

# Maximize Your Chances of Getting Snowflake GES-C01 Certification Exam



As one of the leading brand in the market, our GES-C01 exam materials can be obtained on our website within five minutes. As long as you pay for our GES-C01 study guide successfully, then you will receive it quickly. That is the expression of our efficiency. The amazing quality of our GES-C01 learning questions can totally catch eyes of exam candidates with passing rate up to 98 to 100 percent.

In traditional views, the GES-C01 practice materials need you to spare a large amount of time on them to accumulate the useful knowledge may appearing in the real GES-C01 exam. However, our GES-C01 learning questions are not doing that way. According to data from former exam candidates, the passing rate of our GES-C01 learning material has up to 98 to 100 percent. There are adequate content to help you pass the exam with least time and money.

>> **GES-C01 Valid Test Objectives** <<

## Snowflake GES-C01 Test Questions Vce & GES-C01 Exam Discount Voucher

The PDF version of our GES-C01 learning guide is convenient for reading and supports the printing of our study materials. If client uses the PDF version of GES-C01 exam questions, they can download the demos freely. If clients feel good after trying out our demos they will choose the full version of the test bank to learn our GES-C01 Study Materials. And the PDF version can be printed into paper documents and convenient for the client to take notes.

## Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q232-Q237):

### NEW QUESTION # 232

A data application developer is tasked with creating a multi-turn conversational AI application using Streamlit in Snowflake (SIS), which will leverage Snowflake Cortex LLM functions. Considering the core requirements for building such an interactive chat interface and the underlying Snowflake environment, which of the following actions is a fundamental step in setting up the application for stateful conversations?

- Ensuring the application's role has the `SNOWFLAKE.DOCUMENT_INTELLIGENCE_CREATOR` database role granted to enable multi-turn capabilities.
- Implementing a mechanism to store and retrieve the full conversation history, typically by initializing `st.session_state.messages = []` and appending user/assistant messages for each turn.
- Configuring Snowpark Container Services (SPCS) with a dedicated compute pool and an image repository to host the Streamlit application.
- Setting the `CORTEX_ENABLED_CROSS_REGION` parameter to `TRUE` in the Snowflake account to allow local execution of LLM functions.
- Using the `TRY_COMPLETE (SNOWFLAKE.CORTEX)` function with the `'ON_ERROR': 'SKIP'` configuration to gracefully handle LLM errors in conversations.

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

## Answer: C

### Explanation:

For a multi-turn conversational AI application built with Streamlit, maintaining the conversation history is fundamental. Streamlit's `st.session_state` is the primary way to store and manage state across reruns of the application, which is crucial for remembering past interactions in a chat interface. The typical approach involves initializing `st.session_state.messages` to an empty list and appending messages for each turn. Option A is incorrect because it is a database role specific to Document AI, not general Cortex LLM functions. Option C is not a fundamental step for running a Streamlit application in Snowflake (SiS) itself, as SiS directly hosts the Streamlit app; while models can be served via SPCS, the application itself doesn't inherently require it for basic operation. Option D is related to cross-region inference for LLM functions, which controls where inference requests are processed, not a fundamental step for local execution or conversational state management. Option E suggests a configuration ("ON ERROR:SKIP") that is primarily used with Snowflake ML functions like Anomaly Detection and Time-Series Forecasting to prevent overall training failure for individual series, and is not a direct option for handling errors in `'TRY_COMPLETE` in this manner; `'TRY_COMPLETE` itself returns `NULL` on error.

## NEW QUESTION # 233

A global enterprise has Snowflake accounts in various regions, including a US East (Ohio) account where a critical application is deployed. They need to use `AI_COMPLETE` with the `claude-3-5-sonnet` model for real-time customer support, but this model is not natively available in US East (Ohio) for direct `AI_COMPLETE` usage. The Snowflake administrator considers enabling cross-region inference. Which statements accurately reflect the considerations and characteristics of cross-region inference in Snowflake Cortex?

- A. Setting the `CORTEX_ENABLED_CROSS_REGION` account parameter to '`ANY_REGION`' in the US East (Ohio) account would enable inference requests for `claude-3-5-sonnet` to be processed in any region where it is natively available.
- B. Cross-region inference automatically caches user inputs and generated outputs to reduce latency for subsequent requests to the same model.
- C. The `CORTEX_ENABLED_CROSS_REGION` parameter can be configured at the session level to temporarily enable cross-region inference for specific workloads.
- D. Latency between regions for cross-region inference is negligible and consistently low, irrespective of cloud provider infrastructure.
- E. Cross-region inference is not supported in U.S. SnowGov regions for either inbound or outbound inference requests.

## Answer: A,E

### Explanation:

Option B is correct because setting the parameter to '`ANY_REGION`' enables inference requests to be `CORTEX_ENABLED_CROSS_REGION` processed in a different region from the default, thereby allowing access to models not natively supported in the local region. For example, `claude-` is `3-5-` sonnet available in AWS US East 1 (N. Virginia), which could be accessed from US East (Ohio) via cross-region inference. Option C is `3-5-` sonnet correct as cross-region inference is explicitly not supported in U.S. SnowGov regions. Option A is incorrect because user inputs, service generated prompts, and outputs are not stored or cached during cross-region inference. Option D is incorrect; latency depends on the cloud provider infrastructure and network status, and testing is recommended. Option E is incorrect because `CORTEX_ENABLED_CROSS_REGION` is an account-level parameter, not a session parameter.

## NEW QUESTION # 234

A data engineer is tasked with establishing AI Observability for a generative AI application that integrates with external systems and will undergo continuous improvement. The goal is to compare different iterations of the application efficiently. Which combination of configuration best practices, features, and governance aspects are most relevant for a robust setup of AI Observability within Snowflake for this scenario?

- A. If the AI Observability service is not natively available in the primary region, enable to '`ANY_REGION`' or a specific supported region to allow tracing and evaluation to proceed.
- B. Run the AI Observability project directly within a Snowflake Notebook to leverage its integrated environment for easier debugging and iteration.
- C. To compare different LLMs or prompt configurations, rely on the AI Observability's 'Comparisons' feature, which allows side-by-side analysis of evaluation metrics across multiple evaluations.
- D. For access control, the role used to create and execute runs must be granted the '`SNOWFLAKE.CORTEX_USER`' database role and the '`AI_OBSERVABILITY_EVENTS_LOOKUPS`' application role.
- E. Ensure the Python environment includes '`trulens-core`', '`trulens-connectors-snowflake`', and '`trulens-providers-cortex`'

(version 2.1.2 or later) and set the environment variable `TRULENS_OTEL_TRACING` to 1.

**Answer: A,C,D,E**

Explanation:

Option A is correct because installing the specified TruLens Python packages (version 2.1.2 or later) and setting `STRULENS_OTEL_TRACINGS` to are prerequisites for instrumenting the application and enabling tracing for AI Observability. Option B is correct because the '`CORTEX_USER`' database role and application role are explicitly required for creating and executing runs for AI Observability. Option C is correct as the '`Comparisons`' feature is a core component of AI Observability, designed precisely for assessing and comparing application quality, accuracy, and performance across various LLMs, prompts, and configurations. Option D is correct because AI Observability, like other Cortex LLM Functions, might require '`CORTEX_ENABLED_CROSS_REGION`' to be configured if the service or specific LLMs are not natively available in the primary Snowflake region. Option E is incorrect because the sources explicitly state that you cannot run your project using the TruLens SDK in a Snowflake Notebook for AI Observability.

### NEW QUESTION # 235

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('MISTRAL-LARGE2', 'Hello');`
- B. `SELECT AI_COMPLETE('LLAMA3.1-70B', 'Hello');`
- C. `ALTER ACCOUNT SET CORTEX_MODELS_ALLOWLIST = 'SNOWFLAKE-ARCTIC';`
- D. `SELECT AI_COMPLETE('llama3.1-70b', 'Hello');`
- E. `SELECT AI_COMPLETE('snowflake-arctic', 'Hello');`

**Answer: A,D**

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 # 236

A company is planning to process a large volume of legal documents to generate summaries using `SNOWFLAKE.CORTEX_SUMMARIZE`. Given the scale, they are acutely focused on managing costs and optimizing performance. Which of the following statements are true regarding the cost and performance characteristics of using `SNOWFLAKE.CORTEX_SUMMARIZE?` (Select all that apply)

- A. For `SUWARIZE`, Snowflake adds an internal prompt to the user's input text, which increases the total input token count for billing purposes beyond the raw text length.
- B. The fixed billing rate for the `SUMMARIZE` function is 0.10 Credits per one million Tokens processed.
- C. The `SUWARIZE` function is billed primarily based on the number of output tokens generated in the response, not input tokens.
- D. Snowflake recommends using a larger warehouse (e.g., L or XL) for `SUMMARIZE` function calls to significantly improve processing performance for high-volume tasks.
- E. The context window for the `SWIMARIZE` function is 4,096 tokens, ensuring efficiency for short documents only.

**Answer: A,B**

### Explanation:

Options B and D are correct. - is correct": For 'SUMMARIZE', Snowflake adds an internal prompt to the input text in order to generate the response, which results in a higher input token count for billing than the raw text provided. - is correct": The cost for the 'Summarize' function is 0.10 Credits per one million Tokens processed. - A is incorrect: For functions that generate new text in the response, such as 'SUMMARIZE', both input and output tokens are billable. - C is incorrect: Snowflake recommends executing queries that call a Snowflake Cortex AISQL function, including 'SUMMARIZE', with a smaller warehouse (no larger than MEDIUM), as larger warehouses do not increase performance for these operations. - E is incorrect: The context window for the Snowflake managed model from the 'SUMMARIZE' function is 32,000 tokens, not 4,096 tokens.

## NEW QUESTION # 237

Our company is a professional certificate exam materials provider, we have occupied the field for years, therefore we have rich experiences. GES-C01 training materials of us are compiled by skilled experts, and they are quite familiar with the exam center, and you can pass the exam just one time by using GES-C01 Exam Materials of us. In addition, we offer you free update for 365 days after purchasing, and the update version for GES-C01 training materials will be sent to your email automatically. We have online and offline chat service stuff, if you have any questions, just contact us.

GES-C01 Test Questions Vce: <https://www.exam4labs.com/GES-C01-practice-torrent.html>

If you try on it, you will find that the operation systems of the GES-C01 Dumps exam questions we design have strong compatibility, Snowflake GES-C01 Valid Test Objectives On the other hand, customers who have failed in the exam luckily can ask for full refund or changing other exam files for free, What's more, time witnesses that our GES-C01 test prep have 100% passing rate, Therefore, users can pass GES-C01 exams with very little learning time.

Implement insert, update, delete, and merge GES-C01 data modification strategies. Few see it having a negative impact. If you try on it, you will find that the operation systems of the GES-C01 Dumps exam questions we design have strong compatibility.

## Pass Guaranteed Snowflake - GES-C01 - SnowPro® Specialty: Gen AI Certification Exam Newest Valid Test Objectives

On the other hand, customers who have failed in the exam luckily can ask for full refund or changing other exam files for free. What's more, time witnesses that our GES-C01 test prep have 100% passing rate.

Therefore, users can pass GES-C01 exams with very little learning time. So you need a right training material to help you.

myportal.utt.edu.tt, tjjaabo.dadweynahacilmi.com, thesocraticmethod.in, www.so0912.com, pct.edu.pk, pct.edu.pk,  
Disposable vapes