

# Verified GES-C01 Answers & GES-C01 Interactive Practice Exam

**GES EXAM 2 100% VERIFIED  
CORRECT ANSWERS  
2025/2026 STUDY SET**

THE TRANSFER OF WATER FROM THE LAND AND OCEAN TO THE ATMOSPHERE AND BACK AGAIN TO THE LAND AND OCEAN IN A NEVER ENDING CYCLE IS KNOWN AS \_\_\_\_\_.

HYDROLOGIC CYCLE

THE HYDROLOGIC CYCLE IS A \_\_\_\_\_ SYSTEM

CLOSED

SPECIFIC HUMIDITY

THE ACTUAL QUANTITY OF WATER HELD BY A PARCEL OF AIR (BASED ON TEMPERATURE)

DEW POINT TEMPERATURE

TEMPERATURE AT WHICH AIR WITH A GIVEN HUMIDITY WILL REACH SATURATION WHEN COOLED WITHOUT CHANGING ITS PRESSURE- JUST CHANGING THE TEMPERATURE

EVAPORATION INVOLVES THE \_\_\_\_\_ OF LATENT HEAT AND A CHANGE IN STATE FROM \_\_\_\_\_ PHASE TO THE GASEOUS PHASE

ABSORPTION; LIQUID

CONDENSATION INVOLVES THE \_\_\_\_\_ OF LATENT HEAT AND A CHANGE IN STATE FROM THE \_\_\_\_\_ PHASE TO THE LIQUID PHASE

RELEASE; GASEOUS

\_\_\_\_\_ IS WHEN ICE CHANGES DIRECTLY FROM THE SOLID PHASE TO THE GASES PHASE WITHOUT A TRANSITIONAL LIQUID PHASE

SUBLIMATION

GES-C01 preparation materials will be the good helper for your qualification certification. We are concentrating on providing high-quality authorized GES-C01 study guide all over the world so that you can clear exam one time. GES-C01 reliable exam bootcamp materials contain three formats: PDF version, Soft test engine and APP test engine so that our products are enough to satisfy different candidates' habits and cover nearly full questions & answers of the real GES-C01 test.

As the unprecedented intensity of talents comes in great numbers, what abilities should a talent of modern time possess and finally walk to the success? Well, of course it is GES-C01 exam qualification certification that gives you capital of standing in society. Our GES-C01 preparation materials display a brand-new learning model and a comprehensive knowledge structure on our official exam bank, which aims at improving your technical skills and creating your value to your future. You will be bound to pass the GES-C01 Exam with our advanced GES-C01 exam questions.

>> **Verified GES-C01 Answers <<**

## Download Real Snowflake GES-C01 Exam Questions And Start Your Preparation

DumpStillValid also offers the GES-C01 web-based practice exam with the same characteristics as desktop simulation software but with minor differences. It is online Snowflake Certification Exam which is accessible from any location with an active internet

connection. This SnowPro® Specialty: Gen AI Certification Exam GES-C01 Practice Exam not only works on Windows but also on Linux, Mac, Android, and iOS. Additionally, you can attempt the OMG GES-C01 practice test through these browsers: Opera, Safari, Firefox, Chrome, MS Edge, and Internet Explorer.

## **Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q89-Q94):**

### **NEW QUESTION # 89**

A developer is building a client application that interacts with a Snowflake Cortex Agent using its REST API. They are implementing multi-turn conversation support. Which of the following is the most critical aspect for maintaining conversational context over multiple API calls?

- A. The REST API automatically identifies and reuses context from previous requests based on the Authorization token, eliminating the need to pass full history.
- **B. The client application must include all prior user prompts and assistant/analyst responses in a messages array in each new API request.**
- C. Only the last user prompt and the immediately preceding agent response should be sent to conserve token usage.
- D. The agent internally summarizes long conversations, and the client receives only a condensed summary\_token to pass for subsequent turns.
- E. A specific session\_id parameter, generated at the start of the conversation, must be sent with each request to link turns.

### **Answer: B**

Explanation:

Option A is the most critical and correct method for maintaining conversational state in multi-turn interactions with Cortex LLM functions, including those used by Cortex Analyst and by extension Cortex Agents. The underlying COMPLETE function (and its REST API equivalent) does not retain state from one call to the next; therefore, the client application must explicitly pass all previous user prompts and model (assistant/analyst) responses in chronological order within the 'messages' (or array for each new request to provide a stateful experience. Option B is incorrect as Cortex LLM functions do not automatically retain state across calls. Option C is incorrect; there is no documented 'session\_id' parameter for implicitly managing conversation history for these APIs. Option D is incorrect; passing only partial history would lead to a loss of full conversational context and degrade the quality of follow-up responses. Option E is incorrect; while Cortex Analyst internally uses a summarization agent for long conversations to reframe questions, the client application is still responsible for managing and sending the full conversation history in the 'messages' array via the API.

### **NEW QUESTION # 90**

A data analytics team is building a Retrieval Augmented Generation (RAG) application to provide contextual answers from a vast repository of internal documents stored in Snowflake. They are evaluating different strategies for generating and retrieving text embeddings to optimize the overall RAG pipeline's performance and relevance. Which of the following statements accurately describe performance considerations related to embedding generation and retrieval in this RAG context? (Select all that apply)

Using the snowflake-arctic-embed-1-v2.0-8k model with EMBED\_TEXT\_1024 and longer text chunks (e.g., 2000 tokens) will generally lead to higher retrieval quality due to the model's larger context window.

Deploying a custom Hugging Face embedding model (e.g., SentenceTransformer) on a Snowpark Container Services (SPCS) compute pool with GPUs offers potential for lower latency and higher throughput for embedding generation compared to managed EMBED\_TEXT functions for very high-volume scenarios.

For optimal search results with Cortex Search, Snowflake recommends splitting the text in the search column into chunks of no more than 512 tokens, as this typically results in higher retrieval and downstream LLM response quality.

Executing queries that call managed EMBED\_TEXT functions on a large Snowflake warehouse (e.g., 'X-Large') will significantly improve embedding generation performance compared to a 'Medium' warehouse.

Relying solely on keyword search for document retrieval within Cortex Search will generally provide better contextual relevance for LLM responses than hybrid (vector and keyword) search, resulting in lower latency.

- **A. Option C**

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

**Answer: A,D**

Explanation:

For optimizing RAG pipeline performance and relevance: ' This statement is incorrect. Snowflake's documentation explicitly recommends splitting text into smaller chunks (no more than 512 tokens) for Cortex Search to achieve optimal retrieval and downstream LLM response quality. This holds true even with models that have larger context windows like 'snowflake-arctic-embed-l-v2.0-8k' , because smaller chunks lead to more precise retrieval. \* \*\*\*B:\*\* Deploying custom models like a Hugging Face 'sentenceTransformeN on Snowpark Container Services (SPCS) with GPU compute pools (e.g., \*GPU or \*GPU NV\_M') is optimized for intensive GPU usage scenarios like LLMs/VLMs. This can provide lower latency and higher throughput for embedding generation in very high-volume, custom scenarios, offering more control than managed functions. ' This statement is correct.

Snowflake's documentation clearly states that for best search results with Cortex Search, it is recommended to split the text in the search column into chunks of no more than 512 tokens. This strategy typically results in higher retrieval and better quality responses from downstream LLMs. \* This statement is incorrect. Snowflake explicitly advises executing queries that call Cortex AI SQL functions (including ' EMBED\_TEXT) with a \*smaller\* warehouse (no larger than MEDIUM), as larger warehouses do not increase performance for these specific functions. \* \*E:" This statement is incorrect. Cortex Search powers RAG applications by leveraging \*semantic search\*, which combines both vector and keyword search capabilities, to provide customized, contextualized responses. Relying solely on keyword search would generally yield less contextual relevance for LLM responses than a hybrid approach.

#### NEW QUESTION # 91

A new data analyst, 'DI\_ANALYST', is setting up their Snowflake environment to create and test Document AI model builds through the Snowsight user interface. The database 'document\_ai\_db' and schema 'analysis\_schema' have already been created for this purpose, along with a dedicated virtual warehouse 'di\_compute\_wh'. The 'DI\_ANALYST' role has been granted the SNOWFLAKE.DOCUMENT\_INTELLIGENCE\_CREATOR

database role. Which of the following SQL commands grant the \*strictly necessary USAGE privileges\* required for 'DI\_ANALYST' to successfully initiate a new Document AI model build in Snowsight?

- A. `GRANT OPERATE ON WAREHOUSE di_compute_wh TO ROLE DI_ANALYST;`
- B. `GRANT USAGE ON SCHEMA analysis_schema TO ROLE DI_ANALYST;`
- C. `GRANT USAGE ON WAREHOUSE di_compute_wh TO ROLE DI_ANALYST;`
- D. `GRANT USAGE ON DATABASE document_ai_db TO ROLE DI_ANALYST;`
- E. `GRANT CREATE STAGE ON SCHEMA analysis_schema TO ROLE DI_ANALYST;`

**Answer: B,C,D**

Explanation:

To prepare a Document AI model build, the role must be granted the database role. In addition, it requires 'USAGE' on the database, 'USAGE' on the schema, 'USAGE' on the warehouse, and 'OPERATE' on the warehouse. Options A, B, and D correctly grant the necessary 'USAGE' privileges. Option C grants 'OPERATE' on the warehouse, which is also a necessary privilege but is not a 'USAGE' privilege as specifically asked in the question. Option E grants 'CREATE STAGES', which is required for creating processing pipelines but not strictly for initiating a new model build itself.

#### NEW QUESTION # 92

A data engineering team is tasked with improving the accuracy of a Cortex Analyst solution for a large e-commerce product catalog. Users frequently ask natural language questions involving specific product names, brands, and categories. The team observes that Cortex Analyst sometimes struggles to identify and correctly filter by these literal values in the generated SQL. Which of the following configurations or approaches, within the semantic model, can effectively enhance Cortex Analyst's ability to precisely identify and use literal values for filtering, based on Snowflake's best practices?

For a `product_category` dimension with less than 10 distinct values, setting `is_enum: true` and providing an exhaustive list of `sample_values` in the semantic model YAML.

For a `product_name` dimension, configuring a `cortex_search_service` entry within the dimension, including both the `service` name and the `literal_column` that the Cortex Search Service is indexing.

Using `LIKE` clauses directly within the dimension's `expr` field in the semantic model to enable fuzzy string matching for literal values.

Relying exclusively on the `verified_queries` section of the semantic model to define all possible literal search scenarios with pre-written SQL.

Increasing the `max_tokens` parameter in the Cortex Analyst API request options to allow the underlying LLM to generate more extensive literal value lists.

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

**Answer: A,D**

**Explanation:**

Options A and B are correct. For dimensions with low cardinality (around 1-10 distinct values), setting 'is\_enum: true' and providing an exhaustive 'sample\_values' list ensures Cortex Analyst chooses only from that predefined list, improving literal usage. For higher cardinality dimensions, integrating a Cortex Search Service via the entry, specifying both the 'service' name and the , allows semantic search over the underlying data to find appropriate literal values. Option C is incorrect because Cortex Analyst leverages semantic similarity search or Cortex Search for literal values, not direct 'LIKE' clauses in the 'expr' field. Option D is incorrect because while 'verified\_queries' improve accuracy for specific, known questions, they are not a scalable solution for all possible literal search scenarios and are not the primary mechanism for improving general literal value identification. Option E is incorrect because the 'max\_tokens' parameter controls the length of the LLM's output response, not its ability to identify or filter by literal values.

### NEW QUESTION # 93

A data engineering team is implementing a Document AI solution to automate the extraction of vendor invoice details. They have already published a Document AI model build, 'vendor\_invoice\_model', located in New invoices are uploaded to an internal stage, 'fin\_db.invoice\_processing\_schema.raw\_invoices\_stage'. The 'invoice\_pipeline role' is responsible for executing the

`!PREDICT`

method on these staged documents as part of an automated task. Which of the following USAGE privileges are essential for to successfully execute the

`!PREDICT`

method and process documents from the specified stage, assuming the required

`SNOWFLAKE.DOCUMENT_INTELLIGENCE_CREATOR` role is already granted?

- A. 

```
GRANT READ ON STAGE fin_db.invoice_processing_schema.raw_invoices_stage TO ROLE invoice_pipeline_role;
```
- B. 

```
GRANT SELECT ON TABLE fin_db.invoice_processing_schema.raw_invoices_stage TO ROLE invoice_pipeline_role;
```
- C. 

```
GRANT USAGE ON DATABASE fin_db TO ROLE invoice_pipeline_role;
```
- D. 

```
GRANT USAGE ON WAREHOUSE invoice_compute_wh TO ROLE invoice_pipeline_role;
```
- E. 

```
GRANT USAGE ON SCHEMA fin_db.invoice_processing_schema TO ROLE invoice_pipeline_role;
```

**Answer: C,D,E**

**Explanation:**

For a role to execute Document AI's '`!PREDICT`' method, even if the '`SNOWFLAKE.DOCUMENT_INTELLIGENCE_CREATOR`' database role is granted, it still requires 'USAGE' privileges on the database, schema, and the virtual warehouse used for computation. Options A, B, and C provide these necessary 'USAGE' grants. Option D grants 'READ' on the stage, which is typically needed to read files from a stage (e.g., via '`GET_PRESIGNED_URL`' as used by '`!PREDICT`') but the question specifically asks for 'USAGE' privileges as defined in the Document AI setup documentation, which explicitly lists 'USAGE' on DB, Schema, and Warehouse. Option E refers to 'SELECT ON TABLE', which is not applicable to a stage as a table.

## NEW QUESTION # 94

After passing the Snowflake GES-C01 certification exam, you can take advantage of a number of extra benefits. With the correct concentration, commitment, and GES-C01 exam preparation, you could ace this SnowPro® Specialty: Gen AI Certification Exam GES-C01 test with ease. DumpStillValid is a trusted and leading platform that is committed to preparing the Snowflake GES-C01 exam candidates in a short time period.

GES-C01 Interactive Practice Exam: <https://www.dumpstillvalid.com/GES-C01-prep4sure-review.html>

Choosing DumpStillValid as the GES-C01 exam preparation assistance will be a great help for passing the Aruba Snowflake Snowflake Certification GES-C01 exam. Advantages of Software-Based Network Simulation DumpStillValid GES-C01 Interactive Practice Exam DumpStillValid GES-C01 Interactive Practice Exam is a versatile tool and valuable asset for self-paced learning, classroom instruction and corporate training. Generally, the average person will think the more the better, for example, the more questions the GES-C01 sure exam dumps contain, the better result they will get.

welcome to give us any suggestions, Moving Toward GES-C01 Forging a Political Consensus on Reducing Oil Import Dependency, Choosing DumpStillValid as the GES-C01 Exam Preparation assistance will be a great help for passing the Aruba Snowflake Snowflake Certification GES-C01 exam.

## 2025 Authoritative GES-C01: Verified SnowPro® Specialty: Gen AI Certification Exam Answers

Advantages of Software-Based Network Simulation DumpStillValid DumpStillValid GES-C01 Interactive Practice Exam is a versatile tool and valuable asset for self-paced learning, classroom instruction and corporate training.

Generally, the average person will think the more the better, for example, the more questions the GES-C01 sure exam dumps contain, the better result they will get.

For these latest GES-C01 dumps, you need to visit DumpStillValid, Snowflake Certification logo What's on the GES-C01 Snowflake Solutions Exam?

myportal.utt.edu.tt, ncon.edu.sa, Disposable vapes