

GES-C01 Review Guide & GES-C01 Answers Free

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

In modern society, innovation is of great significance to the survival of a company. The new technology of the GES-C01 practice prep is developing so fast. So the competitiveness among companies about the study materials is fierce. Luckily, our company masters the core technology of developing the GES-C01 Exam Questions. On one hand, our professional experts can apply the most information technology to compile the content of the GES-C01 learning materials. On the other hand, they also design the displays according to the newest display technology.

Before you buy our GES-C01 study questions you can have a free download and tryout and you can have an understanding of our product by visiting our pages of our product on the website. The pages of our GES-C01 guide torrent provide the demo and you can understand part of our titles and the form of our software. On the pages of our GES-C01 exam torrent you can see the version of the product, the updated time, the quantity of the questions and answers, the characteristics and merits of the product, the price of the product and the discounts. The pages also list the details and the guarantee of our GES-C01 Exam Torrent, the methods to contact us, the evaluations of the past client on our product, the related exams and other information about our GES-C01 guide torrent. So before your purchase you can have an understanding of our product and then decide whether to buy our GES-C01 study questions or not.

>> GES-C01 Review Guide <<

GES-C01 Answers Free, New Study GES-C01 Questions

It is a popular belief that only professional experts can be the leading one to do some adept job. And similarly, only high quality and high accuracy GES-C01 Exam Questions like ours can give you confidence and reliable backup to get the certificate smoothly

because our experts have extracted the most frequent-tested points for your reference. Good practice materials like our SnowPro® Specialty: Gen AI Certification Exam study question can educate exam candidates with the most knowledge. Do not make your decisions now will be a pity for good.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q15-Q20):

NEW QUESTION # 15

An ML engineer is designing a Cortex Agent to provide highly accurate and contextualized responses. They intend for the agent to use state-of-the-art LLMs for orchestration and to maintain a specific brand tone in its outputs. Considering the available models and configurations for Cortex Agents, which statement is true?

- A. To ensure the agent's responses adhere to a desired brand and tone, 'Response instructions' can be configured, which guide the agent's output style and persona.
- B. Cortex Agents primarily interact with data through the Snowflake Model Registry API to retrieve and update fine-tuned model parameters during their iterative planning phase.
- C. When an agent utilizes an LLM for orchestration, the system ensures that cross-region inference is automatically enabled without any latency implications, making region selection irrelevant.
- D. Cortex Agents are restricted to using only Snowflake Arctic models for orchestration, due to security and governance requirements, and must operate in the account's default region.
- E. The agent's 'Planning' component is specifically responsible for evaluating results after each tool use and deciding the subsequent steps in the query resolution process, acting as a feedback loop.

Answer: A

Explanation:

Option A is incorrect. Cortex Agents can use a variety of models for orchestration, including

□, and others. Cross-region inference can be used if a model is not available in the local region. Option B is correct. 'Response instructions' allow configuration of the agent's responses to a preferred brand and tone. Option C is incorrect. While cross-region inference can be used for models not available locally, latency between regions depends on cloud provider infrastructure and network status. Therefore, it does have latency implications. Option D is incorrect. Cortex Agents leverage Cortex Analyst and Cortex Search as tools, along with LLMs, to analyze data. While fine-tuned models might be involved, the primary interaction is not described as being through the Snowflake Model Registry API for active data processing or parameter updates during the planning phase. The Model Registry is for logging and managing models. Option E is incorrect. The 'Reflection' component of Cortex Agents is responsible for evaluating results after each tool use to determine the next steps (asking for clarification, iterating, or generating a final response). The 'Planning' component is for orchestrating a plan to address a request.

NEW QUESTION # 16

A data team is deploying a new customer service chatbot using Snowflake Cortex Analyst. To accurately forecast and optimize their costs, the team needs to understand how Cortex Analyst billing works. Which of the following statements accurately describe the cost considerations for Snowflake Cortex Analyst?

- A. Cortex Analyst incurs compute costs based on the total number of input and output tokens processed in each conversation turn.
- B. Virtual warehouse compute is directly billed for Cortex Analyst operations, and its size can be adjusted to optimize query generation performance and cost.
- C. Costs for Cortex Analyst are primarily driven by the number of successful messages processed, with a rate of 0.067 Credits per message.
- D. Cortex Analyst's Cloud Services compute costs are subject to a daily adjustment, where Snowflake only bills if these costs exceed 10% of the daily virtual warehouse cost for the account.
- E. Only successful responses (HTTP 200) from Cortex Analyst are counted towards the credit usage.

Answer: C,E

Explanation:

Option B is correct because Cortex Analyst credit rate usage is based on the number of messages processed, specifically 67 Credits per 1,000 messages or 0.067 Credits per message. Option E is correct as only successful responses (HTTP 200) from Cortex Analyst are counted towards the credit usage. Option A is incorrect because the number of tokens in each message does not affect the cost for Cortex Analyst; this is a billing model for other Cortex LLM functions, but not Cortex Analyst. Option C is incorrect; Cortex Analyst runs on Snowflake-managed compute resources as a Serverless Feature, meaning virtual warehouse compute is not

directly billed for its operations in the same way as a user-managed warehouse. Option D is incorrect. While Cloud Services compute generally has a 10% daily warehouse cost adjustment, Serverless Features (which include Cortex Analyst) do not contribute to this adjustment.

NEW QUESTION # 17

A data engineer is building a robust pipeline to process customer feedback. They need to extract specific sentiment categories (food_quality, food_taste, wait_time, food_cost) from text reviews and ensure the output is always a valid JSON object matching a predefined schema, even for complex reviews. They also want to control the determinism of the LLM responses. Which of the following SQL statements or considerations are correct for achieving this using Snowflake Cortex AI functions?

- A. Using AI_COMPLETE with response_format incurs additional compute cost for the overhead of verifying each token against the supplied JSON schema, in addition to standard token costs.
- B. For the most consistent structured output, especially in complex reasoning tasks, setting the temperature option to 0 when calling AI_COMPLETE is recommended.
- C. To ensure the model explicitly attempts to extract all specified fields, the 'required' array in the JSON schema is critical; AI_COMPLETE will raise an error if any required field cannot be extracted.
- D. The response_format argument with a JSON schema is primarily for OpenAI (GPT) models; for other models like Mistral, a strong prompt instruction such as 'Respond in strict JSON' is generally more effective.
- E. The following SQL statement uses the response_format argument and temperature setting to achieve structured output and determinism

□

Answer: B,C,E

Explanation:

Option A is correct because it demonstrates the proper use of the 'AI_COMPLETE' function with the 'response_format' argument to specify a JSON schema and sets 'temperature' to 0 for consistent output, as per the documentation. Option C is correct as the "required" field in the JSON schema ensures that specific fields must be extracted, and 'COMPLETE' (or 'AI_COMPLETE') will raise an error if these fields cannot be found. Option E is correct because for the most consistent results, setting the 'temperature*' option to 0 is recommended when calling 'COMPLETE' (or 'AI_COMPLETE') with structured outputs, regardless of the task or model. Option B is incorrect because all models supported by support structured output, and specifying the 'response_format' is the direct mechanism to enforce a schema, although for complex tasks, adding 'Respond in JSON' to the prompt can improve accuracy. Option D is incorrect as 'AI_COMPLETE Structured Outputs incurs compute cost based on the number of tokens processed, but it does not incur additional compute cost for the overhead of verifying each token against the supplied JSON schema.

NEW QUESTION # 18

A company is developing a Streamlit application leveraging Snowflake Cortex Analyst for natural language querying over sales data. They want to implement a robust multi-turn conversational experience where users can ask follow-up questions. Which of the following statements accurately describe the design and cost implications of supporting multi-turn conversations in Cortex Analyst? (Select all that apply)

- A. When an LLM judge is used to evaluate the summarization quality for multi-turn conversations, a smaller model like Llama 3.1 8B is generally preferred over Llama 3.1 70B to minimise latency, even if it leads to a slightly higher error rate in rewritten questions.
- B. Cortex Analyst supports multi-turn conversations by simply passing the entire conversation history directly to every LLM call within its agentic workflow, which is the most efficient method for maintaining context.
- C. Developers can manually implement multi-turn conversations in their applications by using the
- D. An internal LLM summarization agent is automatically employed by Cortex Analyst before its original workflow to reframe follow-up questions based on conversation history, optimising LLM processing for each agent.
- E. The cost for Cortex Analyst's multi-turn conversational support is primarily based on the number of messages processed, and the number of tokens within each message does not directly affect the per-message cost.

Answer: C,D,E

Explanation:

Option A is incorrect. Cortex Analyst does not simply pass the entire conversation history to every LLM call. This primitive approach would lead to longer inference times, more non-determinism, and degraded performance. Instead, it employs an LLM summarization agent. Option B is correct. To support multi-turn conversations, Cortex Analyst adds an additional LLM summarization agent before its original workflow to reframe questions based on the conversation history. Option C is incorrect.

While there's a latency-performance tradeoff, Llama 3.1 70B was found to be sufficient for the summarization task in Cortex Analyst, achieving 96.5% good ratings from an LLM judge, significantly outperforming Llama 3.1 8B which had an approximate 5% error rate in rewritten questions. Option D is correct. The credit rate usage for Cortex Analyst is based on the number of messages processed (67 Credits per 1,000 messages or 0.067 Credits per message), and the number of tokens in each message does not affect this per-message cost. Option E is correct. The 'SNOWFLAKE.CORTEX.COMPLETE function, and its 'TRY COMPLETE' variant, can be used to provide a stateful conversational experience by passing an array of objects for the argument, where each object contains a 'role' and 'content' for previous user prompts and model responses.

NEW QUESTION # 19

An operations team at a company is implementing a robust governance framework to monitor and optimize the costs associated with their Snowflake Cortex LLM function usage. They need to identify which functions are driving the highest token consumption and overall credit usage to pinpoint areas for cost reduction. Which of the following monitoring tools or methods are appropriate for gaining these insights into Cortex LLM function costs and token consumption?

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

Answer: A,B,C,D

Explanation:

NEW QUESTION # 20

.....

Fast2test SnowPro® Specialty: Gen AI Certification Exam(GES-C01) Questions have numerous benefits, including the ability to demonstrate to employers and clients that you have the necessary knowledge and skills to succeed in the actual GES-C01 exam. Certified professionals are often more sought after than their non-certified counterparts and are more likely to earn higher salaries and promotions. Moreover, cracking the SnowPro® Specialty: Gen AI Certification Exam(GES-C01) exam helps to ensure that you stay up to date with the latest trends and developments in the industry, making you more valuable assets to your organization.

GES-C01 Answers Free: <https://www.fast2test.com/GES-C01-premium-file.html>

The free demos do honor to the perfection of our latest GES-C01 exam torrent, and also a performance of our considerate after sales services, In addition, your questions about our GES-C01 exam prep: SnowPro® Specialty: Gen AI Certification Exam will be answered completely and correctly, DumpStep Dumps for GES-C01 exam are written to the highest standards of technical accuracy, provided by our certified subject matter experts and published authors for development, We have three version of our GES-C01 exam questions which can let you study at every condition so that you can make full use of your time.

from Cambridge University, and degrees in statistics GES-C01 and music, The basic experience of the world as chaos" is rooted here, The free demos do honor to the perfection of our latest GES-C01 Exam Torrent, and also a performance of our considerate after sales services.

Snowflake GES-C01 Exam | GES-C01 Review Guide - Ensure you Pass GES-C01: SnowPro® Specialty: Gen AI Certification Exam Exam

In addition, your questions about our GES-C01 exam prep: SnowPro® Specialty: Gen AI Certification Exam will be answered completely and correctly, DumpStep Dumps for GES-C01 exam are written to the highest standards of technical GES-C01 Answers Free accuracy, provided by our certified subject matter experts and published authors for development.

We have three version of our GES-C01 exam questions which can let you study at every condition so that you can make full use of your time, As you know the passing rate of them has come up to 98-100 percent with an increasingly trend.

- GES-C01 Latest Exam Answers GES-C01 Best Preparation Materials Study GES-C01 Demo Search for (GES-C01) and download exam materials for free through 「 www.vce4dumps.com 」 Study GES-C01 Plan
- Study Your Snowflake GES-C01 Exam with Pass-Sure GES-C01 Review Guide: SnowPro® Specialty: Gen AI

Certification Exam Efficiently □ Open www.pdfvce.com □ and search for 「 GES-C01 」 to download exam materials for free □ Latest GES-C01 Exam Review