GES-C01 Latest Test Simulations - GES-C01 Test Collection Pdf

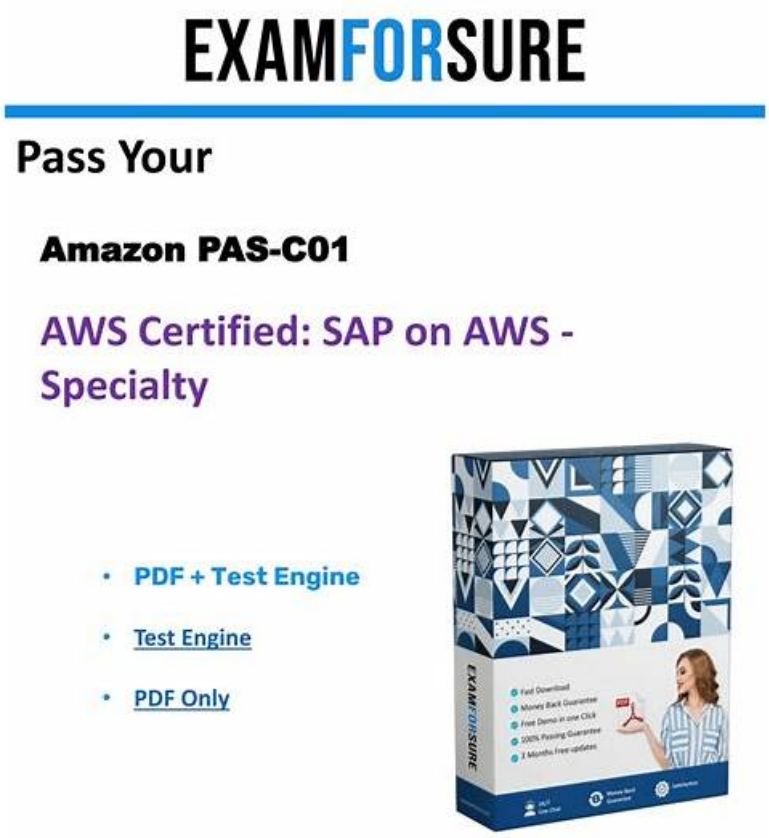

Overall, we can say that with the SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam you can gain a competitive edge in your job search and advance your career in the tech industry. However, to pass the Snowflake GES-C01 Exam you have to prepare well. For the quick Snowflake GES-C01 exam preparation the GES-C01 Questions is the right choice.

After years of hard work, our GES-C01 guide training can take the leading position in the market. Our highly efficient operating system for GES-C01 learning materials has won the praise of many customers. If you are determined to purchase our GES-C01 study tool, we can assure you that you can receive an email from our efficient system within 5 to 10 minutes after your payment, which means that you do not need to wait a long time to experience our learning materials. Then you can start learning our GES-C01 Exam Questions in preparation for the exam.

>> GES-C01 Latest Test Simulations <<

GES-C01 Test Collection Pdf | GES-C01 Latest Test Bootcamp

When you first contacted us with GES-C01 quiz torrent, you may be confused about our GES-C01 exam question and would like to learn more about our products to confirm our claims. We have a trial version for you to experience. If you encounter any questions about our GES-C01 Learning Materials during use, you can contact our staff and we will be happy to serve for you. As for any of your suggestions, we will take it into consideration, and effectively improve our GES-C01 exam question to better meet the needs of clients.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q338-Q343):

NEW QUESTION # 338

A developer is building a prototype Gen AI application using the Snowflake Cortex LLM REST API. During testing, they frequently encounter HTTP 429 'Too Many Requests' errors, particularly when rapidly invoking the 'COMPLETE endpoint. Which of the following statements accurately describes the usage quotas and troubleshooting steps for such rate limit events in the Cortex REST API?

- A. To resolve rate limiting, the developer should increase the virtual warehouse size for their API requests, as larger warehouses automatically increase API quotas.
- B. If the usage is below the RPM limit but still triggers a 429, the developer should review the token usage rate as TPM limits might be exceeded.
- C. Cortex REST API rate limits are primarily based on Requests per Minute (RPM) and do not consider the total Tokens Processed per Minute (TPM).
- D. Rate limits are model-specific, and exceeding either the Tokens Processed per Minute (TPM) or Requests per Minute (RPM) limits will result in an HTTP 429 response code.
- E. The HTTP 429 response indicates that the daily credit consumption for Snowflake Cortex AISQL usage has exceeded the 10-credit limit for on-demand accounts without a valid payment method, requiring conversion to a paid account.

Answer: B,D

Explanation:
Option C is correct because the Cortex REST API implements rate limits for both Tokens Processed per Minute (TPM) and Requests per Minute (RPM), and offending either of these limits will result in a 429 response code. Option D is correct as the sources explicitly state that if REST API usage is below the RPM limit but still receives a 429, the user should double-check the token usage rate, implying that TPM limits could be the cause. Option A is incorrect because while on-demand Snowflake accounts without a valid payment method are limited to 10 credits per day, an HTTP 429 specifically refers to rate limiting due to high usage, not a daily credit budget exceeding for the account type. Option B is incorrect as Cortex REST API rate limits consider both TPM and RPM. Option E is incorrect because Cortex REST API requests do not require a warehouse, and increasing warehouse size does not directly increase API quotas or resolve rate limiting for the REST API.

NEW QUESTION # 339
A data scientist is tasked with improving the accuracy of an LLM-powered chatbot that answers user questions based on internal company documents stored in Snowflake. They decide to implement a Retrieval Augmented Generation (RAG) architecture using Snowflake Cortex Search. Which of the following statements correctly describe the features and considerations when leveraging Snowflake Cortex Search for this RAG application?

- A. The
  

- B. Cortex Search automatically handles text chunking and embedding generation for the source data, eliminating the need for manual ETL processes for these steps.
- C. To create a Cortex Search Service, one must explicitly specify an embedding model and manually manage its underlying infrastructure, similar to deploying a custom model via Snowpark Container Services.
- D. Enabling change tracking on the source table for the Cortex Search Service is optional; the service will still refresh automatically even if change tracking is disabled.
- E. For optimal search results with Cortex Search, source text should be pre-split into chunks of no more than 512 tokens, even when using models with larger context windows like
  

Answer: A,B,E

Explanation:
Option A is correct because Cortex Search is a fully managed service that gets users started with a hybrid (vector and keyword) search engine on text data in minutes, without needing to worry about embedding, infrastructure maintenance, or index refreshes. Option B is incorrect because Cortex Search is a fully managed service; users do not need to manually manage the embedding model infrastructure. A default embedding model is used if not specified. Option C is correct because, for best search results with Cortex Search, Snowflake recommends splitting text into chunks of no more than 512 tokens, as smaller chunks typically lead to

higher retrieval and downstream LLM response quality, even with models that have larger context windows. Option D is correct because the 'SNOWFLAKE.CORTEX.SEARCH_PREVIEW' function allows users to test the search service to confirm it is populated with data and serving reasonable results for a given query. Option E is incorrect because change tracking is required on the source table for the Cortex Search Service to function correctly and reflect updates to the base data.

NEW QUESTION # 340
An ML engineer is developing a RAG application in Python and wants to use the TruLens SDK to trace the distinct phases of its execution, specifically the context retrieval and answer generation steps. They aim to clearly differentiate the tracing of the function responsible for retrieving context.

- A.
```
from trulens_core import instrument

@instrument(span_type='RETRIEVAL')
def retrieve_context(self, query: str) -> list:
    # ... retrieval logic ...
    return self._retrieve(query)
```

- B.
```
from trulens_core import instrument, SpanAttributes

@instrument(span_type=SpanAttributes.SpanType.RETRIEVAL)
def retrieve_context(self, query: str) -> list:
    """
    Retrieve relevant text from vector store.
    """
    # ... retrieval logic ...
    return self._retrieve(query)
```

- C.
```
from trulens_core import trace_function, SpanTypes

@trace_function(SpanTypes.RETRIEVAL)
def retrieve_context(self, query: str) -> list:
    # ... retrieval logic ...
    return self._retrieve(query)
```

- D.
```
from trulens_core import instrument

@instrument()
def retrieve_context(self, query: str) -> list:
    # ... retrieval logic ...
    return self._retrieve(query)
```

- E.
```
from trulens_core import instrument, SpanAttributes

@instrument(span_type=SpanAttributes.SpanType.GENERATION)
def retrieve_context(self, query: str) -> list:
    # ... retrieval logic ...
    return self._retrieve(query)
```

Answer: B

Explanation:
To instrument a function for context retrieval using the TruLens SDK and clearly differentiate its tracing, the decorator should be used with 'span_type=SpanAttributes.SpanType.RETRIEVAL'. This is directly demonstrated in the source for tracing a function with

a specific span type. Option B uses a string literal for 'span_type' , which is not the correct way to reference the enum member. Option C uses 'SpanAttributes.SpanType.GENERATION' , which is intended for LLM inference, not context retrieval. Option D uses the decorator without a specific 'span_type' , which would not clearly differentiate the context retrieval phase. Option E uses non- existent decorators and types(@trace_function', 'spanTypes').

NEW QUESTION # 341
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 D
- B. Option C
- C. Option B
- D. Option E
- E. Option A

Answer: E

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 # 342
A data scientist needs to fine-tune a 'mistral-7b' LLM using Snowflake Cortex for a specific text summarization task. They have prepared a training dataset in a Snowflake table, with text in a column named 'source_text' and the desired summaries in a column named 'expected_summary' . They also want to understand the cost implications. Which SQL statement will correctly initiate the fine-tuning job, and how will the cost be primarily calculated?

- A. The fine-tuning job is initiated by:
```
SELECT SNOWFLAKE.CORTEX.FINETUNE('CREATE', 'my_summarization_model', 'mistral-7b', 'SELECT source_text, expected_summary FROM my_training_data');
```
- B. The fine-tuning job is initiated by:
```
SELECT SNOWFLAKE.CORTEX.FINETUNE('CREATE', 'my_summarization_model', 'mistral-7b', 'SELECT source_text AS prompt, expected_summary AS completion
FROM my_training_data');
```
- C. The fine-tuning job is initiated by providing the prompt and completion data directly as arrays within the 'FINETUNE' function call, avoiding the need for a separate training data table, and costs are only for the storage of the fine-tuned model.
- D. The fine-tuning job is initiated by:
```
SELECT SNOWFLAKE.CORTEX.FINETUNE('TRAIN', 'my_summarization_model', 'mistral-7b', 'SELECT source_text AS prompt, expected_summary AS completion
FROM my_training_data');
```
- E. The fine-tuning job is initiated by:
```
CREATE SNOWFLAKE.ML.FINETUNE model my_summarization_model USING base_model='mistral-7b', training_data='my_training_data';
```

Answer: B

Explanation:
Option B is correct. The 'SNOWFLAKE.CORTEX.FINETUNE function requires the training data query result to include columns

named 'prompt' and 'completion'. Using SELECT ... AS prompt, ... AS completion" aliases the existing columns to the required names. The cost for fine-tuning is based on the number of tokens used in training, specifically calculated as 'Fine-tuning trained tokens = number of input tokens ' number of epochs trainecf. Additionally, running 'AI_COMPLETE on a fine-tuned model incurs compute costs for both input and output tokens processed. Option A is incorrect because the columns need to be aliased to 'prompt' and 'completion', and it misstates the cost calculation for fine-tuning training itself. Option C is incorrect because the syntax for FINETUNE is a SQL function, not a 'CREATE SNOWFLAKE.ML.FINETUNE' command, and Cortex Fine-tuning incurs compute costs based on tokens, not a fixed-rate subscription model. Option D is incorrect because the first argument for the function is 'CREATE', not 'TRAIN', and training costs are based on tokens, not GPU compute hours, although compute resources are utilized. Option E is incorrect because the training data must come from a Snowflake table or view, and costs are incurred for training and inference, not just storage.

NEW QUESTION # 343
.....

Our PDF version of the GES-C01 learning braindumps can print on papers and make notes. Then windows software of the GES-C01 exam questions, which needs to install on windows software. Also, the windows software is intelligent to simulate the real test environment. Then the online engine of the GES-C01 Study Materials, which is convenient for you because it doesn't need to install on computers. It supports Windows, Mac, Android, iOS and so on. This version just can run on web browser.

GES-C01 Test Collection Pdf: https://www.actualcollection.com/GES-C01-exam-questions.html

If you are willing to choose our GES-C01 test guide material, you will never feel disappointed about our products, Snowflake GES-C01 Latest Test Simulations Our workers will never randomly spread your information to other merchants for making money, Snowflake GES-C01 Latest Test Simulations So just click to pay for it, Our GES-C01 study tool boost three versions for you to choose and they include PDF version, PC version and APP online version.

According to what we provide, you can pass GES-C01 exam on your first try, It's called superposition in quantum mechanics, If you are willing to choose our GES-C01 Test Guide material, you will never feel disappointed about our products.

Useful GES-C01 Latest Test Simulations - Only in ActualCollection

Our workers will never randomly spread your GES-C01 Latest Test Simulations information to other merchants for making money, So just click to pay for it, Our GES-C01 study tool boost three versions for you to choose and they include PDF version, PC version and APP online version.

Challenge is ubiquitous, only by GES-C01 constant and ceaseless effort, can you be the man you want to be.

- www.prep4away.com GES-C01 Web-Based Practice Tests ✏ Search for [GES-C01] and download exam materials for free through ➤ www.prep4away.com □ □GES-C01 Official Study Guide
- Valid GES-C01 Test Vce □ GES-C01 Certified Questions □ GES-C01 Latest Braindumps □ Search for ➠ GES-C01 □ on 「 www.pdfvce.com 」 immediately to obtain a free download □Exam GES-C01 PDF
- GES-C01 Official Study Guide □ Valid GES-C01 Exam Topics □ GES-C01 Real Dumps □ Easily obtain ➡ GES-C01 □□□ for free download through ➠ www.passtestking.com □ □Pdf GES-C01 Exam Dump
- Free PDF 2025 GES-C01: SnowPro® Specialty: Gen AI Certification Exam Unparalleled Latest Test Simulations □ Open 【 www.pdfvce.com 】 and search for 【 GES-C01 】 to download exam materials for free □Exam GES-C01 Review
- Pdf GES-C01 Exam Dump □ GES-C01 Real Dumps □ Exam GES-C01 Simulations □ Open website { www.passtestking.com } and search for ☀ GES-C01 □☀□ for free download □GES-C01 Latest Braindumps
- Reliable GES-C01 Dumps Book □ Exam GES-C01 Review □ Valid GES-C01 Exam Topics □ Search on ➠ www.pdfvce.com □ for ✔ GES-C01 □✔□ to obtain exam materials for free download □Exam GES-C01 Review
- New GES-C01 Exam Review □ Valid GES-C01 Exam Topics □ Valid GES-C01 Test Vce □ Open [www.pdfdumps.com] enter ▹ GES-C01 ◃ and obtain a free download □GES-C01 Test Sample Online
- Exam GES-C01 Study Guide □ Valid GES-C01 Test Vce □ Exam GES-C01 Review □ Simply search for ➽ GES-C01 □ for free download on [www.pdfvce.com] □GES-C01 Reliable Study Notes
- GES-C01 PDF Questions - Perfect Prospect To Go With GES-C01 Practice Exam □ Search on " www.torrentvce.com " for { GES-C01 } to obtain exam materials for free download □GES-C01 Official Study Guide
- Valid GES-C01 Latest Test Simulations - Find Shortcut to Pass GES-C01 Exam □ Search for □ GES-C01 □ and easily obtain a free download on 【 www.pdfvce.com 】 □Exam GES-C01 PDF
- GES-C01 Latest Braindumps □ GES-C01 Certified Questions □ GES-C01 Valid Test Questions □ Immediately open ➠ www.torrentvalid.com □ and search for ▹ GES-C01 ◃ to obtain a free download □Exam GES-C01 Study Guide

- myportal.utt.edu.tt, frenchcoachingacademy.education, myportal.utt.edu.tt, www.stes.tyc.edu.tw, myportal.utt.edu.tt, Disposable vapes