


# Test GES-C01 Testking, GES-C01 Dump Check



If you are sure you have learnt all the GES-C01 exam questions, you have every reason to believe it. Itbraindumps's GES-C01 exam dumps have the best track record of awarding exam success and a number of candidates have already obtained their targeted GES-C01 Certification relying on them. They provide you the real exam scenario and by doing them repeatedly you enhance your confidence to GES-C01 questions answers without any hesitation.

It is acknowledged that there are numerous GES-C01 learning questions for candidates for the GES-C01 exam, however, it is impossible for you to summarize all of the key points in so many materials by yourself. But since you have clicked into this website for GES-C01 practice materials you need not to worry about that at all because our company is especially here for you to solve this problem. We have a lot of regular customers for a long-term cooperation now since they have understood how useful and effective our GES-C01 Actual Exam is.

>> Test GES-C01 Testking <<

## Test GES-C01 Testking | 100% Free Accurate SnowPro® Specialty: Gen AI Certification Exam Dump Check

One of the biggest highlights of the GES-C01 exam materials is the availability of three versions: PDF, app/online, and software/pc, each with its own advantages: The PDF version of GES-C01 exam materials has a free demo available for download. You can print exam materials out and read it just like you read a paper. The online version of GES-C01 Exam Materials is based on web browser usage design and can be used by any browser device. At the same time, the first time it is opened on the Internet, it can be used offline next time. You can practice anytime, anywhere.

## Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q132-Q137):

### NEW QUESTION # 132

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. The

SNOWFLAKE.CORTEX.EMBED\_TEXT\_768

function should be used directly within the  
AI\_COMPLETE

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

**Answer: B,E**

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 # 133

A data engineering team is building an automated pipeline in Snowflake to process customer reviews. They need to use AI\_COMPLETE to extract specific details like product, sentiment, and issue type, and store them in a strictly defined JSON format for seamless downstream integration. They aim to maximize the accuracy of the structured output and manage potential model limitations. Which statements accurately reflect the best practices and characteristics when using AI\_COMPLETE with structured outputs for this scenario?

- ☐ Specifying the response\_format argument in AI\_COMPLETE incurs an additional, separate compute cost for the overhead of verifying each token against the provided JSON schema.
- ☐ For complex data extraction tasks requiring high accuracy, it is recommended to explicitly include phrases like 'Respond in JSON' within the prompt, in addition to defining the response\_format schema.
- ☐ When using OpenAI (GPT) models with structured outputs, the JSON schema defined in response\_format must explicitly include "additionalProperties": false in every node and ensure the "required" field lists all properties in the schema.
- ☐ If the model's generated response slightly deviates from the specified JSON schema, AI\_COMPLETE with structured outputs will automatically reformat the output to conform, preventing validation errors.
- ☐ To achieve the most consistent and deterministic results when using AI\_COMPLETE for structured outputs, the temperature option should be set to 0.

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

**Answer: A,C,D**

Explanation:

Option A is incorrect because Structured Outputs do not incur additional compute cost for the overhead of verifying each AI\_COMPLETE token against the supplied JSON schema, though the number of tokens processed (and thus billed) can increase with schema complexity. Option B is correct because for complex reasoning tasks, it is recommended to use the most powerful models and explicitly add 'Respond in JSON' to the prompt to optimize accuracy. Option C is correct as for OpenAI (GPT) models, the schema has specific requirements: response\_format must be set to in every node, and the field must include the names of every property in the schema. Option D additionalProperties false required is incorrect because verifies each generated token against the JSON schema to ensure conformity, and if the model cannot generate a AI\_COMPLETE response that matches the schema, it will result in a validation error. Option E is correct as setting the option to e is recommended for temperature the most consistent results, regardless of the task or model, especially for structured outputs.

#### NEW QUESTION # 134

A data engineering team is tasked with creating vector embeddings for a collection of diverse, multilingual research papers for a semantic search application. They need to use 'SNOWFLAKE.CORTEX.EMBED TEXT 1024' and are considering two models: 'snowflake-arctic-embed-l-v2.0' and 'voyage-multilingual-2'. They also need to ensure the resulting embeddings are stored correctly and understand potential text truncation. Which of the following statements correctly describes the application of the 'EMBED TEXT 1024' function for these models and the characteristics of the generated embeddings?

- A. The 'snowflake-arctic-embed-l-v2.0' model has a larger context window than 'voyage-multilingual-2', making it more suitable for longer research papers without truncation.
- B. Both 'snowflake-arctic-embed-l-v2.0' and 'nv-embed-qa-4' are multilingual models with a 1024-dimension output, making them interchangeable for this use case.
- C. Using either 'snowflake-arctic-embed-l-v2.0' or 'voyage-multilingual-2' will result in a 'VECTOR(FLOAT, 1024)' data type, which is compatible with 'VARIANT' columns.
- D. The function call

```
SNOWFLAKE.CORTEX.EMBED_TEXT_1024(research_paper_text, 'snowflake-arctic-embed-l-v2.0')
```

- E. The query

```
SELECT SNOWFLAKE.CORTEX.EMBED_TEXT_1024('voyage-multilingual-2', research_paper_text) FROM research_papers;
```

**Answer: E**

Explanation:

Option C is correct. The 'EMBED\_TEXT\_1024' function's syntax is 'SNOWFLAKE.CORTEX.EMBED\_TEXT 1024(, y. The 'voyage-multilingual-2' model has a context window of 32000 tokens, and text exceeding the model's context window is truncated before embedding. Option A is incorrect because 'snowflake-arctic-embed-l-v2.0' has a 512-token context window, while 'voyage-multilingual-2' has a 32000-token context window, meaning 'voyage-multilingual-2' has a larger context window. Option B is incorrect because while 'EMBED TEXT 1024' returns a 1024) data type, the 'VECTOR' data type is explicitly not supported in 'VARIANT' columns. Option D is incorrect because 'nv-embed-qa-v' is an English-only model, whereas 'snowflake-arctic-embed-l-v2.0' is multilingual. Option E is incorrect because the model name should be the first argument and the text the second argument in the function call, i.e., 'text\_input'.

#### NEW QUESTION # 135

An AI development team is deploying a new Cortex Agent and needs to ensure optimal performance and adherence to Snowflake's Gen AI principles regarding data governance. Which of the following statements accurately reflect the models supported by Cortex Agents and the data governance considerations?

- ☐ Cortex Agents exclusively use proprietary Snowflake models, such as Snowflake Arctic, ensuring all data remains within Snowflake's governance boundary without exception.
- ☐ Supported LLMs for orchestration by Cortex Agents include llama3.1-70b, mistral-large2, and claude-3.5-sonnet. Cross-region inference may be required if the model is not available in the local region.
- ☐ Customer Data, including inputs and outputs from Cortex Agents, is not used to train or fine-tune models made available to others, and fine-tuned models built with customer data are exclusively for that customer's use.
- ☐ Cortex Agents support all models available through the AI\_COMPLETE function, including those that require explicit opt-in for Azure OpenAI GPT models, with data always staying within Snowflake's governance boundary.
- ☐ To use Cortex Agents, a role must be granted the SNOWFLAKE.CORTEX\_AGENT\_USER database role, which provides access specifically to the Agents feature.

- A. Option E
- B. Option A

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

**Answer: A,C,D**

**Explanation:**

Supported models for orchestration by Cortex Agents include 'llama3.1-70b', 'llama3.3-70b', 'mistral-large2', 'claude-3-5-sonnet', 'claude-3-7-sonnet', 'claude-4-0-sonnet', and 'openai-gpt-5'. Cross-region inference is used if a model is not available in the local region. This makes option B correct. Snowflake's Gen AI principles explicitly state that Customer Data, including inputs and outputs, is NOT used to train, re-train, or fine-tune Models made available to others, and fine-tuned models built using your data are available exclusively for your use. This makes option C correct. To make a request to Cortex Agent via the 'agent:run' API, a role must be granted the 'SNOWFLAKE.CORTEX\_USER' or 'SNOWFLAKE.CORTEX\_AGENT\_USER' role, where 'CORTEX\_AGENT\_USER' specifically provides access to the Agents feature. This makes option E correct. Option A is incorrect because Cortex Agents support a variety of models, not just proprietary Snowflake models, and while data generally stays within Snowflake's governance boundary, there are legacy exceptions for Azure OpenAI models with Cortex Analyst. Option D is incorrect because Cortex Agents support a specific list of models, not all 'AI\_COMPLETE' models, and while Snowflake-hosted LLMs keep data within the boundary, opting for Azure OpenAI models (legacy path) means metadata and prompts are transmitted outside.

### NEW QUESTION # 136

A Snowflake administrator is tasked with monitoring the efficiency and cost-effectiveness of their Cortex Analyst deployments. They need to identify if certain semantic models are generating a high volume of failed or expensive queries. Which of the following approaches or statements are crucial for effectively monitoring and identifying issues with Cortex Analyst usage and associated costs?

Querying the `SNOWFLAKE.LOCAL.CORTEX_ANALYST_REQUESTS` function to access detailed logs, including generated SQL and error/warning messages, for specific semantic models.

}

Analyzing the `CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY` view to track pages processed and credits used, as it aggregates all AI Services activity including Cortex Analyst.

}

Inspecting the `METERING_DAILY_HISTORY` view with a `SERVICE_TYPE` filter of 'AI\_SERVICES' to get a daily aggregate of all AI service credit consumption, including Cortex Analyst.

}

Reviewing the 'prompt\_tokens' and 'completion\_tokens' columns in the query logs to directly assess the token count-based cost incurred by each Cortex Analyst message.

}

Understanding that there is typically a 1-2 minute lag between a Cortex Analyst request being made and its logs becoming visible in the event table.

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

**Answer: B,C,D**

**Explanation:**

Option A is correct because Cortex Analyst logs requests to an event table, and the function can be used to query these logs, which include generated SQL and errors/warnings, helping identify issues. Option C is correct as the 'METERING\_DAILY\_HISTORY' view tracks daily credit usage for services, including 'AI\_SERVICES', which encompasses Cortex Analyst. Option E is correct as there is a reported 1-2 minute lag for Cortex Analyst requests to become visible in the event table logs. Option B is incorrect because 'CORTEX\_DOCUMENT\_PROCESSING\_USAGE\_HISTORY' specifically displays Document AI processing function activity, not all AI Services or Cortex Analyst. Option D is incorrect because Cortex Analyst costs are based on messages, not tokens, so 'prompt\_tokens' and 'completion\_tokens' would not be relevant for direct cost assessment in this context.

### NEW QUESTION # 137

.....

Passing the GES-C01 Exam is a challenging task, but with Itbraindumps Snowflake Practice Test engine, you can prepare yourself for success in one go. The GES-C01 online practice test engine offers an interactive learning experience and includes Snowflake



**GES-C01 Dump Check:** [https://www.itbraindumps.com/GES-C01\\_exam.html](https://www.itbraindumps.com/GES-C01_exam.html)

Digital Voice Signals, Transformers: a necessary evil to allow GES-C01 components to talk to one another, They can simulate the SnowPro® Specialty: Gen AI Certification Exam actual test to feel the real exam in advance.

## Hot Test GES-C01 Testking Offers you Professional Actual Snowflake SnowPro® Specialty: Gen AI Certification Exam Exam Products

- 100% Pass 2025 Snowflake Pass-Sure Test GES-C01 Testking ☐ Immediately open > [www.torrentvalid.com](http://www.torrentvalid.com) < and search for ➤ GES-C01 ☐ to obtain a free download ☐Exam GES-C01 Success
- Latest GES-C01 Exam Practice ☐ Exam GES-C01 Success ☐ New GES-C01 Exam Online ☐ Download [ GES-C01 ] for free by simply searching on [ [www.pdfvce.com](http://www.pdfvce.com) ] ☐GES-C01 Pass Leader Dumps
- 100% Pass Quiz Trustable Snowflake - GES-C01 - Test SnowPro® Specialty: Gen AI Certification Exam Testking ☐ Enter ➤ [www.examcollectionpass.com](http://www.examcollectionpass.com) < and search for ➤ GES-C01 < to download for free ☐Valid GES-C01 Test Sims
- Free PDF Fantastic Snowflake - Test GES-C01 Testking ☐ Search on ✓ [www.pdfvce.com](http://www.pdfvce.com) ☐✓☐ for ⇒ GES-C01 ⇌ to obtain exam materials for free download ☐New GES-C01 Exam Online
- GES-C01 Valid Exam Prep ☐ GES-C01 Exam Objectives ☐ Valid GES-C01 Real Test ☐ The page for free download of ➤ GES-C01 < on ✓ [www.examcollectionpass.com](http://www.examcollectionpass.com) ☐✓☐ will open immediately ☐New GES-C01 Exam Online
- GES-C01 Examcollection Dumps ☐ GES-C01 Examcollection Dumps ☐ Test GES-C01 Vce Free ☐ The page for free download of ✓ GES-C01 ☐✓☐ on ➤ [www.pdfvce.com](http://www.pdfvce.com) ☐ will open immediately ☐GES-C01 Valid Exam Forum
- 100% Pass 2025 Snowflake Pass-Sure Test GES-C01 Testking ☐ Go to website [ [www.examsreviews.com](http://www.examsreviews.com) ] open and search for ( GES-C01 ) to download for free ☐GES-C01 Braindump Free
- 2025 Pass-Sure Test GES-C01 Testking| SnowPro® Specialty: Gen AI Certification Exam 100% Free Dump Check ☐ ➤ [www.pdfvce.com](http://www.pdfvce.com) ☐☐☐ is best website to obtain ➤ GES-C01 ☐ for free download ↘ Valid GES-C01 Test Dumps
- Pass Guaranteed 2025 High Pass-Rate Snowflake Test GES-C01 Testking ☐ Easily obtain ☐ GES-C01 ☐ for free download through ➤ [www.torrentvce.com](http://www.torrentvce.com) ☐ ☐New GES-C01 Cram Materials
- Free PDF Newest Snowflake - GES-C01 - Test SnowPro® Specialty: Gen AI Certification Exam Testking ☐ Go to website ➤ [www.pdfvce.com](http://www.pdfvce.com) < open and search for ✨ GES-C01 ☐ ✨☐ to download for free ☐New GES-C01 Exam Online
- GES-C01 Exam Objectives ☐ GES-C01 Pass Leader Dumps ☐ New GES-C01 Exam Online ☐ ➤ [www.free4dump.com](http://www.free4dump.com) ☐ is best website to obtain ☐ GES-C01 ☐ for free download ☐Valid GES-C01 Real Test
- [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [study.stcs.edu.np](http://study.stcs.edu.np), [shaxianxiaochi.gogreen.top](http://shaxianxiaochi.gogreen.top), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [www.goodgua.com](http://www.goodgua.com), Disposable vapes