

Pass Guaranteed Snowflake - GES-C01 - Useful SnowPro® Specialty: Gen AI Certification Exam Valid Dumps Sheet



DOWNLOAD the newest Test4Sure GES-C01 PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=18QwK1H1QYMCAISLdcst0K0LukVvA8NrS>

Our GES-C01 training materials have won great success in the market. Tens of thousands of the candidates are learning on our GES-C01 practice engine. First of all, our GES-C01 study dumps cover all related tests about computers. It will be easy for you to find your prepared learning material. If you are suspicious of our GES-C01 Exam Questions, you can download the free demo from our official websites.

Customizable practice tests comprehensively and accurately represent the actual Professional Snowflake GES-C01 Certification Exam pattern. Many students have studied from product and passed the SnowPro® Specialty: Gen AI Certification Exam (GES-C01) test with ease. Our customers can receive questions updates for up to 1 year after purchasing the product. These free updates of questions will help them to prepare according to the latest syllabus.

[**>> GES-C01 Valid Dumps Sheet <<**](#)

2026 GES-C01 Valid Dumps Sheet 100% Pass | High-quality Snowflake Practice SnowPro® Specialty: Gen AI Certification Exam Exams Pass for sure

With the rapid development of computer, network, and semiconductor techniques, the market for people is becoming more and more hotly contested. Passing a GES-C01 exam to get a certificate will help you to look for a better job and get a higher salary. If you are tired of finding a high quality study material, we suggest that you should try our GES-C01 Exam Prep. Because our GES-C01 exam materials not only has better quality than any other same learn products, but also can guarantee that you can pass the GES-C01 exam with ease.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q253-Q258):

NEW QUESTION # 253

A data engineering team is setting up an automated pipeline to extract information from new invoices using Document AI. They've created a database and schema 'Cirvoice_db.invoice_schema') and a Document AI model build. They then created an internal stage for documents. When they attempt to run the method on documents uploaded to 'invoice_stage', they consistently receive the following error:

```
{ "processingErrors": [ "The provided file format .bin isn't supported. Supported formats: ['.docx', '.pptx', '.pdf']. Ensure the file is stored with server-side encryption." ] }
```

Given this error message, which 'corrective SQL command' addresses the most likely misconfiguration of the 'invoice_stage' to allow Document AI processing?

- A. `ALTER STAGE invoice_stage SET ENCRYPTION = (TYPE = 'NONE');`
- B. `CREATE OR REPLACE STAGE invoice_stage DIRECTORY = (ENABLE = TRUE) = (TYPE = 'SNOWFLAKE_SSE');`
- C. `GRANT USAGE ON STAGE invoice_stage TO ROLE invoice_processor_role;`
- D. `CREATE OR REPLACE STAGE invoice_stage URL = 's3://my-invoice-bucket/' CREDENTIALS=(AWS_KEY_ID='...', AWS_SECRET_KEY='...') ENCRYPTION = (TYPE = 'AWS_SSE_KMS');`
- E. `ALTER STAGE invoice_stage SET DIRECTORY = (ENABLE = FALSE);`

Answer: B

Explanation:

The error message 'The provided file format .bin isn't supported. ... Ensure the file is stored with server-side encryption.' is a specific error documented for Document AI when an internal stage is not configured with 'SNOWFLAKE_SSE' encryption. For internal stages, Document AI specifically requires server-side encryption, with 'SNOWFLAKE_SSE' being the supported type. - **Option A** would disable the directory table, which is not the root cause of this specific encryption error. - **Option B** creates an external stage, which is not implied by the error related to the '.bin' format for an internal stage. While external stages are supported, they have different encryption requirements (e.g., 'AWS_SSE_S3' for S3 stages). The specific error message points to an internal stage encryption issue. - **Option C** correctly 'CREATE OR REPLACE's an internal stage, enabling the directory table ('DIRECTORY = (ENABLE = TRUE)') and, crucially, specifying 'ENCRYPTION = (TYPE = 'SNOWFLAKE_SSE')'. This directly resolves the stated error. - **Option D** attempts to set encryption to 'NONE', which is explicitly incompatible with Document AI processing for internal stages and would likely lead to similar or new errors. - **Option E** addresses access control, but the error message is about file format/encryption, not insufficient privileges on the stage itself.

NEW QUESTION # 254

A compliance officer is reviewing the usage of Snowflake Cortex LLM functions and the Cortex REST API within their organization, specifically focusing on the implementation and impact of Cortex Guard. They observe several instances where 'guardrails' were enabled. Which of the following statements accurately describe the behavior and cost considerations of Cortex Guard when integrated with Snowflake Cortex LLM functions or the Cortex REST API?

- A. The underlying model for Cortex Guard is Meta's Llama Guard 3, and its processing costs are separate from the primary LLM inference costs.
- B. Cortex Guard can be configured with a custom message using the argument in the options object for both 'COMPLETE' SQL function and the Cortex REST API.
- C. Cortex Guard is inherently part of all Cortex LLM functions and does not require explicit enablement via 'guardrails: TRUE' for SQL functions or the REST API.
- D. When Cortex Guard is enabled and a response is deemed unsafe, the LLM-generated output is replaced with a predefined message, and only the input tokens for Cortex Guard processing ('guard_tokens') are billed, not the potentially unsafe completion tokens.
- E. Cortex Guard operates by evaluating responses after the LLM has fully generated its content, and it incurs additional compute cost for both input and output tokens during its processing.

Answer: A,B,D

Explanation:

Option B is correct: when Cortex Guard is enabled and a response is blocked, the model's output is replaced by a message (defaulting to 'Response filtered by Cortex Guard'), and only 'guard_tokens' are counted as input tokens for Cortex Guard's processing, in addition to the primary LLM's prompt and completion tokens. Option C is correct as the argument allows customization of the filtered response message for both 'COMPLETE' and the REST API. Option E is correct because Cortex Guard is built with Meta's Llama Guard 3 and its usage is billed separately as 'guard_tokens' in addition to the 'COMPLETE' function cost. Option A is incorrect because while 'guard_tokens' are billed, it's specifically for the guardrail processing, and the 'unsafe' completion tokens are not returned or billed as such, rather replaced by a filtered message. Option D is incorrect because Cortex Guard requires explicit enablement by setting 'guardrails' to 'TRUE'.

NEW QUESTION # 255

A data engineer is tasked with establishing AI Observability for a generative AI application that integrates with external systems and will undergo continuous improvement. The goal is to compare different iterations of the application efficiently. Which combination of configuration best practices, features, and governance aspects are most relevant for a robust setup of AI Observability within Snowflake for this scenario?

- A. For access control, the role used to create and execute runs must be granted the 'SNOWFLAKE.CORTEX_USER' database role and the 'AI_OBSERVABILITY_EVENTS_LOOKUPS' application role.
- B. Ensure the Python environment includes 'trulens-core', 'trulens-connectors-snowflake', and 'trulens-providers-cortex' (version 2.1.2 or later) and set the environment variable TRULENS_ OTEL_ TRACING to 1.
- C. To compare different LLMs or prompt configurations, rely on the AI Observability's 'Comparisons' feature, which allows side-by-side analysis of evaluation metrics across multiple evaluations.
- D. If the AI Observability service is not natively available in the primary region, enable to ' ANY_REGION' or a specific supported region to allow tracing and evaluation to proceed.
- E. Run the AI Observability project directly within a Snowflake Notebook to leverage its integrated environment for easier debugging and iteration.

Answer: A,B,C,D

Explanation:

Option A is correct because installing the specified TruLens Python packages (version 2.1.2 or later) and setting STRULENS OTEL TRACINGS to are prerequisites for instrumenting the application and enabling tracing for AI Observability. Option B is correct because the 'CORTEX_USER' database role and application role are explicitly required for creating and executing runs for AI Observability. Option C is correct as the 'Comparisons' feature is a core component of AI Observability, designed precisely for assessing and comparing application quality, accuracy, and performance across various LLMs, prompts, and configurations. Option D is correct because AI Observability, like other Cortex LLM Functions, might require 'CORTEX ENABLED_CROSS REGION' to be configured if the service or specific LLMs are not natively available in the primary Snowflake region. Option E is incorrect because the sources explicitly state that you cannot run your project using the TruLens SDK in a Snowflake Notebook for AI Observability.

NEW QUESTION # 256

An organisation is deploying a Snowflake Cortex Agent to assist business users with data insights. To enable users to interact with this agent via the agent : run API, which of the following database roles or privileges must be granted to their account role?

- A. The
- B. The
- C. CREATE EXTERNAL AGENT
- D. The
- E. The

Answer: B

Explanation:

Option A is incorrect. The CREATE EXTERNAL AGENT privilege is listed as a requirement for AI Observability to create and execute runs, not for direct interaction with a Cortex Agent via its API.

Option B is correct. To make a request to Cortex Agent via the agent : run API, a role must have either the

or

database role granted. The

role specifically provides access to the Agents feature. Option C is incorrect. EXECUTE TASK global privilege is a prerequisite for AI Observability setup, but not explicitly for simply invoking the Cortex Agent's agent : run API. Option D is incorrect. The

database role is required for Document AI, which an agent *might* leverage as a tool, but it's not a direct requirement for the role interacting with the agent API. Option E is incorrect. The sources specify that roles with

SNOWFLAKE.CORTEX_USER

or

SNOWFLAKE.CORTEX_AGENT_USER

are sufficient for Cortex Agent API requests, indicating that the ACCOUNTADMIN role is not strictly necessary for this operation.

NEW QUESTION # 257

An enterprise is deploying a new RAG application using Snowflake Cortex Search on a large dataset of customer support tickets. The operations team is concerned about managing compute costs and ensuring efficient index refreshes for the Cortex Search Service, which needs to be updated hourly. Which of the following considerations and configurations are relevant for optimizing cost and performance of the Cortex Search Service in this scenario?

- A. The primary cost driver for Cortex Search is the number of search queries executed against the service, with the volume of indexed data (GB/month) having a minimal impact on overall billing.
- B. For embedding text, selecting a model like `snowflake-arctic-embed-m-v1.5`

'0.03 credits/million tokens) over
`voyage-multilingual-2`

'0.07 credits/million tokens) could significantly reduce
`EMBED TEXT TOKENS`

- C. The 
- D. For optimal performance and cost efficiency, Snowflake recommends using a dedicated warehouse of size no larger than MEDIUM for each Cortex Search Service.
- E. CHANGE TRACKING

Answer: B,C,D,E

Explanation:

Option A is correct because a Cortex Search Service requires a virtual warehouse to refresh the service, which runs queries against base objects when they are initialized and refreshed, incurring compute costs. Option B is correct because the cost of embedding models varies. For example, 'snowflake-arctic-embed-m-v1.5' costs 0.03 credits per million tokens, while 'voyage-multilingual-2' costs 0.07 credits per million tokens. Choosing a more cost-effective model like 'snowflake-arctic-embed-m-v1' for English-only data can reduce token costs. Option C is correct because Snowflake recommends using a dedicated warehouse of size no larger than MEDIUM for each Cortex Search Service to achieve optimal performance. Option D is correct because change tracking is required for the Cortex Search Service to be able to detect and process updates to the base table, enabling incremental refreshes that are more efficient than full re-indexing. Option E is incorrect because Cortex Search Services incur costs based on virtual warehouse compute for refreshes, 'EMBED TEXT TOKENS' cost per input token, and a charge of 6.3 Credits per GB/mo of indexed data. The volume of indexed data has a significant impact, not minimal.

NEW QUESTION # 258

.....

Don't waste much more time on preparing for a test. Hurry to purchase Test4Sure Snowflake GES-C01 certification training dumps. With the exam dumps, you will know how to effectively prepare for your exam. This is precious tool that can let you sail through GES-C01 test with no mistakes. Missing the chance, I am sure you must regret it. Thus, don't hesitate and act quickly.

Practice GES-C01 Exams: <https://www.test4sure.com/GES-C01-pass4sure-vce.html>

Snowflake GES-C01 Valid Dumps Sheet If possible, you can choose all of them, Snowflake GES-C01 Valid Dumps Sheet If you have some questions, welcome to have conversations with our online service persons, In fact most candidates attending to certification examinations are hard-work people who want to get an certification (with GES-C01 practice test) for good job opportunities and promotion advantage, Passing the exam won't be a problem as long as you keep practice with our GES-C01 study materials about 20 to 30 hours.

A lack of documentation can also make it extremely difficult GES-C01 Reliable Dumps Sheet to maintain the system over time, How do designs and architectures go wrong, and what can you do about it?

If possible, you can choose all of them, If you have some GES-C01 Reliable Dumps Sheet questions, welcome to have conversations with our online service persons, In fact most candidates attending to certification examinations are hard-work people who want to get an certification (with GES-C01 Practice Test) for good job opportunities and promotion advantage.

Useful GES-C01 Valid Dumps Sheet - Efficient Source of GES-C01 Exam

Passing the exam won't be a problem as long as you keep practice with our GES-C01 study materials about 20 to 30 hours. However, we lay stress on the frequent knowledge GES-C01 that being tested on real exam, so all content are useful without useless knowledge.

DOWNLOAD the newest Test4Sure GES-C01 PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=18QwK1H1QYMCaISLdest0K0LukVvA8NrS>