


Exam GES-C01 Tips, New GES-C01 Test Simulator



Our GES-C01 PDF format is also an effective format to do test preparation. In your spare time, you can easily use the GES-C01 dumps PDF file for study or revision. The PDF file of Snowflake GES-C01 real questions is convenient and manageable. These Snowflake GES-C01 Questions are also printable, giving you the option of paper study since some Snowflake GES-C01 applicants prefer off-screen preparation rather than on a screen.

This free SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam questions demo download facility is available in all three GES-C01 exam dumps formats. Just choose the best SurePassExams SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam demo questions format and download it quickly. If you think that GES-C01 exam dumps can work for you then take your buying decision.

>> Exam GES-C01 Tips <<

GES-C01 Practice Test Training Materials - GES-C01 Test Prep - SurePassExams

Our goal is to help you save both time and money by providing you with the GES-C01 updated exam questions. Keep up the good work on preparing for the Snowflake GES-C01 test with our actual Snowflake GES-C01 Dumps. We are so confident that you will succeed on the first try that we will return your money according to the terms and conditions if you do not.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q128-Q133):

NEW QUESTION # 128

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?

- ▷ The developer should store the entire conversation history in a temporary table in Snowflake and query it with each new turn, passing only the latest user message to the COMPLETE function.
- ▷ Snowflake automatically manages conversational context for COMPLETE within the session, so the developer only needs to pass the current user prompt as a string.
- ▷ The conversation history must be explicitly managed within the Streamlit application's state, typically by initializing `st.session_state.messages = []` and appending each user and assistant message as an object with 'role' and 'content' keys, then passing the full list to the `prompt_or_history` argument of COMPLETE.
- ▷ The developer should concatenate all previous user prompts and assistant responses into a single, long string, and pass this as the `<prompt>` argument to COMPLETE for each turn.
- ▷ The COMPLETE function has an optional 'conversation id' parameter that automatically retrieves and manages conversation history when provided.

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

Answer: C

NEW QUESTION # 129

An AI developer is testing a new RAG application in Snowflake. The application uses

`SNOWFLAKE.CORTEX.EMBED_TEXT_168`

to create embeddings for documents and then

`SNOWFLAKE.CORTEX.TRY_COMPLETE`

with a retrieval-augmented prompt to answer user questions. The developer observes that sometime

returns

NULL

even for seemingly valid prompts. Which of the following, if true, could explain a

NULL

return from

`TRY_COMPLETE`

in this scenario?

☐ The model specified in the `TRY_COMPLETE` call, such as `claude-3-5-sonnet`, is restricted from being used by the current role due to an active `CORTEX_MODELS_ALLOWLIST` parameter.

☐ The combined token count of the input prompt (including retrieved context) and the desired output exceeds the context window of the specified LLM model.

☐ The `temperature` parameter in the `options` argument was set to 0, making the model's output too deterministic to provide a response.

☐ The `prompt_or_history` argument was provided as a simple string instead of an array of role-content objects, which `TRY_COMPLETE` requires for all calls.

☐ The virtual warehouse executing the query was an X-Small size, which is insufficient for any `TRY_COMPLETE` operations.

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

Answer: C,E

Explanation:

Option A is correct. If access to a specific model is restricted by the `CORTEX_MODELS_ALLOWLIST`

parameter, the LLM operation would fail, and `TRY_COMPLETE` would return `NULL`

instead of an error. Option B is correct. LLMs have context window limitations, meaning the combined length of input tokens (prompt + retrieved context) and output tokens cannot exceed a certain limit for a given model (e.g., `claude-3-5-sonnet`

has a 200,000 token context window). If this limit is exceeded, the operation cannot be performed, and `TRY_COMPLETE`

would return `NULL`.

Option C is incorrect; setting 'temperature' to 0 makes the output more deterministic and focused, which generally improves adherence to instructions and is recommended for consistent results, not causes failures. Option D is incorrect. The `prompt_or_history`

argument can be a simple string if the `options`

argument is not present, or an array of objects if `options`

is present. The question implies `options`

might be used for RAG, but it does not mandate an array for all calls. Option E is incorrect. Snowflake recommends executing Cortex AI function queries, including

`TRY_COMPLETE`, with smaller warehouses (no larger than MEDIUM), meaning an X-Small warehouse is generally sufficient for executing the function itself, though it might impact performance for very large datasets if used for other operations.

NEW QUESTION # 130

A data engineering team is building an automated pipeline within Snowflake to process newly ingested documents. This pipeline needs to classify each document's sentiment (positive, neutral, negative) and summarise its content using Cortex LLM functions, then store the results in a table. The pipeline is orchestrated using Streams and Tasks. Which considerations are paramount for implementing and monitoring this AI-infused data pipeline?

- ☐ The `SENTIMENT (SNOWFLAKE.CORTEX)` and `SUMMARIZE (SNOWFLAKE.CORTEX)` functions can be directly invoked within SQL statements as part of the pipeline's transformation logic for each new document from the stream.
- ☐ To handle potential LLM processing failures gracefully without stopping the entire pipeline, the team should embed `TRY_COMPLETE` calls within a User-Defined Function (UDF) that is then called by the Snowflake Task.
- ☐ Effective prompt engineering, such as defining a 'financial analyst' persona or encouraging an 'inner monologue' for the LLM, is critical to maximise the accuracy of sentiment classification and summarization outputs.
- ☐ To monitor the credit consumption specifically for these LLM function calls within the pipeline, the team should query the `SNOWFLAKE.ORGANIZATION_USAGE.METERING_DAILY_HISTORY` view, filtering by `SERVICE_TYPE = 'AI_SERVICES'`.
- ☐ The `SENTIMENT` and `SUMMARIZE` functions only bill for input tokens, making them highly cost-effective regardless of the output length generated.

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

Answer: C,D,E

Explanation:

Option A is correct. Cortex LLM functions like `SENTIMENT (SNOWFLAKE.CORTEX)` and `SUMMARIZE (SNOWFLAKE.CORTEX)` (or their updated `AI_SENTIMENT` and `AI_SUMMARIZE_AGG` versions) are SQL functions that can be directly called within SQL statements to process text from a table. This aligns with building AI-infused pipelines directly in Snowflake using SQL. Option B is correct. Using `TRY_COMPLETE` (or its older version `TRY_COMPLETE (SNOWFLAKE.CORTEX)`) is a helper function that returns `NULL` instead of raising an error if the LLM operation cannot be performed. Embedding this in a UDF or direct task logic helps in creating robust pipelines that can handle individual document processing failures without stopping the entire task. Option C is correct. Prompt engineering principles, such as defining a persona, clearly outlining tasks, and encouraging 'inner monologue,' are crucial for maximizing the effectiveness and accuracy of AI models, especially in tasks requiring nuance like sentiment analysis and summarization. Option D is incorrect. While `METERING_DAILY_HISTORY` does show `AI_SERVICES` usage, for granular usage information grouped by specific LLM functions and models, the `SNOWFLAKE.ACCOUNT_USAGE.CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY` view is more appropriate. Option E is incorrect. For functions that generate new text in the response, such as `AI_COMPLETE`, `AI_CLASSIFY`, `AI_FILTER`, `AI_AGG`, `AI_SUMMARIZE`, and `TRANSLATE` (and their previous versions like `SENTIMENT` and `SUMMARIZE`), both input and output tokens are billable.

NEW QUESTION # 131

A data engineering team is designing a pipeline in Snowflake to translate a continuous stream of multi-language customer support tickets into English using `SNOWFLAKE.CORTEX.TRANSLATE`. They are concerned about potential language identification issues and the overall cost implications. Which of the following statements are true regarding the use of `SNOWFLAKE.CORTEX.TRANSLATE` for this scenario? (Select all that apply)

- A. Snowflake Cortex functions, including `TRANSLATE`, add an internal prompt to the user's input text, which increases the total input token count for billing purposes beyond the raw text length.
- B. For cost efficiency, Snowflake recommends using a larger warehouse (e.g., XL or 2XL) for executing queries that call `TRANSLATE` functions, as this significantly reduces the per-token processing cost.
- C. The fixed billing rate for the `TRANSLATE` function is 1.50 Credits per one million Tokens processed.
- D. The `TRANSLATE` function is exclusively billed based on the number of input tokens, as it primarily analyzes existing text rather than generating new content.
- E. If the source language of a ticket is unknown or contains mixed languages (e.g., 'Spanglish'), the function can still process it by specifying an empty string for the `source_language` argument.

Answer: A,C,E

Explanation:

Option A is correct because `SNOWFLAKE.CORTEX.TRANSLATE` can handle mixed-language input or unknown source languages by specifying an empty string for the `'source_language'` argument. Option B is incorrect; for functions that generate new text, such as `TRANSLATE`, both input and output tokens are billable. Option C is incorrect; Snowflake recommends executing queries that call Cortex AISQL functions with a smaller warehouse (no larger than MEDIUM), as larger warehouses do not increase performance. Option D is correct, as `TRANSLATE` (among other Cortex functions) adds an internal prompt to the input text,

resulting in a higher billed input token count than the raw text provided. Option E is correct, as the cost for the 'Translate' function is 1.50 Credits per one million Tokens processed.

NEW QUESTION # 132

A data engineering team is setting up an automated pipeline in Snowflake to process call center transcripts. These transcripts, once loaded into a raw table, need to be enriched by extracting specific entities like the customer's name, the primary issue reported, and the proposed resolution. The extracted data must be stored in a structured JSON format in a processed table. The pipeline leverages a SQL task that processes new records from a stream. Which of the following SQL snippets and approaches, utilizing Snowflake Cortex LLM functions, would most effectively extract this information and guarantee a structured JSON output for each transcript?

- ☐ Use `SNOWFLAKE.CORTEX.EXTRACT_ANSWER()` multiple times with separate questions for 'customer name', 'issue', and 'resolution', then use SQL JSON functions to combine the results into a single JSON object.
- ☐ Use `SNOWFLAKE.CORTEX.COMPLETE()` with a prompt like 'Extract customer name, issue, and resolution as a JSON object: [transcript_text]' and rely on the LLM's natural ability to generate JSON.
- ☒ Use `AI_COMPLETE()` with a single prompt and explicitly define a `response_format` argument containing a JSON schema for 'customer_name', 'issue', and 'resolution' fields.
- ☐ Use `SNOWFLAKE.CORTEX.SUMMARIZE()` on the transcript and then manually parse the summary text using regular expressions within SQL to extract the desired entities.
- ☐ Create a Python UDF that calls an external LLM API to extract the entities and return a JSON string, then integrate this UDF into the SQL task.

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

Answer: C

Explanation:

To guarantee a structured JSON output for entity extraction, (the updated version of 'COMPLETE()') with the `response_format` argument and a specified JSON schema is the most effective approach. This mechanism enforces that the LLM's output strictly conforms to the predefined structure, including data types and required fields, significantly reducing the need for post-processing and improving data quality within the pipeline. Option A requires multiple calls and manual JSON assembly, which is less efficient. Option B relies on the LLM's 'natural ability' to generate JSON, which might not be consistently structured without explicit 'response_format'. Option D uses , which is for generating summaries, not structured entity extraction. Option E involves external LLM API calls and Python UDFs, which, while possible, is less direct than using native 'AI_COMPLETE' structured outputs within a SQL pipeline in Snowflake Cortex for this specific goal.

NEW QUESTION # 133

.....

There are also free demos of our GES-C01 study materials on the website that you can download before placing the orders. Taking full advantage of our GES-C01 practice guide and getting to know more about them means higher possibility of winning. And our GES-C01 Exam Quiz is a bountiful treasure you cannot miss. Not only the content is the latest and valid information, but also the displays are varied and interesting. Just have a try and you will love them!

New GES-C01 Test Simulator: <https://www.surepassexams.com/GES-C01-exam-bootcamp.html>

Efficient purchase, In peacetime, you may take months or even a year to review a professional exam, but with GES-C01 exam guide, you only need to spend 20-30 hours to review before the exam, and with our GES-C01 study materials, you will no longer need any other review materials, because our GES-C01 study materials has already included all the important test points, We mainly provide GES-C01 actual test questions for the industry certification personnel exam (examination reference), and our database is a software type, after you purchase pass-for-sure GES-C01 test torrent, it will be delivered online email to you.

Protecting and Using Your Devices in the Kitchen, By Jose Chinchilla, Stacia Varga, Efficient purchase, In peacetime, you may take months or even a year to review a professional exam, but with GES-C01 exam guide, you only need to spend 20-30 hours to review before the exam, and with our GES-C01 Study Materials, you will no longer need any other review materials, because our GES-C01 study materials has already included all the important test points.

Providing You Useful Exam GES-C01 Tips with 100% Passing Guarantee

We mainly provide GES-C01 actual test questions for the industry certification personnel exam (examination reference), and our database is a software type, after you purchase pass-for-sure GES-C01 test torrent, it will be delivered online email to you.

No one can deny the great significance of our GES-C01 pass-sure torrent materials in helping more and more candidates achieve their success with less time and higher efficiency in their life career.

And the most convenient thing about this type of GES-C01 practice exam is that you don't have to install any software as it is a GES-C01 web-based practice exam.

- GES-C01 Top Dumps □ GES-C01 Preparation Store □ Real GES-C01 Torrent □ Download ➡ GES-C01 □ for free by simply searching on > www.examcollectionpass.com < □ GES-C01 Exam Reviews
- Fast, Hands-On GES-C01 Exam-Preparation Questions □ Immediately open □ www.pdfvce.com □ and search for [GES-C01] to obtain a free download □ GES-C01 Preparation Store
- Interactive GES-C01 Practice Exam □ Pdf Demo GES-C01 Download □ GES-C01 Exam Tests ➡ Search for [GES-C01] and easily obtain a free download on > www.testsimulate.com □ □ Pdf Demo GES-C01 Download
- GES-C01 Preparation Store □ New GES-C01 Exam Question □ GES-C01 Best Vce □ Download { GES-C01 } for free by simply searching on ⇒ www.pdfvce.com ⇐ □ New GES-C01 Test Preparation
- Pdf Demo GES-C01 Download □ GES-C01 Top Dumps □ Pdf Demo GES-C01 Download □ Search for “GES-C01 ” and download it for free immediately on 【 www.prep4sures.top 】 □ GES-C01 Valid Study Plan
- GES-C01 Test-king File - GES-C01 Practice Materials - GES-C01 Torrent Questions □ Copy URL { www.pdfvce.com } open and search for > GES-C01 □ to download for free □ GES-C01 Valid Study Plan
- GES-C01 Exam Tests □ Interactive GES-C01 Practice Exam □ GES-C01 Best Vce □ Download □ GES-C01 □ for free by simply entering □ www.examcollectionpass.com □ website □ GES-C01 Valid Real Exam
- 100% Pass Quiz GES-C01 - Accurate Exam SnowPro® Specialty: Gen AI Certification Exam Tips □ Search for > GES-C01 □ on { www.pdfvce.com } immediately to obtain a free download □ Pdf Demo GES-C01 Download
- 2025 Exam GES-C01 Tips - High-quality Snowflake New GES-C01 Test Simulator: SnowPro® Specialty: Gen AI Certification Exam □ Download > GES-C01 < for free by simply searching on { www.examsreviews.com } □ Pdf Demo GES-C01 Download
- 100% Pass Quiz GES-C01 - Accurate Exam SnowPro® Specialty: Gen AI Certification Exam Tips □ Easily obtain free download of ➡ GES-C01 □ by searching on □ www.pdfvce.com □ ☞ Interactive GES-C01 Practice Exam
- GES-C01 Top Dumps □ Interactive GES-C01 Practice Exam □ GES-C01 Best Vce □ Download 「 GES-C01 」 for free by simply entering 「 www.dumps4pdf.com 」 website □ Online GES-C01 Bootcamps
- myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, graphicschoolacademy.com, bonich.org, thestartuptribe.biz, digitalchakku.com, pct.edu.pk, www.stes.tyc.edu.tw, joyrulez.com, www.stes.tyc.edu.tw, vxlxemito123.bloginwi.com, Disposable vapes