

Snowflake GES-C01 Questions 2025 - All Subjects Covered



BONUS!!! Download part of UpdateDumps GES-C01 dumps for free: https://drive.google.com/open?id=1u8uoXo_y758flksy9x7vuF6uTnoa5TsC

Most people define GES-C01 study tool as regular books and imagine that the more you buy, the higher your grade may be. It is true this kind of view make sense to some extent. However, our GES-C01 real questions are high efficient priced with reasonable amount, acceptable to exam candidates around the world. Our GES-C01 practice materials comprise of a number of academic questions for your practice, which are interlinked and helpful for your exam. Just hold the supposition that you may fail the exam even by the help of our GES-C01 Study Tool, we can give full refund back or switch other versions for you to relieve you of any kind of losses. What is more, we offer supplementary content like updates for one year after your purchase.

UpdateDumps examines it regularly for new updates so that you always get new SnowPro® Specialty: Gen AI Certification Exam (GES-C01) practice questions. Since it is a printable format, you can do a paper study. The SnowPro® Specialty: Gen AI Certification Exam (GES-C01) PDF Dumps document is accessible from every location at any time. This SnowPro® Specialty: Gen AI Certification Exam (GES-C01) software has a simple-to-use interface. By using the SnowPro® Specialty: Gen AI Certification Exam (GES-C01) practice exam software, you can evaluate your mistakes at the end of every take and overcome them. Our software helps you to get familiar with the format of the original SnowPro® Specialty: Gen AI Certification Exam (GES-C01) test.

>> GES-C01 Valid Test Vce <<

Pass-Sure GES-C01 Valid Test Vce - Updated Source of GES-C01 Exam

After years of operation, our platform has accumulated a wide network of relationships, so that we were able to learn about the

changes in the exam at the first time. This is a benefit that students who have not purchased GES-C01 exam guide can't get. The team of experts hired by SnowPro® Specialty: Gen AI Certification Exam study questions constantly updates and supplements the contents of study materials according to the latest syllabus and the latest industry research results. We also have dedicated staff to maintain GES-C01 Exam Material every day, and you can be sure that compared to other test materials on the market, SnowPro® Specialty: Gen AI Certification Exam study questions are the most advanced.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q221-Q226):

NEW QUESTION # 221

A machine learning team is leveraging the Snowflake Model Registry to manage diverse models, including a custom Python utility for data preprocessing that they wish to make available as a model method. Which of the following statements accurately describe capabilities or considerations when logging models and their associated artifacts and methods in the Model Registry?

- ☐ The Snowflake Model Registry supports built-in types such as Scikit-learn, XGBoost, and PyTorch, but does not allow logging custom Python objects or processing code directly as models.
- ☐ To include additional local files, such as configuration files or custom scripts, with a logged model, the `user_files` argument must be used in `log_model`, mapping stage subdirectories to local file paths.
- ☐ Once a model version is logged, its methods can be invoked using either `mv.run()` in Python or through service functions named `<service_name>.<method_name>` in SQL, after the model has been deployed to SPCS.
- ☐ The `function_type` option within `method_options` in `log_model` allows specifying whether a model method should be exposed as a `FUNCTION` or `TABLE_FUNCTION` in SQL, influencing how data is processed.
- ☐ The maximum total model size for models deployed to a Snowflake warehouse is 5 GB, whereas models deployed to SPCS have no such size limitations.

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

Answer: A,B,D,E

Explanation:

Option B is correct because the `user_files` argument in `log_model` is specifically designed for including additional local files (e.g., configuration, scripts) with the model, by mapping them to stage subdirectories. Option C is correct because after a model is deployed to SPCS Snowflake Model Serving creates service functions named `<service_name>.<method_name>` that allow invoking the model's methods from SQL. Model methods can also be called via `mv.run()` in Python. Option D is correct because the `function_type` option within `method_options` allows developers to specify if a model method should be registered as a `FUNCTION` or `TABLE_FUNCTION` when exposed in SQL, affecting its input/output processing. Option E is correct. The maximum total model size for warehouse-deployed models is 5 GB. For models running on Snowpark Container Services, these size restrictions are eased or eliminated, allowing for much larger models. Option A is incorrect; the Model Registry is flexible enough to support not only built-in types (like Scikit-learn, XGBoost, PyTorch) but also custom processing code and previously-trained models.

NEW QUESTION # 222

An administrator has configured the 'CORTEX_MODELS_ALLOWLIST' parameter to only permit the 'mistral-large?' model at the account level. A user with the 'PUBLIC' role, which has been granted 'SNOWFLAKE.CORTEX USER' and 'SNOWFLAKE."CORTEX- MODEL-ROLE-LLAMA3.1-70B"', attempts to execute several 'AI COMPLETE' queries. Which of the following queries will successfully execute?

- A.

```
SELECT AI_COMPLETE('SNOWFLAKE-MODELS."LLAMA3.1-70B"', 'Hello');
```
- B.

```
SELECT AI_COMPLETE('SNOWFLAKE-MODELS."CORTEX-MODEL-ROLE-LLAMA3.1-70B"', 'Hello');
```
- C.

```
SELECT AI_COMPLETE('MISTRAL-LARGE2', 'Hello');
```
- D.

```
SELECT AI_COMPLETE('MISTRAL-LARGE2', 'Hello');
```
- E.

```
SELECT AI_COMPLETE('MISTRAL-LARGE2', 'Hello');
```

Answer: B,C

Explanation:

Option A is correct. The query directly references 'MISTRAL-LARGE2', which is explicitly in the account-level, so it will succeed. Option B is correct. When a model name is provided as a string argument, Cortex first treats it as an identifier for a schema-level

model object. If found, RBAC is applied. The user's role has 'SNOWFLAKE.'CORTEX-MODEL-ROLE-LLAMA3.1-70B' granted, which provides access to the 'LLAMA3.1-70B' model object in 'SNOWFLAKE.MODELS', regardless of the 'CORTEX MODELS_ALLOWLIST' setting for plain model names. Option C is incorrect because 'llama3.1-70b' as a plain model name is not in the 'CORTEX_MODELS_ALLOWLIST'. Although the user has access to the model object, a plain string like 'llama3.1-70b' will be looked up in the allowlist after failing to match a model object by that plain name, and the allowlist only has 'MISTRAL-LARGE2'. Option D is incorrect. 'snowflake-arctic' is neither in the 'CORTEX MODELS_ALLOWLIST' nor does the user have a specific application role granting access to a 'snowflake-arctic' model object. Option E is incorrect because 'ALTER ACCOUNT' operations can only be performed by the 'ACCOUNTADMIN' role, not by a 'SPUBLIC' user role.

NEW QUESTION # 223

A Gen AI developer is implementing a Document AI solution to extract key fields from thousands of diverse PDF reports, which vary significantly in length and complexity. They use the 'PREDICT' method with 'GET_PRESIGNED_URL' to process documents from an external stage. After initial testing, they observe two distinct types of errors in the query results:

```
1.
{ "__processingErrors": [ "Received HTTP 403 response for presigned URL. URL may be expired." ] }
for some documents.
2.
{ "__processingErrors": [ "Document has too many pages. Actual: 150. Maximum: 125." ] }
```

for other, lengthy PDF files. Which two of the following actions should the developer take to resolve these issues?

- A. Grant the 'CREATE STREAM'

and

'CREATE TASK

privileges to the role executing the 'PREDICT

'MATERIALIZED VIEW'.

- B. Redesign the input documents to ensure they do not exceed 125 pages per file, or preprocess by splitting overly long documents into multiple smaller files.
- C. Increase the virtual warehouse size to a Large or X-Large to speed up processing and prevent URL expiration.
- D. Implement a mechanism to process documents in smaller batches or extend the expiration time for the presigned URLs to ensure timely access by Document AI.
- E. Reconfigure the external stage to use

Answer: B,D

Explanation:

The first error, 'Received HTTP 403 response for presigned URL. URL may be expired.', indicates that the presigned URLs generated by 'GET_PRESIGNED_URL' are expiring before the Document AI model can process the documents. The 'GET_PRESIGNED_URL' function has a default expiration time of 60 minutes. A recommended solution is to use several queries to process documents, effectively breaking the workload into smaller batches to complete within the URL's lifespan. The second error, 'Document has too many pages. Actual: 150. Maximum: 125.', directly indicates that some documents exceed Document AI's page limit of 125 pages per document. The solution is to ensure documents comply with this limit, possibly by splitting them. Option A is incorrect because scaling up the warehouse does not increase the speed of query processing for Document AI; X-Small, Small, or Medium warehouses are recommended, and larger warehouses may result in unnecessary costs. Option D relates to setting up continuous processing pipelines but does not directly address these specific extraction errors. Option E is applicable for internal stages requiring encryption, but the scenario specifies an external stage, which would have different encryption configurations (e.g., SAWS SSE SS for S3).

NEW QUESTION # 224

A financial analyst is concerned about the rising costs of their Document AI pipeline, which uses 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.

- Query the 'SNOWFLAKE.ACCOUNT_USAGE.METERING_DAILY_HISTORY' view, filtering by 'SERVICE_TYPE = 'WAREHOUSE_METERING'', to understand Document AI's specific credit consumption.
- Scaling down the warehouse to 'X-SMALL', 'SMALL', or 'MEDIUM' is recommended, as larger warehouses do not increase Document AI query processing speed and incur unnecessary costs.
- Document AI's '!PREDICT' method costs are primarily based on the number of tokens processed for each document, so reducing document length will be the most impactful cost-saving measure.
- Replace the 'invoice_model!PREDICT' function with 'AI_PARSE_DOCUMENT' as it is a newer, more cost-efficient function for document text extraction.
- The 'SNOWFLAKE.ACCOUNT_USAGE.CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY' view tracks only 'AI_EXTRACT' calls, making it unsuitable for monitoring '!PREDICT' function usage.

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

Answer: E

Explanation:

Snowflake explicitly recommends using an X-Small, Small, or Medium warehouse for Document AI. Scaling up the warehouse does not increase the speed of query processing for Document AI but can lead to unnecessary costs. This directly addresses the scenario of a 'LARGE' warehouse running continuously and contributing to high bills. Option A is incorrect because while 'METERING_DAILY_HISTORY' is used for cost tracking, Document AI's service-side usage appears under 'AI_SERVICES', not 'WAREHOUSE_METERING' for the AI service component itself. 'WAREHOUSE_METERING' would show general warehouse costs, not specifically tied to Document AI's compute portion. Option C is incorrect because Document AI (using '!PREDICT') incurs 'AI Services compute' costs based on 'time spent actually using these resources' (8 Credits per hour of compute), not per token. Option D is not necessarily accurate guidance; 'SAI_PARSE_DOCUMENT' is a separate Cortex AI SQL function for document processing, billed per page, while Document AI's '!PREDICT' is part of a Document AI model build. Replacing it without a full re-evaluation of the workflow might not be optimal or directly cost-efficient for an established pipeline. Option E is incorrect because the view tracks Document AI processing activity, including '!PREDICT' calls.

NEW QUESTION # 225

A team of data application developers is leveraging Snowflake Copilot to streamline the creation of analytical SQL queries within their Streamlit in Snowflake application. They observe that Copilot sometimes struggles with complex joins or provides suboptimal queries when dealing with a newly integrated, deeply nested dataset. Based on Snowflake's best practices and known limitations, which actions or considerations would help improve Copilot's performance in this scenario?

- A. Grant Copilot direct access to the raw data using ACCOUNTADMIN privileges, allowing it to infer schema relationships more effectively from data content.
- B. Break down complex requests into simpler, multi-turn questions, as Copilot is designed to build complex queries through conversational refinement and follow-up questions.
- C. Enable the CORTEX_MODELS_ALLOWLIST parameter to restrict Copilot to only use the largest available LLMs, thereby guaranteeing higher accuracy for complex queries.
- D. Ensure that a database and schema are explicitly selected for the current session, and that column names are meaningful, to provide Copilot with better context for query generation.
- E. Implement curated views with descriptive and easy-to-understand names for views and columns, appropriate data types, and pre-define common/complex joins to simplify the underlying schema for Copilot.

Answer: B,D,E

Explanation:

To improve Snowflake Copilot's performance, creating curated views with descriptive names, appropriate data types, and capturing common/complex joins is a key best practice. Copilot can build complex queries through a conversation by asking follow-up questions. It also uses the names of databases, schemas, tables, and columns, and their data types to determine available data, so ensuring these are meaningful and correctly set for the session is crucial for relevant responses. Option B is incorrect because CORTEX_MODELS_ALLOWLIST controls access to specific LLMs but doesn't guarantee higher accuracy for Copilot's SQL generation. Option D is incorrect as Snowflake Copilot does not have access to the data inside tables; it operates on metadata. Granting privileges would not change this fundamental operational principle and is ACCOUNTADMIN against best practices for

least privilege.

NEW QUESTION # 226

• • • • •

Our GES-C01 study materials have a high quality which is mainly reflected in the pass rate. Our product can promise a higher pass rate than other study materials. 99% people who have used our GES-C01 study materials passed their exam and got their certificate successfully, it is no doubt that it means our GES-C01 Study Materials have a 99% pass rate. So our product will be a very good choice for you. If you are anxious about whether you can pass your exam and get the certificate, we think you need to buy our GES-C01 study materials as your study tool, our product will lend you a good helping hand.

Guaranteed GES-C01 Success: <https://www.updatedumps.com/Snowflake/GES-C01-updated-exam-dumps.html>

The most important thing is that the GES-C01 exam questions are continuously polished to be sold, so that users can enjoy the best service that our products bring. Our GES-C01 exam prep file has won good reputation among numerous candidates and peers in the industry through continuous GES-C01 innovation and creation. In particular, our experts keep the GES-C01 real test the latest version, they check updates every day and send them to your e-mail in time, making sure that you know the latest news.

All contents of GES-C01 training guide are being explicit to make you have explicit understanding of this exam, Create the Bankruptcies Database Control, The most important thing is that the GES-C01 Exam Questions are continuously polished to be sold, so that users can enjoy the best service that our products bring.

Pass Guaranteed Quiz GES-C01 - SnowPro® Specialty: Gen AI Certification Exam –Reliable Valid Test Vce

Our GES-C01 exam prep file has won good reputation among numerous candidates and peers in the industry through continuous GES-C01 innovation and creation, In particular, our experts keep the GES-C01 real test the latest version, they check updates every day and send them to your e-mail in time, making sure that you know the latest news.

Spending little money is to do great things, We insist to keep our GES-C01 customers' information secret and never share the information to any other third part without the permission of the customer.

- [illegible]

myportal.utt.edu.tt, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, lms.ait.edu.za, www.stes.tyc.edu.tw, Disposable vapes

BONUS!!! Download part of UpdateDumps GES-C01 dumps for free: https://drive.google.com/open?id=1u8uoXo_v758flksy9x7vuF6uTnoa5TsC