

# Snowflake GES-C01 Practice Test Prepare for Success



BONUS!!! Download part of BraindumpsIT GES-C01 dumps for free: <https://drive.google.com/open?id=1qqFRM-ZXDpJnRYYf2CQHE3dS9uVfsQm5>

If you want to become a future professional person in this industry, getting qualified by Snowflake certification is necessary. Now, pass your GES-C01 actual exam in your first time by the help of BraindumpsIT study material. Our GES-C01 pdf torrent contains the best relevant questions and verified answers which exactly matches with the GES-C01 Actual Exam and surely helps you to pass the exam. Besides, one year free update of GES-C01 practice torrent is available after purchase.

If you are still struggling to prepare for passing Snowflake real exam at this moment, our BraindumpsIT GES-C01 vce dumps can help you preparation easier and faster. Our website can provide you Valid GES-C01 Exam Cram with high pass rate to help you get certification, and then you will become a good master of certification exam.

>> Valid GES-C01 Guide Files <<

## GES-C01 Practice Test | Braindumps GES-C01 Downloads

A considerable amount of effort goes into our products. So in most cases our GES-C01 study materials are truly your best friend. On one hand, our GES-C01 study materials are the combination of the latest knowledge and the newest technology, which could constantly inspire your interest of study. On the other hand, our GES-C01 Study Materials can predicate the exam correctly. Therefore you can handle the questions in the real exam like a cork. Through highly effective learning method and easily understanding explanation, you will pass the GES-C01 exam with no difficulty.

## Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q252-Q257):

### NEW QUESTION # 252

A data engineer is integrating a custom application with Snowflake Cortex to leverage the 'COMPLETE' function via its REST API. They are preparing a 'curl' request to send a prompt to the 'mistral-large?' model. Which of the following 'curl' command configurations correctly specifies the "mandatory" authentication header and a valid token type for accessing the Cortex REST API?

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

**Answer: C**

Explanation:

The Cortex REST API requires the 'Authorization: Bearer ' header for authentication. The token can be a JSON web token (JWT), OAuth token, or programmatic access token. Option B correctly uses 'Authorization: Bearer ' and includes the other required headers ('Content-Type' and 'Accept') for a 'COMPLETE request. Options A, C, D, and E use incorrect header names or token

types for the 'Authorization' header according to the documentation.

### NEW QUESTION # 253

A data team is refining their Cortex Analyst semantic model to improve the accuracy of responses for specific, frequently asked questions and to enable better literal value searches. Consider a semantic model being developed to address these requirements. Which two configurations or features are directly relevant and correctly applied in the semantic model YAML for these purposes?

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

**Answer: B,E**

Explanation:

Option A is correct. Cortex Search Services can be integrated into a dimension's definition (using the field with 'service' and fields) to improve literal matching by performing semantic search over the underlying column, which enhances Cortex Analyst's ability to find literal values for filtering. Option B is correct. The 'verified\_queries' section allows pre-defining accurate SQL queries for specific natural language questions. Setting 'use\_as\_onboarding\_question true' ensures these queries are used when relevant and presented as suggested questions to users. Option C is incorrect; while metrics can reference logical columns, 'relationships' between logical tables are necessary for defining joins, especially across different underlying base tables. Option D is incorrect; 'custom\_instructions' are provided at the model level to give general context to the LLM for SQL query generation, not embedded within individual dimension definitions. Option E is incorrect; the 'sample\_values' field is recommended for dimensions with relatively low-cardinality (approximately 1-10 distinct values) to aid in semantic search for literals, not for high-cardinality dimensions.

### NEW QUESTION # 254

A data application developer is building a Streamlit chat application within Snowflake. This application uses a RAG pattern to answer user questions about a knowledge base, leveraging a Cortex Search Service for retrieval and an LLM for generating responses. The developer wants to ensure responses are relevant, concise, and structured. Which of the following practices are crucial when integrating Cortex Search with Snowflake Cortex LLM functions like AI\_COMPLETE for this RAG chatbot?

- A. The
- B. Using the
- C. For performance and cost optimization, it is always recommended to query Cortex Search and the LLM function within a single
- D. To maintain conversational context in a multi-turn chat, the developer should pass all previous user prompts and model responses in the
- E. The retrieved context from Cortex Search should be directly concatenated with the user's prompt as input to the

**Answer: B,D**

Explanation:

Option A is incorrect. The user's query is typically embedded (e.g., using ' to perform a similarity search against the Cortex Search Service. The 'retrieved documents' (context) are then passed to the function, not the embedding function itself. Option B is correct because to provide a stateful, conversational experience, all previous user prompts and model responses should be passed in the array to the 'COMPLETE' or function. Option C is incorrect. While concatenation is a method, for better accuracy and control, the retrieved context should be integrated into a well-engineered prompt, often using tags or specific instructions, rather than just raw concatenation, to guide the LLM's response. Option D is correct because 'AI\_COMPLETE Structured OutputS allows you to supply a JSON schema that completion responses must follow, reducing the need for post-processing and enabling seamless integration with systems requiring deterministic responses. Option E is incorrect. While keeping processing within Snowflake is good for data governance, complex RAG pipelines often involve multiple distinct steps (query embedding, search, retrieval, LLM completion) that may benefit from a staged approach rather than a single monolithic SQL statement. The optimal approach depends on the specific complexity and performance requirements, and a single 'SELECT for the 'entire' RAG flow might not always be the most efficient or practical solution.

### NEW QUESTION # 255

A Snowflake developer, 'AI\_ENGINEER', is creating a Streamlit in Snowflake (SiS) application that will utilize a range of Snowflake Cortex LLM functions, including SNOWFLAKE.CORTEX.COMPLETE, SNOWFLAKE.CORTEX.CLASSIFY\_TEXT, and SNOWFLAKE.CORTEX.EMBED\_TEXT 768. The application also needs to access data from tables within a specific database and schema. 'AI\_ENGINEER' has created a custom role, for the application to operate under. Which of the following privileges or roles are absolutely necessary to grant to for the successful execution of these Cortex LLM functions and interaction with the specified database objects? (Select all that apply.)

- A.
- B. The CREATE COMPUTE POOL privilege to provision resources for the Streamlit application.
- C.
- D. The ACCOUNTADMIN role to ensure unrestricted access to all Snowflake Cortex features.
- E. The USAGE privilege on the specific database and schema where the Streamlit application and its underlying data tables are located.

**Answer: C,E**

Explanation:

To execute Snowflake Cortex AI functions such as 'SNOWFLAKE.CORTEX.COMPLETE', 'SNOWFLAKE.CORTEX.CLASSIFY\_TEXT', and (or their 'AL' prefixed counterparts), the role used by the application in this case) must be granted the 'SNOWFLAKE.CORTEX.USER' database role. Additionally, for the application to access any database or schema objects (like tables for data input/output or storing the Streamlit app itself), the 'USAGE' privilege must be granted on those specific database and schema objects. Option B, 'CREATE SNOWFLAKE.ML.DOCUMENT\_INTELLIGENCE', is a privilege specific to creating Document AI model builds and is not required for general Cortex LLM functions. Option D, 'ACCOUNTADMIN', grants excessive privileges and is not a best practice for application roles. Option E, 'CREATE COMPUTE POOL', is a privilege related to Snowpark Container Services for creating compute pools, which is not directly required for running a Streamlit in Snowflake application that consumes Cortex LLM functions.

### NEW QUESTION # 256

A data team is deploying a new customer service chatbot using Snowflake Cortex Analyst. To accurately forecast and optimize their costs, the team needs to understand how Cortex Analyst billing works. Which of the following statements accurately describe the cost considerations for Snowflake Cortex Analyst?

- A. Costs for Cortex Analyst are primarily driven by the number of successful messages processed, with a rate of 0.067 Credits per message.
- B. Cortex Analyst incurs compute costs based on the total number of input and output tokens processed in each conversation turn.
- C. Virtual warehouse compute is directly billed for Cortex Analyst operations, and its size can be adjusted to optimize query generation performance and cost.
- D. Cortex Analyst's Cloud Services compute costs are subject to a daily adjustment, where Snowflake only bills if these costs exceed 10% of the daily virtual warehouse cost for the account.
- E. Only successful responses (HTTP 200) from Cortex Analyst are counted towards the credit usage.

**Answer: A,E**

Explanation:

Option B is correct because Cortex Analyst credit rate usage is based on the number of messages processed, specifically 67 Credits per 1,000 messages or 0.067 Credits per message. Option E is correct as only successful responses (HTTP 200) from Cortex Analyst are counted towards the credit usage. Option A is incorrect because the number of tokens in each message does not affect the cost for Cortex Analyst; this is a billing model for other Cortex LLM functions, but not Cortex Analyst. Option C is incorrect; Cortex Analyst runs on Snowflake-managed compute resources as a Serverless Feature, meaning virtual warehouse compute is not directly billed for its operations in the same way as a user-managed warehouse. Option D is incorrect. While Cloud Services compute generally has a 10% daily warehouse cost adjustment, Serverless Features (which include Cortex Analyst) do not contribute to this adjustment.

### NEW QUESTION # 257

.....

