A. Views
- B. Tables
- C. Stage
- D. Virtual warehouses

Answer: (SHOW ANSWER)

A schema in Snowflake is a logical container inside a database that groups related objects. It can contain tables, views, stages, file formats, sequences, streams, tasks, and other database objects. Schemas allow structured organization, privilege management, and separation of workloads.

* Tables store structured data.

* Views are virtual tables defined by queries on other objects.

* Stages store data files for loading/unloading.

Virtual warehouses are compute resources, not database objects, and therefore exist outside schemas.

Warehouses belong to the account-level namespace and are managed independently of databases and schemas.

Schemas play a key role in RBAC and object naming through fully qualified names: `database.schema.object`.

NEW QUESTION: 66

What is the primary benefit of the Snowflake data cloud?

- A. It eliminates the need for data governance.
- B. It provides direct access to underlying infrastructure.
- C. It replaces traditional data warehouses with on-premises solutions.
- D. It enables organizations to unite and share their data.

Answer: (SHOW ANSWER)

The Snowflake Data Cloud allows organizations to seamlessly share, access, and collaborate on data across departments and external partners, without copying or moving data. Through secure data sharing, listings, and data clean rooms, Snowflake eliminates data silos and dramatically improves data collaboration.

It does not eliminate the need for governance—Snowflake enhances governance via RBAC, masking policies, and centralized controls. It does not provide access to underlying cloud infrastructure; Snowflake abstracts that. It is not an on-premises solution; Snowflake is fully cloud-native.

Thus, the primary benefit is unifying and securely sharing data across the ecosystem.

NEW QUESTION: 67

What is a benefit of using an external stage to load data into Snowflake?

- A. External stages reduce data storage costs because data is stored outside Snowflake.

- B. External stages provide automatic data purging after successful loads.
- C. External stages are more secure than internal stages for sensitive data.
- D. External stages reduce the number of objects in a database.

Answer: (SHOW ANSWER)

External stages point to files in external cloud storage (S3, Azure Blob, GCS). Because the data is not stored inside Snowflake, the user avoids Snowflake storage charges, which can significantly reduce cost for large staging datasets.

External stages do not automatically delete files, are not inherently more secure than internal stages, and still count as database objects even though they reference external storage.

NEW QUESTION: 68

What does a red checkmark next to a cell in a Snowflake notebook signify?

- A. The cell has errors
- B. The cell is currently running
- C. The cell has completed execution successfully
- D. The cell is paused

Answer: (SHOW ANSWER)

A red checkmark in Snowflake Notebooks indicates that the cell encountered an error during execution. This visual indicator helps users quickly identify issues so they can inspect the error message and adjust code accordingly. Successful execution shows a green checkmark, while running cells show a spinner. Snowflake notebooks do not include a "paused" state for notebook cells.

NEW QUESTION: 69

How can Python variable substitution be performed in Snowflake notebooks?

- A. By manually editing data files
- B. By configuring network settings
- C. By using placeholders in SQL queries
- D. By writing HTML code

Answer: (SHOW ANSWER)

In Snowflake Notebooks, Python and SQL operate in an integrated environment. Python variable substitution is performed using string placeholders or f-string interpolation within SQL statements. This allows dynamic query construction where Python values are embedded directly into SQL code executed through the Snowpark session.

Example:

```
table_name = "CUSTOMERS"  
session.sql(f"SELECT * FROM {table_name}").show()
```

This method enables parameterized, context-aware SQL generation-ideal for iterative development, automation, and pipeline construction. It avoids manual text editing and provides consistency across Python and SQL execution layers.

Incorrect responses:

- * Network settings and HTML are unrelated to variable handling.
- * Editing data files does not apply to dynamic query parameterization.

Variable substitution is essential for combining logic, iteration, and conditional flows within notebooks.

NEW QUESTION: 70

How is the initial activation of the Snowflake Data Exchange feature, which is necessary for establishing private exchanges, typically accomplished?

- A.** Any Snowflake account administrator can enable Data Exchange through the web interface with a single click.
- B.** Activation of the Data Exchange feature generally requires a request to Snowflake Support, who then enable it for the specific account.
- C.** Data Exchange is automatically enabled for all Snowflake accounts upon creation.
- D.** Activation of the Data Exchange feature requires a specific SQL command to be executed by the ACCOUNTADMIN role.

Answer: B (LEAVE A REPLY)

The Data Exchange feature is not automatically enabled for all Snowflake accounts. Instead, enabling this capability-particularly for private exchanges-typically requires submitting a request to Snowflake Support.

Snowflake evaluates account type, region, and compliance requirements before enabling the Data Exchange functionality. Snowflake administrators cannot independently enable this through Snowsight, nor is there a SQL command available to turn it on. Once activated, account administrators can create private exchanges, manage listings, and define which Snowflake accounts can participate. This activation step ensures secure, intentional use of data-sharing capabilities at an account level.

NEW QUESTION: 71

Which of the following are numeric data types supported by Snowflake? (Choose any 2 options)

- A.** INTEGER
- B.** BOOLEAN
- C.** VARCHAR
- D.** FLOAT

Answer: (SHOW ANSWER)

Snowflake supports several numeric data types, including INTEGER (and its aliases INT, BIGINT, SMALLINT, TINYINT) and FLOAT (a floating-point type). INTEGER represents exact whole numbers, and FLOAT supports approximate numerical values with fractional components. BOOLEAN is a logical data type used for TRUE/FALSE values, not numeric. VARCHAR is a string data type, also not numeric. Thus, INTEGER and FLOAT are the correct numeric data types supported by Snowflake.

NEW QUESTION: 72

What are the benefits of using the Snowsight data loading interface? (Select TWO).

- A. It creates permanent file formats that can be used to load data in the future.
- B. It allows a user to insert the records of a supported file into a table.
- C. It will try to detect data types.
- D. It allows a user to optimize data loading into a table.
- E. It lets a user merge file rows into the table records.

Answer: (SHOW ANSWER)

Snowsight's file-loading interface:

*Automatically detects data types by inspecting column values.

*Loads (inserts) file contents directly into a table, either a new table or an existing one.

Snowsight does not automatically create permanent file formats, does not merge rows, and does not optimize warehouse-level performance.

NEW QUESTION: 73

Which edition of Snowflake offers the most comprehensive security and compliance features, including HIPAA and PCI DSS compliance?

- A. Standard Edition
- B. Business Critical Edition
- C. Enterprise Edition
- D. Virtual Private Snowflake (VPS) Edition

Answer: (SHOW ANSWER)

Business Critical Edition includes advanced compliance certifications such as HIPAA, PCI DSS, HITRUST, and additional encryption and security controls.

Standard and Enterprise editions lack these compliance guarantees.

VPS provides even stronger isolation but builds on Business Critical features.

NEW QUESTION: 74

Which SQL command is used to create a new SCHEMA in Snowflake? (Choose any 3 options)

- A. CREATE SCHEMA Schema_name;
- B. USE SCH Schema_name;
- C. USE SCHEMA Schema_name;
- D. CREATE SCH Schema_name;

Answer: (SHOW ANSWER)

The only valid Snowflake SQL command among the options for creating a schema is:

CREATE SCHEMA schema_name;

This command creates a schema within the currently active database. Optional modifiers like IF NOT EXISTS or OR REPLACE can be used to control behavior.

Other options:

- * USE SCHEMA switches to an existing schema; it does not create one.
- * USE SCH is invalid because SCH is not a Snowflake keyword.
- * CREATE SCH is invalid; SCH is not a valid abbreviation.

Thus, the only valid schema creation command is CREATE SCHEMA.

NEW QUESTION: 75

Which option can be used to specify the delimiter for a CSV file in the COPY INTO <TABLE> command?

- A. DELIMITER
- B. SEPARATOR
- C. FIELD_DELIMITER
- D. COLUMN_DELIMITER

Answer: (SHOW ANSWER)

The correct Snowflake parameter for defining a CSV delimiter during data loading is FIELD_DELIMITER.

This option can be set within a named file format or specified inline in a COPY INTO command, such as FIELD_DELIMITER = ',' for comma-separated files or FIELD_DELIMITER = '\t' for tab-delimited data.

Snowflake does not recognize DELIMITER, SEPARATOR, or COLUMN_DELIMITER as valid parameters.

FIELD_DELIMITER ensures that Snowflake accurately splits each row into separate fields based on the chosen delimiter. Without the right delimiter definition, the COPY INTO operation would treat entire rows as a single string, causing misalignment with the table structure.

FIELD_DELIMITER is essential when handling CSV or TSV files, particularly when ingesting data from external systems, cloud storage, or partner feeds. Snowflake's flexibility with file formatting—combined with parallel loading and automatic error-handling—makes COPY INTO an optimized solution for structured data ingestion.

NEW QUESTION: 76

What does the SELECT * statement do in a Snowflake query?

- A. Retrieves all columns from a specified table or view
- B. Retrieves only distinct values
- C. Retrieves only the primary key column
- D. Retrieves a limited number of rows

Answer: A (LEAVE A REPLY)

The SELECT * statement instructs Snowflake to return all columns from the referenced table or view. This is commonly used during data exploration, debugging, initial data profiling, and validation steps. It allows users to quickly view the complete dataset structure without manually specifying each column name.

However, while SELECT * retrieves all columns, it does not limit the number of rows. To restrict rows, developers must include a LIMIT clause (e.g., SELECT * FROM table LIMIT 10;).

The query does not automatically apply DISTINCT or primary key filtering—Snowflake returns all rows exactly as stored unless additional filtering, WHERE conditions, or ordering are provided.

Though SELECT * is convenient, Snowflake best practices recommend explicitly selecting columns in production workloads to optimize performance and avoid unnecessary scanning of unused fields.

Valid SOL-C01 Dumps shared by ExamDiscuss.com for Helping Passing SOL-C01 Exam! ExamDiscuss.com now offer the **newest SOL-C01 exam dumps**, the ExamDiscuss.com SOL-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com SOL-C01 dumps with Test Engine here: <https://www.examdiscuss.com/Snowflake/exam/SOL-C01/premium/> (218 Q&As Dumps, **35%OFF Special Discount Code: freecram**)

NEW QUESTION: 77

Where is unstructured data stored in Snowflake?

- A. In the Cloud Services layer
- B. In internal or external stages
- C. In external tables
- D. In tables with a single VARCHAR column

Answer: (SHOW ANSWER)

Unstructured data such as PDF files, images, and other binary documents is stored in stages in Snowflake.

These stages may be internal stages, which Snowflake manages directly, or external stages, which reference external cloud storage such as Amazon S3, Azure Blob Storage, or Google Cloud Storage. Stages are the designed mechanism for storing and accessing unstructured files so that they can be processed with functions like PARSE_DOCUMENT or accessed via directory tables.

External tables are used to query structured or semi-structured data (for example, Parquet or JSON) stored in external locations, not to store raw unstructured binary content. The Cloud Services layer coordinates metadata, security, and query services; it does not store user data. Tables with a single VARCHAR column might be used as an improvised approach for small text blobs, but this is not the native or recommended method for managing unstructured data at scale.

NEW QUESTION: 78

What are characteristics of the Snowflake Platform? (Select TWO).

- A. Snowflake supports column-level security not row-level security.
- B. Snowflake is responsible for all data security.
- C. Snowflake handles platform maintenance, management, and upgrades.
- D. There is no infrastructure to provision and maintain.
- E. Snowflake can run on private infrastructures on-premises.

Answer: (SHOW ANSWER)

Snowflake is a fully managed, cloud-native platform that automates nearly all infrastructure tasks. Two defining characteristics are:

1. Snowflake handles platform maintenance, management, and upgrades.

Snowflake automatically performs software updates, scaling operations, tuning, optimization, and patching.

Users never need to upgrade hardware, install software, or manage clusters.

2. There is no infrastructure for customers to provision and maintain.

All compute, storage, and services run on public clouds (AWS, Azure, GCP). Snowflake abstracts the underlying infrastructure completely, eliminating the need for customers to manage VMs, disks, clusters, or networking.

Incorrect options:

* Snowflake supports both column-level security (via masking policies) and row-level security (via row access policies).

* Snowflake is not responsible for all data security—customers must configure RBAC, masking, and governance.

* Snowflake does not run on-premises; it is exclusively cloud-based.

Thus, options C and D correctly describe Snowflake platform characteristics.

NEW QUESTION: 79

Query History can display data up to how many days?

A. 28

B. 14

C. 21

D. 7

Answer: (SHOW ANSWER)

In Snowsight, the Query History user interface retains and displays executed query information for up to 14 days. This includes query text, status, start/end timestamps, duration, warehouse usage, and result statistics.

Snowsight also offers grouping by session, user, warehouse, and time windows for detailed analytics and troubleshooting.

It is important to distinguish Snowsight UI limits from SQL-based history functions. The table function `INFORMATION_SCHEMA.QUERY_HISTORY()` returns only 7 days of history, whereas Snowsight visually retains 14 days. Snowflake's Account Usage views (e.g., `QUERY_HISTORY`) retain data for longer periods but are meant for programmatic auditing rather than UI display.

Incorrect options:

* 7 days refers only to Information Schema.

* 21 or 28 days are not supported by Snowsight's UI.

Thus, the correct visibility window in the Snowsight interface is 14 days.

NEW QUESTION: 80

What is the name of Snowflake's default web-based interface?

- A. SnowSQL
- B. Snowpark
- C. Snowsight
- D. Snowflake CLI

Answer: (SHOW ANSWER)

Snowsight is the default web-based interface for Snowflake. It provides a graphical, browser-based environment to write and run SQL and Python, manage and explore database objects, create visualizations and dashboards, monitor query and warehouse activity, and collaborate via worksheets. It has replaced the Classic Console as the default UI for new Snowflake accounts. SnowSQL is a command-line client used to connect to Snowflake from a terminal and execute SQL statements. Snowpark is a developer framework that supports writing data applications in languages like Python, Java, and Scala, but it is not a UI. "Snowflake CLI" refers to command-line tooling and automation utilities, again not a graphical web interface.

NEW QUESTION: 81

What Snowflake object provides a secure connection to external cloud storage?

- A. A directory table
- B. An external table
- C. An external stage
- D. A named file format

Answer: (SHOW ANSWER)

An external stage is the Snowflake object that encapsulates a secure connection to external cloud storage such as Amazon S3, Azure Blob Storage, or Google Cloud Storage. It stores the location (URL or bucket path) and, where required, credentials or role-based access configuration, and may also reference a file format. External stages are used as the source or target for COPY INTO operations when loading from or unloading to external storage.

An external table provides a logical SQL interface to data stored externally but relies on a stage for connectivity; it does not itself define the connection. A directory table exposes metadata about files stored in a stage, not the connection. A named file format defines parsing rules (type, delimiter, compression) but has no knowledge of or connection to a specific external storage location.

NEW QUESTION: 82

What happens when you execute the following command in Snowflake?

```
DROP SCHEMA IF EXISTS marketing_schema;
```

- A. An error is returned if the schema does not exist
- B. The schema "marketing_schema" is created if it does not exist
- C. The schema "marketing_schema" is dropped only if it exists
- D. The schema "marketing_schema" is dropped even if it contains tables

Answer: (SHOW ANSWER)

The SQL command `DROP SCHEMA IF EXISTS marketing_schema;` instructs Snowflake to drop the specified schema only if it already exists in the database. The inclusion of the `IF EXISTS` clause prevents Snowflake from returning an error when the schema is not found. If the schema does exist, Snowflake drops it along with all objects it contains, unless a `RESTRICT` clause is enforced (`RESTRICT` is the default when not explicitly specified, but Snowflake also supports `CASCADE` for forcing deletion of all objects). Importantly, this command never creates a schema; it strictly performs a conditional drop. In addition, Snowflake supports Time Travel, so the dropped schema can typically be restored using `UNDROP SCHEMA` within the defined retention window. Therefore, the correct interpretation is that the schema is dropped only if it exists, and no error is thrown otherwise.

NEW QUESTION: 83

What is the primary benefit of the separation of storage and compute in Snowflake?

- A. It allows independent scaling of resources.
- B. It reduces network latency.
- C. It simplifies data governance.
- D. It eliminates the need for data loading.

Answer: ([SHOW ANSWER](#))

Snowflake's architecture separates storage and compute, enabling:

- * Compute scaling (up/down or multi-cluster) without changing storage
- * Storage expansion without affecting compute
- * Cost optimization by paying for compute only when needed

This separation does not impact governance, latency, or data loading requirements.

NEW QUESTION: 84

What options are available under the settings for Python worksheets? (Choose any 2 options)

- A. Handler
- B. Return Value
- C. Return Type
- D. Output

Answer: ([SHOW ANSWER](#))

In Snowflake Python worksheets, settings include the `Handler` and `Return Type`. The `Handler` identifies the specific Python function that serves as the entry point for execution-essential when deploying code as stored procedures or user-defined functions. The `Return Type` defines how Snowflake interprets the output, such as `STRING`, `VARIANT`, or `TABLE`, ensuring proper result rendering. "Return Value" is not a worksheet setting but rather the outcome produced by the code itself, and "Output" is not a configuration option.

NEW QUESTION: 85

What is the purpose of the `COMPLETE` function in Snowflake Cortex?

- A. To parse documents

- B. To translate languages
- C. To classify text
- D. To generate text completions based on a given prompt

Answer: ([SHOW ANSWER](#))

The COMPLETE function enables natural language generation within Snowflake Cortex. It accepts a prompt and produces a text-based completion using an underlying large language model (LLM). COMPLETE can generate single-turn or multi-turn responses, depending on how the prompt is structured. It is used for tasks such as drafting summaries, generating recommendations, writing structured output such as JSON, or extending existing text. It is not a classification tool-that role is served by CLASSIFY_TEXT-and it does not perform linguistic translation (TRANSLATE) or document extraction (PARSE_DOCUMENT). COMPLETE allows developers and analysts to embed AI-driven generative capabilities directly within SQL pipelines or Snowpark processes, eliminating the need for external LLM services.

NEW QUESTION: 86

What are Snowflake customers responsible for?

- A. Managing the hardware on which Snowflake is installed
- B. Software installation and updates
- C. Managing table metadata
- D. Loading, unloading, and managing data

Answer: ([SHOW ANSWER](#))

As a fully managed cloud data platform, Snowflake is responsible for infrastructure provisioning, hardware, software installation, platform upgrades, scaling, and internal metadata management such as micro-partitions and statistics. Customers do not manage physical hardware or install Snowflake software.

Customers are responsible for their data and its lifecycle within Snowflake. This includes loading data into tables from internal and external sources, unloading data when required, organizing data structures (databases, schemas, tables), defining access controls, and managing how data is used, transformed, and governed. They design schemas and workloads but do not manage the underlying engine. Therefore, "Loading, unloading, and managing data" correctly describes the customer's responsibility.

NEW QUESTION: 87

What is a fully qualified name in Snowflake used for?

- A. To store data in tables
- B. To manage user permissions
- C. To configure network settings
- D. To uniquely identify database objects

Answer: ([SHOW ANSWER](#))

A fully qualified name uniquely identifies Snowflake objects by specifying database.schema.object. This prevents ambiguity when multiple schemas or databases contain objects with identical

names. Fully qualified names ensure that SQL statements operate on the intended object. They are not used for storing data, managing permissions, or configuring network settings. Their core purpose is precise object identification.

NEW QUESTION: 88

What type of storage format does Snowflake use for its data?

- A. Row-based format
- B. Compressed columnar format
- C. Compressed row-store format
- D. Uncompressed row-store format

Answer: (SHOW ANSWER)

Snowflake stores data in a compressed columnar format, which is optimal for analytical workloads. Columnar storage allows Snowflake to read only the relevant columns needed for query execution, improving performance significantly for aggregation and filtering operations. Compression further reduces storage consumption, enabling cost efficiency. This approach differs from row-store formats used in OLTP systems and provides better analytic processing. Snowflake automatically manages partitioning, micro-partitioning, and metadata without user intervention.

NEW QUESTION: 89

Which of the following is a key feature of Snowflake's Data Storage Layer? (Select 2 options)

- A. In-memory storage
- B. Internal optimized, compressed, columnar format
- C. Elastic storage
- D. Elastic Cloud

Answer: (SHOW ANSWER)

The Data Storage Layer in Snowflake stores all data in a compressed, encrypted, columnar format designed for analytical efficiency. Snowflake automatically reorganizes ingested data into micro-partitions optimized for pruning and high-performance scans. The compression and structuring are fully automated, requiring no tuning or indexing by the user.

Another foundational feature is elastic storage, made possible by Snowflake's cloud-native architecture.

Storage automatically scales up or down based on data volume without user intervention, and is separated from compute. This ensures cost efficiency and immediate availability of capacity for large datasets.

Incorrect options:

* In-memory storage is not a Snowflake feature; long-term data is persisted in cloud object storage.

* Elastic Cloud is not a Snowflake architectural term.

The Data Storage Layer is designed for durability, scalability, and optimal query performance.

Valid SOL-C01 Dumps shared by ExamDiscuss.com for Helping Passing SOL-C01 Exam!
ExamDiscuss.com now offer the **newest SOL-C01 exam dumps**, the ExamDiscuss.com SOL-C01 exam **questions have been updated** and **answers have been corrected** get the **newest** ExamDiscuss.com SOL-C01 dumps with Test Engine here:

<https://www.examdiscuss.com/Snowflake/exam/SOL-C01/premium/> (218 Q&As Dumps,

35%OFF Special Discount Code: **freecram**)