

# Free PDF Snowflake - GES-C01 - SnowPro® Specialty: Gen AI Certification Exam–Valid Exams Collection

---

## Snowflake GES-C01 Exam

### SnowPro® Specialty: Gen AI Certification Exam

<https://www.passquestion.com/ges-c01.html>



Pass GES-C01 Exam with PassQuestion GES-C01 questions and answers in the first attempt.

<https://www.passquestion.com/>

---

1 / 23

BONUS!!! Download part of PassLeader GES-C01 dumps for free: [https://drive.google.com/open?id=1Xi8pA5gRJ9DEu\\_MwI6vcLVNp1P9sodmx](https://drive.google.com/open?id=1Xi8pA5gRJ9DEu_MwI6vcLVNp1P9sodmx)

GES-C01 certification can demonstrate your mastery of certain areas of knowledge, which is internationally recognized and accepted by the general public as a certification. GES-C01 certification is so high that it is not easy to obtain it. It requires you to invest time and energy. If you are not sure whether you can strictly request yourself, our GES-C01 Exam Training can help you. Help is to arrange time for you and provide you with perfect service. What are the advantages of our GES-C01 test guide? I hope you can take a moment to find out.

Our GES-C01 exam prep is elaborately compiled and highly efficiently, it will cost you less time and energy, because we shouldn't waste our money on some unless things. The passing rate and the hit rate are also very high, there are thousands of candidates choose to trust our GES-C01 Guide Torrent and they have passed the exam. We provide with candidate so many guarantees that they can purchase our study materials no worries. The GES-C01 exam prep we provide can help you realize your dream to pass exam and then own a GES-C01 exam torrent.

>> GES-C01 Exams Collection <<

**GES-C01 Exams Collection Pass-Sure Questions Pool Only at PassLeader**

Our GES-C01 exam guide has high quality of service. We provide 24-hour online service. If you have any questions in the course of using the GES-C01 exam questions, you can contact us by email. We will provide you with excellent after-sales service with the utmost patience and attitude. And we will give you detailed solutions to any problems that arise during the course of using the GES-C01 practice torrent. And our GES-C01 study materials welcome your supervision and criticism. With the company of our GES-C01 study materials, you will find the direction of success.

## Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q14-Q19):

### NEW QUESTION # 14

An operations team at a company is implementing a robust governance framework to monitor and optimize the costs associated with their Snowflake Cortex LLM function usage. They need to identify which functions are driving the highest token consumption and overall credit usage to pinpoint areas for cost reduction. Which of the following monitoring tools or methods are appropriate for gaining these insights into Cortex LLM function costs and token consumption?

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

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

Explanation:

### NEW QUESTION # 15

A data engineering team is onboarding a new client whose workflow involves extracting critical financial data from thousands of daily scanned PDF receipts. They decide to use Snowflake Document AI and store all incoming PDFs in an internal stage name. After deploying their pipeline, they observe intermittent failures and varying error messages in the output, specifically:

Which two of the following actions are most likely required to resolve these processing errors?

- A. Change the virtual warehouse size from an X-Small to a Large to improve Document AI processing speed.
- B. Grant the 'SNOWFLAKCORTEX\_USER' database role to the role executing the '!PREDICT' function.
- C. Ensure the internal stage is configured with 'ENCRYPTION = (TYPE = 'SNOWFLAKE\_SSE')'.
- D. Increase the 'max\_tokens' parameter within the '!PREDICT' function options to accommodate longer document processing.
- E. Split any PDF documents exceeding 125 pages into smaller, compliant files, or reject them if splitting is not feasible.

**Answer: C,E**

Explanation:

The first error message, 'cannot identify image file', is a known error that occurs when an internal stage used for Document AI is not configured with 'SNOWFLAKE\_SSE' encryption. Therefore, option A is a direct solution. The second error message, 'Document has too many pages. Actual: 130. Maximum: 125.', indicates that some documents exceed Document AI's page limit of 125 pages per document. Option B directly addresses this limitation. Option C is incorrect because 'max\_tokens' is relevant for LLM output length, not document input page/size limits. Option D is incorrect because scaling up the warehouse for Document AI does not increase query processing speed and is not recommended for cost efficiency; X-Small, Small, or Medium warehouses are typically sufficient for Document AI. Option E is incorrect because is the required database role for Document AI, not 'SNOWFLAKE.CORTEX\_USER'.

### NEW QUESTION # 16

An AI engineer is building an automated pipeline in Snowflake that processes various types of textual data using Cortex AI functions. To ensure the pipeline's stability and avoid failures due to exceeding LLM context windows, they integrate SNOWFLAKE.CORTEX.COUNT\_TOKENS and TRY\_COMPLETE

. Consider the following code snippets and statements about context window management in Snowflake Cortex.

- A. Option B
- B. Option A

- C. Option D
- **D. Option E**
- E. Option C

**Answer: B,D,E**

Explanation:

Option A is correct.

□ functions are explicitly stated as not being subject to context window limitations. This means input length is less of a concern for truncation for these specific functions, though

□ has a 200,000-token context window, mistral -7b has a 32,000-token context window) before invoking the LLM function. Option D is incorrect. The

□ controls the maximum number of \*output\* tokens the model can generate, not the input tokens. Option E is correct.

TRY\_COMPLETE is a helper function that returns NULL instead of raising an error if the operation cannot be performed.

Integrating COUNT\_TOKENS before TRY\_COMPLETE can proactively identify potential token overflow issues, thus helping to prevent the operation from failing or returning NULL, and enhancing pipeline stability.

### NEW QUESTION # 17

A developer is instrumenting a RAG application using the TruLens SDK within Snowflake AI Observability. The application has distinct functions for retrieving context and generating a completion. To ensure clear tracing and readability, which span\_type should ideally be used for the function responsible for retrieving relevant text from the vector store?

- A. □
- **B. □**
- C. □
- D. □
- E. No specific span\_type is needed; the default instrumentation is sufficient for all functions.

**Answer: B**

Explanation:

The TruLens SDK allows for specifying span\_type to improve the readability and understanding of traces. For a RAG application, RETRIEVAL the span type is explicitly recommended for search services or retrievers (functions that retrieve context).

GENERATION is used for LLM inference calls that generate answers, and RECORD\_ROOT identifies the entry point method of the application.

### NEW QUESTION # 18

A data application developer is tasked with building a multi-turn conversational AI application using Streamlit in Snowflake (SIS) that leverages the COMPLETE (SNOWFLAKE.CORTEX) LLM function. To ensure the conversation flows naturally and the LLM maintains context from previous interactions, which of the following is the most appropriate method for handling and passing the conversation history?

□

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

**Answer: D**

Explanation:

To provide a stateful, conversational experience with the 'COMPLETE (SNOWFLAKE.CORTEX)' function, all previous user prompts and model responses must be passed as part of the argument. This argument accepts an array of objects, where each object represents a turn and contains a 'role' ('system', 'user', or 'assistant') and a 'content' key, presented in chronological order. In Streamlit, st.session\_states is the standard and recommended mechanism for storing and managing data across reruns of the application, making it ideal for maintaining chat history. Option A is inefficient and incorrect because 'COMPLETE does not inherently manage history from external tables. Option B is incorrect as 'COMPLETE does not retain state between calls; history must be explicitly managed. Option D is less effective than structured history, as it loses the semantic role distinction and can be less



