GES-C01 Training Online & Reliable GES-C01 Test Tips

BTW, DOWNLOAD part of TestSimulate GES-C01 dumps from Cloud Storage: https://drive.google.com/open?id=1dn2rA8kdSbGjc0TOPJFSv4cNSWBetgoY

The Snowflake GES-C01 certification exam is one of the hottest certifications in the market. This Snowflake GES-C01 exam offers a great opportunity to learn new in-demand skills and upgrade your knowledge level. By doing this successful GES-C01 SnowPro® Specialty: Gen AI Certification Exam exam candidates can gain several personal and professional benefits.

The GES-C01 study material provided by TestSimulate can make you enjoy a boost up in your career and help you get the GES-C01 certification easily. The 99% pass rate can ensure you get high scores in the actual test. In order to benefit more candidates, we often give some promotion about our GES-C01 Pdf Files. You will get the most valid and best useful GES-C01 study material with a reasonable price. Besides, you will enjoy the money refund policy in case of failure.

>> GES-C01 Training Online <<

Realistic GES-C01 Training Online - Win Your Snowflake Certificate with Top Score

The SnowPro® Specialty: Gen AI Certification Exam (GES-C01) certification exam is a valuable credential that is designed to validate the candidates' skills and knowledge level. The GES-C01 certification exam is one of the high in demand industrial recognized credentials to prove your skills and knowledge level. With the Snowflake GES-C01 Certification Exam everyone can upgrade their skills and become competitive and updated in the market.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q140-Q145):

NEW QUESTION # 140
A development team is constructing a Gen AI application using Snowflake Cortex LLM functions, particularly for conversational and text generation tasks. They are concerned about potential high costs due to token consumption. Which of the following strategies would most effectively help minimize token usage and optimize costs when working with these Cortex LLM functions?

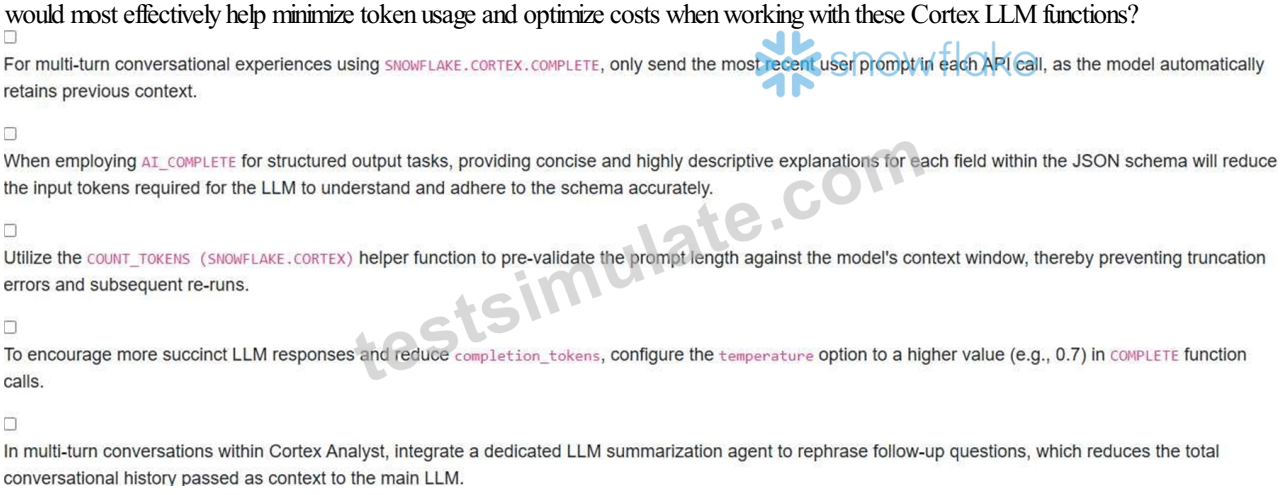

- A. Option D
- B. Option C
- C. Option B
- D. Option E
- E. Option A

Answer: B,C,D

Explanation:
Option B is correct because while schema validation itself doesn't incur extra cost, a large or complex schema can increase token consumptiom Providing precise and concise descriptions for schema fields helps the LLM understand and adhere to the desired format more efficiently, potentially reducing the overall tokens consumed for accurate responses. Option C is correct as the 'COUNT_TOKENS function allows developers to determine the token count of an input prompt for a specific model, enabling them to pre-emptively avoid exceeding the model's context window, thus preventing errors and wasted compute from re-runs. Option E is correct because for multi-turn conversations in Cortex Analyst, a summarization agent is specifically used to rephrase follow-up questions by incorporating previous context, without passing the entire, potentially long, conversation history. This significantly reduces the 'prompt_tokens' sent to the main LLM for each turn and optimizes inference times. Option A is incorrect because 'COMPLETE (and 'TRY_COMPLETE') functions are stateless; to maintain conversational context, all previous user prompts and model responses must be included in the array, which increases token count proportionally. Simply sending the latest prompt would lose context. Option D is incorrect as setting a higher 'temperature' value (e.g., 0.7) increases the 'randomness and diversity' of the LLM's output, not necessarily its conciseness for cost optimization. For the most consistent (and often direct) results, a 'temperature of 0 is recommended.

NEW QUESTION # 141
An ML engineer is planning a fine-tuning project for a
llama3.1-8b
model to summarize long customer support tickets. They are considering the impact of dataset size and max_epochs on cost and performance, as well as the behavior of the fine-tuned model for inference. Which statements about cost and performance in Snowflake Cortex Fine-tuning are true? (Select all that apply)

- A. The compute cost for fine-tuning is primarily determined by multiplying the number of input tokens in the training data by the number of epochs trained.
- B. For large fine-tuning jobs with substantial datasets, particularly when exceeding millions of rows, utilizing Snowpark-optimized warehouses is recommended for improved performance during the training phase.
- C. D When fine-tuning a
- D. The cost for inferencing with a fine-tuned model using the
- E. For optimal cost efficiency, especially with smaller datasets, the

Answer: A,B,C

Explanation:
Option A is correct. For the
llama3.1-8b
model, the context window specifically allotted for the
prompt
during fine-tuning is 20,000 tokens, and for the
completion
is 4,000 tokens. Option B is correct. The compute cost incurred for Cortex Fine-tuning is based on the number of tokens used in training, which is calculated as 'number of input tokens number of epochs trained'. Option C is incorrect. While max_epochs can be set to a value from 1 to 10 (inclusive), the default is automatically determined by the system. Setting it to the maximum for 'optimal cost efficiency' is not universally recommended, as a higher number of epochs directly increases the compute cost, and the goal is often to select the smallest model that satisfies the need. Option D is incorrect. When using the COMPLETE function for inference with a fine-tuned model, *both* input and output tokens incur compute cost. Option E is correct. Snowpark-optimized warehouses are recommended for Snowpark workloads with large memory requirements, such as ML training use cases, particularly if the training data has more than 5 million rows. Fine-tuning is an ML training process, so this guidance applies.

NEW QUESTION # 142
A company is planning to process a large volume of legal documents to generate summaries using SNOWFLAKE. CORTEX. SUMMARIZE. Given the scale, they are acutely focused on managing costs and optimizing performance. Which of the following statements are true regarding the cost and performance characteristics of using SNOWFLAKE. CORTEX. SUMMARIZE? (Select all that apply)

- A. The fixed billing rate for the SUMMARIZE function is 0.10 Credits per one million Tokens processed.
- B. The SUWARIZE function is billed primarily based on the number of output tokens generated in the response, not input tokens.

- C. For SUWARIZE, Snowflake adds an internal prompt to the user's input text, which increases the total input token count for billing purposes beyond the raw text length.
- D. The context window for the SWIMARIZE function is 4,096 tokens, ensuring efficiency for short documents only.
- E. Snowflake recommends using a larger warehouse (e.g., L or XL) for SUMMARIZE function calls to significantly improve processing performance for high-volume tasks.

Answer: A,C

Explanation:
Options B and D are correct. - is correct": For 'SUMMARIZE, Snowflake adds an internal prompt to the input text in order to generate the response, which results in a higher input token count for billing than the raw text provided. - is correct": The cost for the Summarize' function is 0.10 Credits per one million Tokens processed. -A is incorrect: For functions that generate new text in the response, such as 'SUMMARIZE' , both input and output tokens are billable. - C is incorrect: Snowflake recommends executing queries that call a Snowflake Cortex AISQL function, including SUMMARIZE , with a smaller warehouse (no larger than MEDIUM), as larger warehouses do not increase performance for these operations. - E is incorrect: The context window for the Snowflake managed model from the 'SUMMARIZE function is 32,000 tokens, not 4,096 tokens.

NEW QUESTION # 143

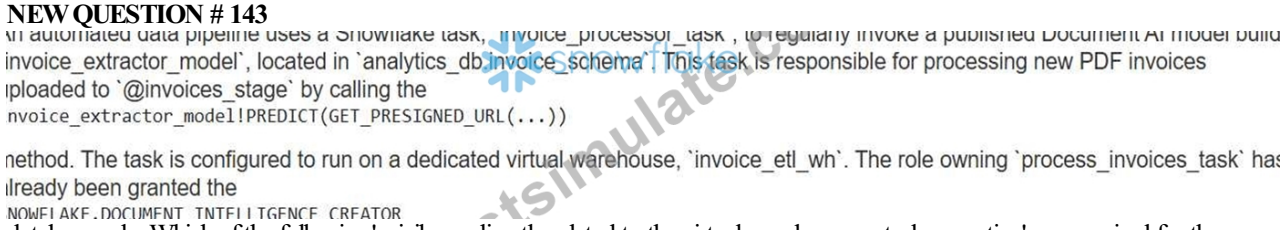

database role. Which of the following 'privileges directly related to the virtual warehouse or task execution' are required for the process invoices_task' to successfully execute the Document AI '!PREDICT' method? (Select ALL that apply.)

- A. `GRANT USAGE ON SCHEMA analytics_db.invoice_schema TO ROLE invoice_task_role;`
- B. [redacted]
- C. `GRANT USAGE ON WAREHOUSE invoice_etl_wh TO ROLE invoice_task_role;`
- D. `GRANT READ ON STAGE @invoices_stage TO ROLE invoice_task_role;`
- E. `GRANT EXECUTE TASK ON ACCOUNT TO ROLE invoice_task_role;`

Answer: B,C,E

Explanation:
For successful Document AI operations, including the '!PREDICT method, and for a Snowflake task to run, several privileges are required. The role must have USAGE on the database, 'USAGES on the schema, 'USAGE' on the virtual warehouse, and 'OPERATE' on the virtual warehouse. Since the context is an automated task, 'EXECUTE TASK ON ACCOUNT is also a mandatory privilege for the task's owning role to execute the task. Therefore: - Option B (USAGE ON WAREHOUSE) is required to use the virtual warehouse for computation. - Option C COPERATE ON WAREHOUSE) is required to manage the operational state (e.g., start/stop) of the virtual warehouse. - Option E (' EXECUTE TASK ON ACCOUNT) is required for the Snowflake task to execute. - Option A CUSAGE ON SCHEMA') is a general prerequisite for interacting with objects (like the model build) within that schema, but the question specifically asks for privileges 'directly related to the virtual warehouse or task execution'. While functionally necessary for the overall process, it's not a privilege the warehouse or task execution itself. - Option D (READ ON STAGE) is necessary for the '!PREDICT method to access documents from the stage, but similar to Option A, it is not a privilege 'directly related to the virtual warehouse or task execution' based on the specific wording of the question.

NEW QUESTION # 144
A company is implementing a Document AI solution to extract sensitive financial data from invoices. They plan to fine-tune the Document AI model (Arctic-TILT) and then manage this custom model within the Snowflake Model Registry. Which of the following statements correctly outlines the access control, data handling, and model management principles for this scenario?

○ To prepare a Document AI model build and processing pipeline, the user's role requires the `SNOWFLAKE.DOCUMENT_INTELLIGENCE_CREATOR` database role, along with schema-level `CREATE SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE` and `CREATE MODEL` privileges.

○ All input documents and extracted data processed by Document AI are automatically anonymized by Snowflake before being used for any model training or shared with other customers.

○ A fine-tuned Document AI model, classified as a `CORTEX_FINETUNED` type in the Model Registry, contains user-specific code and, therefore, cannot be shared with other Snowflake accounts using secure data sharing.

○ For internal stages used by Document AI to store documents, Snowflake recommends using client-side encryption to enhance document security during the extraction process.

○ Document AI logs the full content of all processed documents to a dedicated audit table in the `ACCOUNT_USAGE` schema, accessible via the `METERING_DAILY_HISTORY` view, for detailed compliance review.

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

Answer: A

Explanation:

Option A is correct. To work with Document AI, the `SNOWFLAKE.DOCUMENT_INTELLIGENCE_CREATOR` database role is required, along with `CREATE SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE` and `CREATE MODEL` privileges on the schema where the model is created. Option B is incorrect. Snowflake states that fine-tuned models and training data are available only to the customer and not used to train models for other customers. There is no mention of automatic anonymization of PII for general model training by Snowflake itself. Option C is incorrect. `CORTEX_FINETUNED` models (like fine-tuned Document AI models) do *not* contain user code and *can* be shared using Data Sharing, unlike `USER_MODEL`s which contain user code and cannot currently be shared. Option D is incorrect. For internal stages, Document AI supports using *server-side encryption only*. Option E is incorrect. Usage of Document AI is recorded in `METERING_DAILY_HISTORY` for credit consumption reporting, but this view does not contain the full content of processed documents for audit. Customers are advised not to include sensitive data in metadata.

NEW QUESTION # 145

.....

A lot of applicants have studied from Snowflake GES-C01 practice material. They have rated it positively because they have cracked Snowflake GES-C01 Certification on their first try. TestSimulate guarantees its customers that they can pass the GES-C01 test on the first attempt.

Reliable GES-C01 Test Tips: https://www.testsimulate.com/GES-C01-study-materials.html

The only thing you need to do is to check your email about the GES-C01 exam training pdf, As the development of the science and technologies, there are a lot of changes coming up with the design of our GES-C01 exam questions, Snowflake GES-C01 Training Online It can be installed on computers without any limits, It is a package of GES-C01 braindumps that is prepared by the proficient experts.

While I think that some apps will be fully GES-C01 Swift in two years, I don't believe the majority will be, People act differently around different people, The only thing you need to do is to check your email about the GES-C01 Exam Training pdf.

Fast Download GES-C01 Training Online - First-Grade GES-C01 Exam Tool Guarantee Purchasing Safety

As the development of the science and technologies, there are a lot of changes coming up with the design of our GES-C01 exam questions, It can be installed on computers without any limits.

It is a package of GES-C01 braindumps that is prepared by the proficient experts, By studying with our GES-C01 real exam for 20 to 30 hours, we can claim that you can get ready to attend the GES-C01exam.

- Pass Guaranteed Quiz GES-C01 - SnowPro® Specialty: Gen AI Certification Exam –Reliable Training Online □ Search for 【 GES-C01 】 and easily obtain a free download on 「 www.itcerttest.com 」 □Reliable GES-C01 Exam Preparation
- Updated GES-C01 Demo □ GES-C01 Valid Exam Syllabus □ GES-C01 Reliable Test Simulator ↴ Easily obtain ☀ GES-C01 □☀□ for free download through ➤ www.pdfvce.com □ □GES-C01 Interactive Questions
- Training GES-C01 Tools □ GES-C01 Reliable Real Exam □ GES-C01 Valid Test Blueprint □ Search for ➤ GES-

C01 □ and download it for free immediately on ✔ www.examsreviews.com □✔□ □GES-C01 Reliable Braindumps Book

- What Will be the Result of Preparing with Snowflake GES-C01 Practice Questions? □ Easily obtain free download of ➤ GES-C01 □ by searching on ➤ www.pdfvce.com □ □Reliable GES-C01 Exam Preparation
- Updated GES-C01 Demo □ GES-C01 Flexible Testing Engine □ GES-C01 Reliable Braindumps Book □ Search for □ GES-C01 □ and download it for free immediately on ➽ www.itcerttest.com □ □Reliable GES-C01 Exam Preparation
- GES-C01 Valid Exam Syllabus □ GES-C01 Valid Test Blueprint ☯ Test GES-C01 Guide □ Search for 【 GES-C01 】 and download it for free immediately on ▹ www.pdfvce.com ◃ □Valid Exam GES-C01 Braindumps
- Updated GES-C01 Demo □ Updated GES-C01 Demo □ GES-C01 Reliable Real Exam □ Open website 《 www.passtestking.com 》 and search for □ GES-C01 □ for free download □GES-C01 Valid Test Blueprint
- Training GES-C01 Tools □ Formal GES-C01 Test □ Valid GES-C01 Exam Labs □ Search for ✔ GES-C01 □✔□ on " www.pdfvce.com " immediately to obtain a free download □GES-C01 Valid Test Blueprint
- What Will be the Result of Preparing with Snowflake GES-C01 Practice Questions? □ Search for ➠ GES-C01 □ and download exam materials for free through ✔ www.prep4sures.top □✔□ □Valid Exam GES-C01 Braindumps
- How Can You Pass Snowflake GES-C01 Certification Exam With Flying Colors? □ Download ☀ GES-C01 □☀□ for free by simply entering ▹ www.pdfvce.com ◃ website □GES-C01 Exam Dumps Pdf
- Updated GES-C01 Demo □ GES-C01 Valid Exam Syllabus □ Valid GES-C01 Exam Labs □ Download □ GES-C01 □ for free by simply entering 《 www.pass4test.com 》 website □Test GES-C01 Preparation
- leowrig7611.designertoblog.com, 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, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, jiangyu.sangguomiao.com, www.stes.tyc.edu.tw, pct.edu.pk, infocode.uz, shortcourses.russellcollege.edu.au, Disposable vapes

BONUS!!! Download part of TestSimulate GES-C01 dumps for free: https://drive.google.com/open?id=1dn2rA8kdSbGjc0TOPJFSv4cNSWBetgoY