

100% Pass 2026 High Pass-Rate Snowflake GES-C01: SnowPro® Specialty: Gen AI Certification Exam Latest Test Dumps



BTW, DOWNLOAD part of VCE4Plus GES-C01 dumps from Cloud Storage: https://drive.google.com/open?id=1fYAtAqIr-RIHUZwddLILNNfKUKy3g_H

Our company has applied the latest technologies to the design of our GES-C01 exam material not only on the content but also on the displays. So you are able to keep pace with the changeable world and remain your advantages with our GES-C01 Study Guide. Besides, you can consolidate important knowledge for you personally and design customized study schedule or to-do list on a daily basis with our GES-C01 learning questions.

Failure makes people depressed especially for working engineers. If your test score effects your work and you make mistakes, it is lost than gained. The best method for working people is to purchase valid Snowflake GES-C01 test questions and answers. It only takes you a little money to solve a big difficult for you. Also once you pass this subject, the certification is coming to you. Our passing rate of GES-C01 Test Questions and answers is normally 100% just one shot. It is worth buying.

>> GES-C01 Latest Test Dumps <<

HOT GES-C01 Latest Test Dumps - Snowflake SnowPro® Specialty: Gen AI Certification Exam - The Best Reliable GES-C01 Exam Question

We have accommodating group offering help 24/7. It is our responsibility to aid you through those challenges ahead of you. So instead of focusing on the high quality GES-C01 latest material only, our staff is genial and patient to your questions of our GES-C01 real questions. It is our obligation to offer help for your trust and preference. Besides, you can have an experimental look of demos and get more information of GES-C01 Real Questions. The customer-service staff will be with you all the time to smooth your acquaintance of our GES-C01 latest material.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q286-Q291):

NEW QUESTION # 286

A Snowflake administrator needs to implement a granular access control strategy for LLMs. The general policy is to restrict access to a select few models via an account-level allowlist. However, a specific data science team (using role 'DATA SCIENCE TEAM ROLE') requires access to the 'claude-3-5-sonnet' model, which should not be available to other users or globally via the allowlist. Given this scenario, which set of commands would correctly establish this access control while adhering to the specified requirements?

- A.
- B.
- C.
- D.

- E.

Answer: E

Explanation:

Option A is correct. This sequence of commands first sets an account-level allowlist for 'mistral-large?' and 'snowflake-arctic', thereby restricting general access to other models for plain-name string lookups. The 'CALL' ensures the changes are applied. It then explicitly grants the DATA SCIENCE_TEAM ROLES access to the 'claude-3-5-sonnet' model object using its dedicated application role 'SNOWFLAKE."CORTEX-MODEL-ROLE-CLAUDE-3-5-SONNET"'. This ensures 'claude-3-5-sonnet' is accessible only to that specific role and not globally through the allowlist, fulfilling the granular access requirement. Option B is incorrect because 'ALTER ACCOUNT' operations require the 'ACCOUNTADMIN' role, not 'SYSADMIN'. Additionally, setting to 'claude-3-5-sonnet' would make it globally available, contradicting the requirement for restricted access. Option C is incorrect because model-level RBAC for base models in 'SNOWFLAKE.MODELS' is primarily applied using application roles (e.g., 'CORTEX-MODEL-ROLE'), not directly with 'GRANT USAGE ON MODEL'. Option D is incorrect. While clearing the allowlist is a valid part of a strategy, 'GRANT USAGE ON ALL MODELS IN SCHEMA SNOWFLAKE.MODELS' would grant access to 'all' models in that schema, which contradicts the requirement for 'claude-3-5-sonnet' to be exclusive to the data science team and not generally available. Option E is incorrect because 'ALTER ACCOUNT' requires the 'ACCOUNTADMIN' role, not 'SECURITYADMIN', and setting the allowlist to 'claude-3-5-sonnet' would make it generally available, violating the isolation requirement.

NEW QUESTION # 287

An enterprise is deploying a new RAG application using Snowflake Cortex Search on a large dataset of customer support tickets. The operations team is concerned about managing compute costs and ensuring efficient index refreshes for the Cortex Search Service, which needs to be updated hourly. Which of the following considerations and configurations are relevant for optimizing cost and performance of the Cortex Search Service in this scenario?

- A. The primary cost driver for Cortex Search is the number of search queries executed against the service, with the volume of indexed data (GB/month) having a minimal impact on overall billing.
- B. CHANGE_TRACKING
- C. For embedding text, selecting a model like
- D. For optimal performance and cost efficiency, Snowflake recommends using a dedicated warehouse of size no larger than MEDIUM for each Cortex Search Service.
- E. The

Answer: B,C,D,E

Explanation:

Option A is correct because a Cortex Search Service requires a virtual warehouse to refresh the service, which runs queries against base objects when they are initialized and refreshed, incurring compute costs. Option B is correct because the cost of embedding models varies. For example, 'snowflake-arctic-embed-m-v1.5' costs 0.03 credits per million tokens, while 'voyage-multilingual-2' costs 0.07 credits per million tokens. Choosing a more cost-effective model like 'snowflake-arctic-embed-m-v1.5' for English-only data can reduce token costs. Option C is correct because Snowflake recommends using a dedicated warehouse of size no larger than MEDIUM for each Cortex Search Service to achieve optimal performance. Option D is correct because change tracking is required for the Cortex Search Service to be able to detect and process updates to the base table, enabling incremental refreshes that are more efficient than full re-indexing. Option E is incorrect because Cortex Search Services incur costs based on virtual warehouse compute for refreshes, 'EMBED_TEXT_TOKENS' cost per input token, and a charge of 6.3 Credits per GB/mo of indexed data. The volume of indexed data has a significant impact, not minimal.

NEW QUESTION # 288

A data analyst is tasked with identifying customers who purchased items with similar feature vectors. They have a table products with an

to measure similarity. Which of the following statements correctly describe aspects of defining and using vector types or functions in this scenario? (Select all that apply)

- A. If
- B. Comparing two
- C. To correctly define a column to store 768-dimensional float embeddings, the SQL statement

-
- D. When inserting literal arrays as vectors into a table for comparison, explicit casting, e.g.,
-
- E. The Snowpark Python library provides native support for calling

Answer: C,D,E

Explanation:

Option A is correct. The syntax for specifying a

□) are byte-wise lexicographic and do not produce semantically expected results for number comparisons; dedicated vector similarity functions should be used instead. Option D is correct. The Snowpark Python library supports the VECTOR data type and vector similarity functions. While the sources specifically mention VECTOR_L2_DISTANCE in a Snowpark Python example, VECTOR_L1_DISTANCE is listed as one of the four core vector similarity functions provided by Snowflake Cortex, implying similar support in Snowpark Python. Option E is correct. SQL examples demonstrate the necessity of explicit casting when using array literals as vectors, such as .

NEW QUESTION # 289

A financial analyst is concerned about the rising costs of their Document AI pipeline, which uses 'invoice_model!PREDICT' to extract data from daily financial reports. They observe that their assigned 'LARGE' virtual warehouse is running continuously, even during periods of low document ingestion, contributing significantly to their bill. They want to investigate how to reduce costs effectively for their existing Document AI setup.

-
- A. Option D
- B. Option C
- C. Option E
- D. Option A
- E. Option B

Answer: E

Explanation:

NEW QUESTION # 290

A data analyst is using Snowflake Copilot in Snowsight to generate SQL queries for a new dataset containing customer PII. Which of the following statements accurately describes how Snowflake Copilot operates with respect to data access, governance, and model interaction?

- A. Snowflake Copilot requires explicit column-level grants for direct data access, similar to how a human analyst would query specific data points.
- B. To protect sensitive information, Snowflake Copilot transmits sampled PII data to an external LLM for schema understanding before generating SQL.
- C. While Snowflake Copilot generates SQL based on metadata, the generated SQL queries are executed in an isolated environment that does not respect existing Snowflake RBAC policies.
- D. Snowflake Copilot directly accesses and processes the raw data within customer tables to understand its content and generate SQL.
- E. Snowflake Copilot is powered by a fine-tuned model that runs securely inside Snowflake Cortex, leveraging only database/schema/table/column names and data types, ensuring data remains within Snowflake's governance boundary and respects RBAC.

Answer: E

Explanation:

Snowflake Copilot is an LLM-powered assistant that is powered by a model fine-tuned by Snowflake, running securely inside Snowflake Cortex. It ensures that your enterprise data and metadata always stay securely inside Snowflake and fully respects RBAC. Crucially, Snowflake Copilot does not have access to the data inside your tables; it generates responses based on the names of your databases, schemas, tables, and columns, and their data types. Options A and C are incorrect because Copilot does not directly access or transmit customer data. Option D is incorrect as it implies direct data access, which Copilot does not perform. Option E is incorrect because Copilot fully integrates with Snowflake's RBAC policies.

P.S. Free & New GES-C01 dumps are available on Google Drive shared by VCE4Plus: https://drive.google.com/open?id=1fYAtAqIr-RIHUZwddLLNNffKUKY3g_H