

# GES-C01 Valid Test Camp, GES-C01 Valid Exam Testking

---

## 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

P.S. Free & New GES-C01 dumps are available on Google Drive shared by GetValidTest: [https://drive.google.com/open?id=1omtAUOB0uapz764\\_7DMLVpeMuvUbPPA-](https://drive.google.com/open?id=1omtAUOB0uapz764_7DMLVpeMuvUbPPA-)

Under the hatchet of fast-paced development, we must always be cognizant of social long term goals and the direction of the development of science and technology. Adapt to the network society, otherwise, we will take the risk of being obsoleted. Although our GES-C01 exam dumps have been known as one of the world's leading providers of exam materials, you may be still suspicious of the content. For your convenience, we especially provide several demos for future reference and we promise not to charge you of any fee for those downloading. Therefore, we welcome you to download to try our GES-C01 Exam for a small part. Then you will know whether it is suitable for you to use our GES-C01 test questions. There are answers and questions provided to give an explicit explanation. We are sure to be at your service if you have any downloading problems.

As is known to us, getting the newest information is very important for all people to pass the exam and get the certification in the shortest time. In order to help all customers gain the newest information about the GES-C01 exam, the experts and professors from our company designed the best GES-C01 test guide. The experts will update the system every day. If there is new information about the exam, you will receive an email about the newest information about the GES-C01 Learning Materials. We can promise that you will never miss the important information about the GES-C01 exam.

>> GES-C01 Valid Test Camp <<

## Get Success in GES-C01 by Using GES-C01 Valid Test Camp

The number of questions of the GES-C01 preparation questions you have done has a great influence on your passing rate. And we update the content as well as the number of the GES-C01 exam braindumps according to the exam center. As for our GES-C01 Study Materials, we have prepared abundant exercises for you to do. You can take part in the real GES-C01 exam after you have memorized all questions and answers accurately. And we promise that you will get a 100% pass guarantee.

### Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q321-Q326):

#### NEW QUESTION # 321

A project team is preparing to deploy a Document AI solution to process scanned customer feedback forms. They have created a dedicated role, 'customer\_feedback\_processor', and successfully granted it the SNOWFLAKE.

DOCUMENT\_INTELLIGENCE\_CREATOR database role. The environment consists of 'feedback database', 'forms schema' schema, and 'ai workload warehouse'. However, when the attempts to prepare a Document AI model build in Snowsight, they encounter a 'permission denied' error. Which of the following missing 'USAGE' grants could be the direct cause of this error?

- A. 

```
GRANT USAGE ON DATABASE feedback_db TO ROLE customer_feedback_processor;
```
- B. 

```
GRANT USAGE ON SCHEMA forms_schema TO ROLE customer_feedback_processor;
```
- C. 

```
GRANT OPERATE ON WAREHOUSE ai_workload_wh TO ROLE customer_feedback_processor;
```
- D. 

```
GRANT OPERATE ON WAREHOUSE ai_workload_wh TO ROLE customer_feedback_processor;
```
- E. 

```
GRANT CREATE SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE ON SCHEMA forms_schema TO ROLE customer_feedback_processor;
```

Answer: A,B,C

Explanation:

To prepare a Document AI model build, the role requires specific privileges in addition to the 'SNOWFLAKE.DOCUMENT\_INTELLIGENCE\_CREATOR' database role. These include 'USAGE' on the database, 'USAGE' on the schema, 'USAGE' on the virtual warehouse, and 'OPERATE' on the virtual warehouse. Options A, B, and E represent the missing 'USAGE' grants that would directly lead to a 'permission denied' error when trying to prepare a Document AI model build. Option C grants 'OPERATE' on the warehouse, which is necessary but not a 'USAGE' privilege. Option D grants 'CREATE SNOWFLAKE.ML.DOCUMENT\_INTELLIGENCE', which is also a required privilege for model build creation but not a 'USAGE' grant.

#### NEW QUESTION # 322

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 retrieved context from Cortex Search should be directly concatenated with the user's prompt as input to the
- B. Using the `response_format` option within `AI_COMPLETE` or `COMPLETE`
- C. To maintain conversational context in a multi-turn chat, the developer should pass all previous user prompts and model responses in the

```
prompt_or_history
array to the
AI_COMPLETE
function for each turn.
```

- D. The SNOWFLAKE.CORTEX.EMBED\_TEXT\_768

function should be used directly within the  
AI COMPLETE

- E. For performance and cost optimization, it is always recommended to query Cortex Search and the LLM function within a single

**Answer: B,C**

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 'AI COMPLETE 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 # 323

A data engineer is configuring a Document AI pipeline to process scanned PDF invoices stored in an internal stage named 'invoice\_docs\_stage'. After uploading the PDF files, they execute an extracting query using '!PREDICT. The query consistently returns the error:

```
{ "__processingErrors": [ "File extension does not match actual mime type. Mime-Type: application/octet-stream" ] }
```

Which of the following is the most likely cause of this error?

- A. The PDF documents exceed the maximum allowed file size of 50 MB.
- B. The documents contain non-English text, which is not fully supported by Document AI for optimal results.
- C. The 'GET\_PRESIGNED\_URL' function used in the '!PREDICT query has an expired URL.
- **D. The internal stage was not created with 'SNOWFLAKE\_SSE' encryption enabled.**
- E. The Document AI model build is attempting to process more than 1000 documents in a single query.

**Answer: D**

Explanation:

The error message 'File extension does not match actual mime type. Mime-Type: application/octet-stream' is a specific error documented for DocumentAI when internal stages are not created with 'SNOWFLAKE\_SSE' encryption. For internal stages, Document AI requires server-side encryption to be enabled. Options A, C, and D would typically result in different error messages or behaviors. Option E refers to language support, which might impact accuracy but is not the cause of a file format identification error.

### NEW QUESTION # 324

A data engineering team has developed a Python-based generative AI application and instrumented its key functions using the TruLens SDK. Their next step is to register this application with Snowflake AI Observability to initiate evaluation runs and capture

application traces within Snowflake.

- The `test_app` parameter, which is an instance of the user-defined application that will be invoked for evaluation, is responsible for managing the connection to Snowflake.
- The `app_name` parameter specifies a unique identifier for the application within Snowflake, which also dictates the name of the underlying Snowflake table where traces are stored.
- The `connector` parameter, an instance of `SnowflakeConnector`, is explicitly designed to manage the Snowpark session and Snowflake database connection, facilitating the export of traces.
- The `main_method` parameter defines a mandatory entry point for the application that must be instrumented with `SpanAttributes.SpanType.RECORD_ROOT` for traces to be exported correctly.
- The `app_version` parameter, while optional, directly controls the pricing model for the AI Observability evaluation runs by setting the LLM-as-a-judge cost per token.

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

**Answer: E**

Explanation:

To register a generative AI application in Snowflake for capturing traces and evaluations, a 'TruApp' object is created. The 'connector' parameter within 'TruApp' is a 'SnowflakeConnector' instance, specifically a wrapper class that manages the Snowpark session and Snowflake database connection to export traces to Snowflake. - Option A is incorrect because 'test\_app' is an instance of the user-defined application, not responsible for managing the connection. - Option B is incorrect; is an arbitrary name for the application but the source does not state it dictates the name of an underlying table for traces. The event table contains logs, but its naming convention is not directly tied to 'app\_name' in this manner. - Option D is incorrect because 'main\_method' is optional if another method is instrumented with 'RECORD\_ROOT'. It doesn't state it's mandatory, and the responsibility for correct trace export lies with the 'connector'. - Option E is incorrect. 'app\_version' is for experiment tracking and comparison, not for controlling the pricing model for evaluation runs. LLM judge costs are based on Cortex Complete function calls.

### NEW QUESTION # 325

A development team is building an AI-powered data pipeline in Snowflake. The pipeline involves extracting text from documents, generating embeddings using

`EMBED_TEXT_768`

,and then performing similarity searches using

`VECTOR_INNER_PRODUCT`

to find related documents. They plan to manage this pipeline using Snowflake tasks and want to integrate with a Python application for some custom processing. Considering this scenario, which of the following statements about implementing this pipeline are true?

- A. `VECTOR_INNER_PRODUCT`
- B. Snowflake
- C. When using Snowflake tasks to automate the embedding generation and similarity search,
- D. If the team wants to use the Snowpark Python library to call
- E. To generate document embeddings, the

**Answer: A**

Explanation:

Option A is incorrect. Snowflake recommends executing queries that call Cortex AI SQL functions like `EMBED_TEXT_768` with a smaller warehouse (no larger than `MEDIUM`), as larger warehouses do not increase performance. Snowpark-optimized warehouses are recommended for workloads with large memory requirements or specific CPU architectures, typically for ML training, not for general Cortex AI function calls. Option B is incorrect. The Snowpark Python library explicitly states that it does not support the `VECTOR_COSINE_SIMILARITY` function, meaning it does not 'fully support all vector similarity functions'. Option C is incorrect. The `VECTOR` data type is not supported as clustering keys. Option D is correct. After generating embeddings (e.g., storing them in a `VECTOR` column like `issue_vec`), vector similarity functions can be effectively used in SQL queries with `ORDER BY` and `LIMIT` clauses to retrieve the most similar results, as demonstrated with `VECTOR_COSINE_SIMILARITY` in a RAG example. This pattern applies to `VECTOR_INNER_PRODUCT` as well. Option E is incorrect. The `VECTOR` data type is not supported for use with dynamic tables. Additionally, Snowflake Cortex functions (including `EMBED_TEXT_768`) do not support dynamic tables.



DOWNLOAD the newest GetValidTest GES-C01 PDF dumps from Cloud Storage for free: [https://drive.google.com/open?id=1omtAUOB0uapz764\\_7DMLVpeMuvUbPPA-](https://drive.google.com/open?id=1omtAUOB0uapz764_7DMLVpeMuvUbPPA-)