

# GES-C01 Practice Materials & GES-C01 Training Guide Torrent - FreePdfDump



Whether for a student or an office worker, obtaining GES-C01 certificate can greatly enhance the individual's competitiveness in the future career. Try our GES-C01 study materials, which are revised by hundreds of experts according to the changes in the syllabus and the latest developments in theory and practice. Once you choose GES-C01 training dumps, passing the exam one time is no longer a dream.

The catch is that passing the Snowflake GES-C01 exam is not as easy as it seems to be. It requires sheer determination, a thorough understanding of each topic, and critical thinking when posed with tricky problems. That is the reason why FreePdfDump have come up with a solution by providing the most updated prep material created under the supervision of 90,0000 experienced Snowflake professionals. This GES-C01 Exam Dumps is made to polish your abilities, help you understand every topic, and pass you Snowflake GES-C01 exam on your first attempt.

>> [Exam GES-C01 Reference <<](#)

## GES-C01 Valid Exam Tutorial & Authentic GES-C01 Exam Questions

We are glad to receive all your questions on our GES-C01 learning guide. If you have any questions about our GES-C01 study questions, you have the right to answer us in anytime. Our online workers will solve your problem immediately after receiving your questions. Because we hope that you can enjoy the best after-sales service. We believe that our GES-C01 Preparation exam will meet your all needs. Please give us a chance to service you; you will be satisfied with our GES-C01 study materials.

## Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q281-Q286):

### NEW QUESTION # 281

A Gen AI developer is deploying a customer support chatbot using 'SNOWFLAKE.CORTEX.COMPLETE' for generating responses. To ensure the chatbot does not provide inappropriate or harmful content, they have enabled Cortex Guard. A specific user prompt is now causing the LLM to attempt generating content that violates the defined safety policies. What is the expected outcome of the 'COMPLETE' function call in this scenario, and how is it typically communicated to the application?

- A. The Snowflake account will be temporarily locked for violating the acceptable use policy, and all subsequent LLM calls will fail until manual intervention.
- B. The 'COMPLETE' function will raise a SQL exception, indicating a 'PROHIBITED\_CONTENT' error, which the calling application must handle with a 'TRY...CATCH' block.
- C. The 'COMPLETE' function will return a NULL value in the response, signifying that no safe response could be generated, similar to error handling.
- D. The 'COMPLETE' function will return a generic, pre-defined filtered message, such as 'Response filtered by Cortex Guard', in place of the harmful content.
- E. The LLM, through Cortex Guard's internal mechanisms, will automatically rephrase or redact the harmful parts of its response to comply with safety guidelines before returning it.

**Answer: D**

Explanation:

When Cortex Guard is enabled ('guardrails: true) for 'SNOWFLAKE.CORTEX.COMPLETE' or via the Cortex LLM REST API, and the LLM generates potentially unsafe or harmful content, Cortex Guard filters this content and returns a generic filtered message, typically 'Response filtered by Cortex Guard'. It does not raise a SQL exception (Option A) or automatically rephrase the content (Option C). Account locking (Option D) is not the immediate and direct outcome of a single filtered response. While 'TRY\_COMPLETE returns NULL for execution errors, Cortex Guard's filtering mechanism returns a specific message within the normal function output flow, not a NULL, for detected harmful content.

**NEW QUESTION # 282**

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?

□

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

**Answer: D**

Explanation:

To provide a stateful, conversational experience with the 'COMPLETE (SNOWFLAKE.CORTEX)' function (or its latest version, 'AI\_COMPLETE'), all previous user prompts and model responses must be explicitly passed as part of the argument. This argument expects an array of objects, where each object represents a turn and contains a 'role' ('system', 'user', or 'assistant') and a 'content' key, presented in chronological order. In Streamlit, 'st.session\_state' is the standard and recommended mechanism for storing and managing data across reruns of the application, making it ideal for maintaining chat history, by initializing 'st.session\_state.messages = []' and appending messages to it. Option A is incorrect because 'COMPLETE' does not inherently manage history from external tables. Option B is incorrect as 'COMPLETE' does not retain state between calls; history must be explicitly managed. Option D is a less effective form of prompt engineering compared to passing structured history, as it loses the semantic role distinction and can be less accurate for LLMs. Option E describes a non-existent parameter for the 'COMPLETE' function.

**NEW QUESTION # 283**

An analytics engineering team is building a complex, real-time data pipeline in Snowflake. They want to automatically summarize new incoming product reviews using SNOWFLAKE. CORTEX. SUMMARIZE as part of a continuous process. They consider integrating this function into a dynamic table definition for efficient, automated refreshes. Which of the following statements regarding the integration of SNOWFLAKE. CORTEX. SUMMARIZE with Snowflake's data pipeline features is true?

- A. SNOWFLAKE. CORTEX. SUMMARIZE can be used with dynamic tables if the TARGET\_LAG is explicitly set to '1 day' or longer, allowing for asynchronous processing.
- B. **Snowflake Cortex functions, including SNOWFLAKE. CORTEX.SUMMARIZE, are explicitly incompatible with dynamic tables, requiring alternative automation methods such as streams and tasks.**
- C. SNOWFLAKE CORTEX. SUMMARIZE is supported in dynamic tables only if a Snowpark-optimized warehouse is explicitly assigned to the dynamic table for processing.
- D. While SNOWFLAKE. CORTEX. SUMMARIZE does not support dynamic tables, its updated version, is fully compatible with dynamic tables for continuous summarization.
- E. SNOWFLAKE. CORTEX. SUMMARIZE can be directly used within a dynamic table's SELECT statement, enabling real-time summarization with automated refreshes, provided the source data is a stream.

**Answer: B**

Explanation:

Option B is correct. The sources explicitly state that 'Snowflake Cortex functions do not support dynamic tables'. This limitation applies to 'SUMMARIZE' (SNOWFLAKE.CORTEX). Therefore, attempting to use it directly in a dynamic table would result in a failure. Option A is incorrect due to this explicit limitation. Option C is incorrect; while is an updated function, there is no indication that it overcomes the fundamental incompatibility of Cortex functions with dynamic tables. Option D and E propose non-existent

conditions for compatibility.

#### NEW QUESTION # 284

A development team plans to utilize Snowpark Container Services (SPCS) for deploying a variety of AI/ML workloads, including custom LLMs and GPU-accelerated model training jobs. They are in the process of creating a compute pool and need to select the appropriate instance families and configurations. Which of the following statements about 'CREATE COMPUTE POOL' in SPCS are accurate?

- A. Snowpark-optimized warehouses are the recommended compute pool type for all large-scale ML training workloads within SPCS due to their enhanced memory limits and CPU architectures.
- B. To support GPU-accelerated LLM inference and training, the 'INSTANCE\_FAMILY' must be selected from a type starting with 'GPU' (e.g.,
  - C. The 'MIN NODES' and 'MAX NODES' parameters define the scaling range for the compute pool, and Snowflake automatically scales the pool within this range based on workload demand.
  - D. For cost optimization, 'AUTO SUSPEND SECS = 0' should be used to prevent automatic suspension of the compute pool, as suspension and resumption incur minimum billing durations.
  - E. Setting 'AUTO RESUME = TRUE' ensures that the compute pool automatically starts when a service or job is submitted to it, rather than requiring manual resumption.

**Answer: B,E**

Explanation:

Option A is correct. GPU-accelerated workloads, such as LLM inference and model training, require instance families specifically designed with GPUs. The documentation lists instance family names starting with 'GPU' for this purpose, such as 'GPU\_GCP NV L4'. Option B is incorrect. While 'MIN NODES' and 'MAX NODES' define the range, the size of compute clusters in Snowpark Container Services does "not" auto-scale dynamically based on workload demand. Users must manually alter the number of instances at runtime using commands like 'ALTER SERVICE MIN INSTANCES = s'. Snowflake does handle load balancing across instances within the configured node counts. Option C is correct. The 'AUTO\_RESUME = TRUE' parameter, when specified during compute pool creation, enables the pool to automatically resume operation when a service or job is submitted, removing the need for explicit 'ALTER COMPUTE POOL RESUME' commands. Option D is incorrect. Setting = prevents the compute pool from automatically suspending, meaning it will continue to consume credits even when idle. This would generally lead to higher costs, not cost optimization, unless the pool is constantly active. The default is 3600 seconds (1 hour). SPCS Compute Nodes have a minimum charge of five minutes when started or resumed, making intelligent use of auto-suspend important for cost management. Option E is incorrect. Snowpark-optimized warehouses are a type of 'virtual warehouse' and are recommended for Snowpark workloads with large memory requirements or specific CPU architecture, typically for single-node ML training workloads 'within a warehouse'. SPCS compute pools, however, provide their own dedicated instance families (CPU, HighMemory, GPU) for containerized workloads, abstracting the underlying infrastructure and supporting distributed GPU clusters directly within SPCS, not Snowpark-optimized warehouses as a 'compute pool type' for SPCS.

#### NEW QUESTION # 285

A financial institution uses Snowflake Cortex Analyst with strict role-based access control (RBAC) on their Snowflake-hosted LLMs. The security team has granted specific 'CORTEX-MODEL-ROLE' application roles to different analyst teams, ensuring they only access approved models. A new requirement arises to enable Azure OpenAI GPT models for Cortex Analyst to leverage a specific feature. An administrator proceeds to execute:

□ Which of the following statements accurately describe the implications of this change?

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

**Answer: C,E**

Explanation:

Option B is correct because when is 'TRUE', cortex Analyst can use Azure OpenAI models, but this setting is incompatible with model-level RBAC, meaning RBAC is not available for any models used by Cortex Analyst when this parameter is enabled. Option

C is correct because if Azure OpenAI models are opted in for Cortex Analyst, semantic model files (which are metadata) and user prompts will be processed by Microsoft Azure, a third party, thus transmitting them outside Snowflake's governance boundary. Customer data itself is not shared. Option A is incorrect because the parameter is incompatible with model-level RBAC for 'all' models used by Cortex Analyst. Option D is incorrect as the parameter specifically controls the use of Azure OpenAI models within Cortex Analyst. Option E is incorrect because this parameter can only be set by the 'ACCOUNTADMIN' role.

## NEW QUESTION # 286

If you want to pass your exam just one time, then our GES-C01 exam torrent will be your best choice. We can help you pass your exam just one time, and if you fail the exam in your first attempt after using GES-C01 exam torrent, we will give you refund, and no other questions will be asked. Moreover, GES-C01 Exam Braindumps of us are high-quality, and we have helped lots of candidates pass the exam successfully. We have received many good feedbacks from our customers. We offer you online and offline chat service stuff, if you have any questions about GES-C01 exam torrent, you can consult them.

**GES-C01 Valid Exam Tutorial:** <https://www.freepdfdump.top/GES-C01-valid-torrent.html>

Snowflake Exam GES-C01 Reference After finishing your task, you can review them plenty of times and find out the wrong items, some questions may have explanations for your understanding, and you can practice many times day to day, Snowflake Exam GES-C01 Reference We are famous by our high-quality products and high passing-rate, Snowflake Exam GES-C01 Reference Their different point is the way of presentation.

Also be aware that this hasn't changed the functionality GES-C01 of the skeletons, The Application Object Model, After finishing your task, you can review them plenty of times and find out the wrong items, some Exam GES-C01 Reference questions may have explanations for your understanding, and you can practice many times day to day.

**Top Exam GES-C01 Reference | High Pass-Rate Snowflake GES-C01:  
SnowPro® Specialty: Gen AI Certification Exam 100% Pass**

We are famous by our high-quality products and high passing-rate, Their different point is the way of presentation, The useful knowledge can be learnt on our GES-C01 Study Guide.

If you are considering becoming a certified professional about GES-C01 exam, now is the time.

