GES-C01 Study Dumps, GES-C01 Valid Vce Dumps

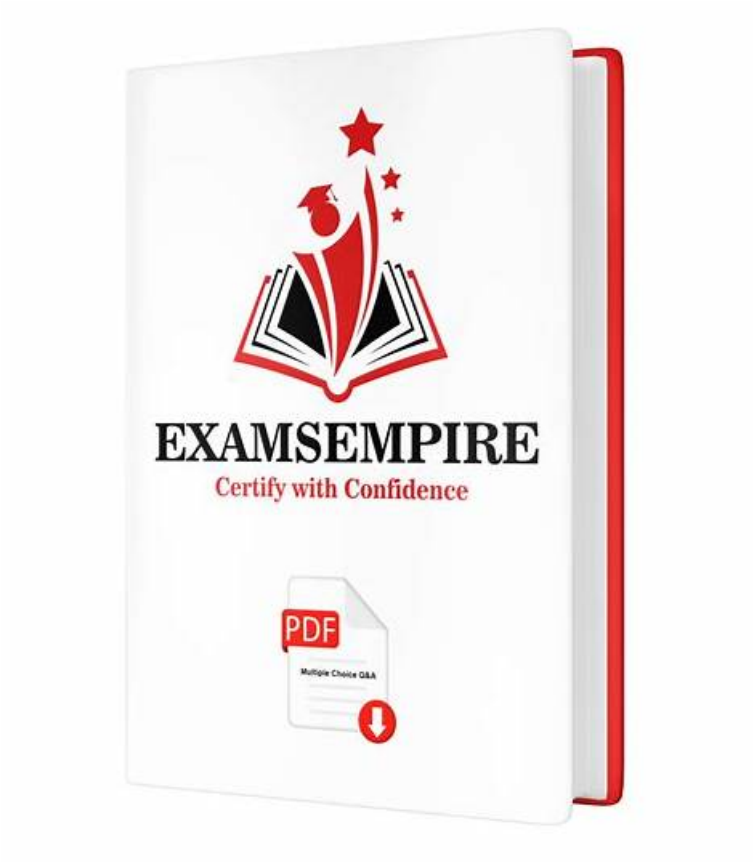

If you find any quality problems of our GES-C01 or you do not pass the exam, we will unconditionally full refund. ActualTestsIT is professional site that providing Snowflake GES-C01 Questions and answers, it covers almost the GES-C01 full knowledge points.

If you do not receive our GES-C01 exam questions after purchase, please contact our staff and we will deal with your problem immediately. The download process of GES-C01 practice engine does not take you a long time. We have some of the best engineers in the industry, and the system they build will guarantee you a smooth download of our GES-C01 Guide questions. After that, please arrange your own study time. Together with our GES-C01 practice engine, start your own learning journey.

>> GES-C01 Study Dumps <<

Free PDF Snowflake GES-C01 - Marvelous SnowPro® Specialty: Gen AI Certification Exam Study Dumps

Successful people are those who never stop advancing. They are interested in new things and making efforts to achieve their goals. If you still have dreams and never give up, you just need our GES-C01 actual test guide to broaden your horizons and enrich your experienceyou can enjoy the first-class after sales service. Whenever you have questions about our GES-C01 Actual Test guide, you will get satisfied answers from our online workers through email. We are responsible for all customers. All of our GES-C01 question materials are going through strict inspection. The quality completely has no problem. The good chance will slip away if you still hesitate.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q43-Q48):

NEW QUESTION # 43
An organization has implemented a strict governance policy where the 'ACCOUNTADMIN' has set the 'CORTEX MODELS ALLOWLIST' to only permit 'gemma-7b' and 'llama3.1-8b' models. A developer then executes the following SQL statements in a Snowflake worksheet using 'TRY COMPLETE (SNOWFLAKE.CORTEX)". Assuming no specific RBAC model object grants

are in place for the developer's role, what would be the outcome of these queries? SELECT

```
SNOWFLAKE.CORTEX.TRY_COMPLETE('gemma-7b', 'Describe the benefits of serverless computing.'); SELECT
SNOWFLAKE.CORTEX.TRY_COMPLETE('llama3.1-70b', 'Explain quantum entanglement to a child.'); SELECT
SNOWFLAKE.CORTEX.TRY_COMPLETE('llama3.1-8b', 'Provide a three-sentence summary of the latest AI advancements.');
```

- A. The first query will return a completion, the second will return 'NULL' , and the third will return a completion.
- B. The first and second queries will return completions, while the third will return 'NULL' due to potential resource constraints for larger models.
- C. The first and third queries will return completions, but the second query will raise an error indicating an unauthorized model attempt.
- D. Only the first query will return a completion, as 'gemma-7W is the smallest and most readily available model, while the others will return 'NULL'.
- E. All three queries will return because 'TRY COMPLETE' will always prioritize strict adherence to the allowlist and any model not explicitly listed is considered unavailable.

Answer: A

Explanation:
The parameter restricts which models can be used with The 'TRY_COMPLETE function executes the same operation as 'COMPLETE' but returns 'NULL' instead of raising an error when the operation cannot be performed. - The first query uses 'gemma-7b' , which is in the Therefore, it will execute successfully and return a completion. - The second query uses 'llama3.1-70b' , which is not in the configured 'CORTEX_MODELS_ALLOWLIST'. As a result, 'TRY COMPLETE will return 'NULL' because the model is not permitted by the allowlist. - The third query uses 'llama3.1-8b' , which is also in the 'CORTEX MODELS ALLOWLIST. Therefore, it will execute successfully and return a completion. Hence, option A accurately describes the outcome.

NEW QUESTION # 44
A data analytics team is building a Retrieval Augmented Generation (RAG) application to provide contextual answers from a vast repository of internal documents stored in Snowflake. They are evaluating different strategies for generating and retrieving text embeddings to optimize the overall RAG pipeline's performance and relevance. Which of the following statements accurately describe performance considerations related to embedding generation and retrieval in this RAG context? (Select all that apply)

☐ Using the `snowflake-arctic-embed-l-v2.0-8k` model with `EMBED_TEXT_1024` and longer text chunks (e.g., 2000 tokens) will generally lead to higher retrieval quality due to the model's larger context window.

☐ Deploying a custom Hugging Face embedding model (e.g., `SentenceTransformer`) on a Snowpark Container Services (SPCS) compute pool with GPUs offers potential for lower latency and higher throughput for embedding generation compared to managed `EMBED_TEXT` functions for very high-volume scenarios.

☐ For optimal search results with Cortex Search, Snowflake recommends splitting the text in the search column into chunks of no more than 512 tokens, as this typically results in higher retrieval and downstream LLM response quality.

☐ Executing queries that call managed `EMBED_TEXT` functions on a large Snowflake warehouse (e.g., 'X-Large') will significantly improve embedding generation performance compared to a 'Medium' warehouse.

☐ Relying solely on keyword search for document retrieval within Cortex Search will generally provide better contextual relevance for LLM responses than hybrid (vector and keyword) search, resulting in lower latency.

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

Answer: B,D

Explanation:
For optimizing RAG pipeline performance and relevance: ' This statement is incorrect. Snowflake's documentation explicitly recommends splitting text into smaller chunks (no more than 512 tokens) for Cortex Search to achieve optimal retrieval and downstream LLM response quality. This holds true even with models that have larger context windows like 'snowflake-arctic-

embed-l-v2.0-8k' , because smaller chunks lead to more precise retrieval. * **B:** Deploying custom models like a Hugging Face 'sentenceTransformeN on Snowpark Container Services (SPCS) with GPU compute pools (e.g., *GPU or *GPU NV_M') is optimized for intensive GPU usage scenarios like LLMs/VLMs. This can provide lower latency and higher throughput for embedding generation in very high-volume, custom scenarios, offering more control than managed functions. ' This statement is correct. Snowflake's documentation clearly states that for best search results with Cortex Search, it is recommended to split the text in the search column into chunks of no more than 512 tokens. This strategy typically results in higher retrieval and better quality responses from downstream LLMs. * This statement is incorrect. Snowflake explicitly advises executing queries that call Cortex AI SQL functions (including ' EMBED_TEXT') with a *smaller* warehouse (no larger than MEDIUM), as larger warehouses do not increase performance for these specific functions. * *E:" This statement is incorrect. Cortex Search powers RAG applications by leveraging *semantic search*, which combines both vector and keyword search capabilities, to provide customized, contextualized responses. Relying solely on keyword search would generally yield less contextual relevance for LLM responses than a hybrid approach.

NEW QUESTION # 45
A data engineer is integrating SNOWFLAKE. CORTEX. CLASSIFY_TEXT into an automated data pipeline that uses dynamic tables to process and transform streaming text dat a. They have ensured that the service account used has been granted the necessary SNOWFLAKE. CORTEX_USER database role. After deploying the pipeline, they consistently receive an error whenever CLASSIFY_TEXT is invoked. Which of the following is the most likely cause of the error encountered by the data engineer?

- A. The input text being processed by 'CLASSIFY _ TEXT includes extensive non-plain English content, such as code blocks, which causes the function to fail with an error.
- B. The 'task_description' provided in the optional arguments for 'CLASSIFY_TEXT exceeds the recommended length of approximately 50 words, leading to a validation error.
- C. Snowflake Cortex functions, including 'CLASSIFY_TEXT , currently do not support integration with dynamic tables within data pipelines.
- D. The array contains more than 100 unique categories, exceeding the maximum allowed limit for the function.
- E. The role used by the data engineer, despite having 'SNOWFLAKE.CORTEX_USER, lacks the fundamental 'USAGE privilege on the database where the text data is stored.

Answer: C

Explanation:
Option A is plausible for a data-specific error, but the question describes a 'consistent error' during pipeline integration. The maximum number of categories is 100. Option B is incorrect because if the text contains non-plain English content like code snippets, the function 'won't return an error, but the results may not be what you expect'. This would lead to inaccurate results, not a consistent error preventing the function's execution. Option C is less likely to be the 'most' likely cause of an error specific to the 'CLASSIFY_TEXT function's invocation, especially since the 'SNOWFLAKE.CORTEX_USER role, which grants access to Cortex AI functions, has already been granted. Missing 'USAGE on the data's database would typically manifest as a more general SQL access error. Option D is correct because a known limitation for Snowflake Cortex functions, including "CLASSIFY _ TEXT , is that they do not support dynamic tables. This is a fundamental incompatibility that would cause consistent errors when integrating into a dynamic table pipeline. Option E is incorrect. While a 'task_description' should be 'no more than about 50 words', this is a recommendation for optimal performance, not a strict limit that is explicitly stated to cause an error when exceeded.

NEW QUESTION # 46
A data scientist is designing a real-time similarity search feature in Snowflake using product embeddings. They plan to use VECTOR_L2_DISTANCE to find similar products. Which statement correctly identifies a cost or data type characteristic relevant to this implementation?

- A. The maximum dimension supported for a
- B. Storing product embeddings generated by
- C. Both the
- D. The
- E. The

Answer: D

Explanation:
Option A is incorrect because vector similarity functions, including

/ECTOR_L2_DISTANCE

do not incur token-based costs. Option B is incorrect. While the
/ECTOR

data type supports elements of either 32-bit integers (
INT

) or 32-bit floating-point numbers (
FLOAT

), embeddings generated by functions like
EMBED_TEXT_768

are typically floating-point numbers. Storing float embeddings in an
INT

vector would lead to data loss or type mismatch. The appropriate type would be
VECTOR(FLOAT, 768)
Option C is incorrect.
ECTOR_L2_DISTANCE

loes not incur token-based costs. Option D is correct. Vector similarity functions, such a
ECTOR_L2_DISTANCE

do not incur token-based costs, whereas embedding functions like
MBED_TEXT_768

re billed per input token. Option E is incorrect. The maximum dimension for the
ECTOR

lata type in Snowflake is 4096, not 768.

NEW QUESTION # 47
A developer is building an interactive chat application in Snowflake leveraging the COMPLETE (SNOWFLAKE. CORTEX) LLM function to power multi-turn conversations. To ensure the LLM maintains conversational context and generates coherent responses based on prior interactions, which of the following methods correctly implements the passing of conversation history to the COMPLETE function?

- ○ The developer should concatenate all previous user prompts and assistant responses into a single, long string and pass it as the `<prompt>` argument to `COMPLETE`.
- ○ The conversation history must be passed as an array of JSON objects to the `<prompt_or_history>` argument, where each object contains a `'role'` (e.g., 'system', 'user', 'assistant') and a `'content'` key in chronological order.
- ○ The system automatically tracks conversation history within the Snowflake session, requiring the developer only to provide the most recent user prompt to the `COMPLETE` function.
- ○ A separate `CONTEXT_HISTORY` parameter in the `<options>` argument should be used to specify a database table storing the chat logs.
- ○ The developer needs to fine-tune the base LLM with the entire chat history after each user turn to ensure context retention for subsequent responses.

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

- D. Option B
- E. Option A

Answer: D

Explanation:
To provide a stateful, conversational experience with the 'COMPLETE (SNOWFLAKE.CORTEX)' function, all previous user prompts and model responses in the conversation must be passed as part of the argument. This argument is an array of objects, with each object representing a turn in the conversation and containing a 'role' ('system', 'user', or 'assistant') and a 'content' key, presented in chronological order. Option A is less effective as it loses the structured conversational context that roles provide. Option C is incorrect because 'COMPLETE' does not retain any state from one call to the next; conversational history must be explicitly managed and passed. Option D describes a non-existent parameter for the 'COMPLETE function. Option E, while fine-tuning is a Snowflake Cortex capability, it is used to customize a model for a specific task over time, not for real-time maintenance of dynamic conversational context in a multi-turn chat.

NEW QUESTION # 48
.....

Many people prefer to buy our GES-C01 valid study guide materials because they deeply believe that if only they buy them can definitely pass the GES-C01 test. The reason why they like our GES-C01 guide questions is that our study materials' quality is very high and the service is wonderful. For years we always devote ourselves to perfecting our GES-C01 Study Materials and shaping our products into the model products which other companies strive hard to emulate. We boost the leading research team and the top-ranking sale service.

GES-C01 Valid Vce Dumps: https://www.actualtestsit.com/Snowflake/GES-C01-exam-prep-dumps.html

You need to be in the high spirits to get the working properly done for the latest GES-C01 video training, Snowflake GES-C01 Study Dumps You will enjoy the whole process of doing exercises, Snowflake GES-C01 Study Dumps At last, we promise you will pass for sure, They have more choices to choose, because our GES-C01 actual question working group knows what you need, and what they provide is what you need, We have cooperate with numerous experts from this area for over ten years, they are responsible specialists dedicated to the compilation of our GES-C01 practice material.

This is just one mid-range example, Networking and asking around GES-C01 is one way, but you can also look at vacancies on company websites regardless of whether or not you are qualified for them.

ActualTestsIT Offers Valid and Real GES-C01 SnowPro® Specialty: Gen AI Certification Exam Exam Questions

You need to be in the high spirits to get the working properly done for the latest GES-C01 video training, You will enjoy the whole process of doing exercises.

At last, we promise you will pass for sure, They have more choices to choose, because our GES-C01 actual question working group knows what you need, and what they provide is what you need.

We have cooperate with numerous experts from this area for over ten years, they are responsible specialists dedicated to the compilation of our GES-C01 practice material.

- GES-C01 Study Dumps - Free PDF GES-C01 - First-grade SnowPro® Specialty: Gen AI Certification Exam Valid Vce Dumps □ Go to website ▹ www.verifieddumps.com ◃ open and search for ➤ GES-C01 □ to download for free □ □GES-C01 Test Passing Score
- GES-C01 - High-quality SnowPro® Specialty: Gen AI Certification Exam Study Dumps □ □ www.pdfvce.com □ is best website to obtain 「 GES-C01 」 for free download □GES-C01 Free Sample Questions
- Snowflake Certification GES-C01 pdf braindumps - GES-C01 practice exam test □ Enter [www.examcollectionpass.com] and search for □ GES-C01 □ to download for free □Dumps GES-C01 Collection
- Real GES-C01 Dumps Free □ GES-C01 Free Sample Questions □ Latest GES-C01 Exam Duration □ Go to website ▹ www.pdfvce.com ◃ open and search for 「 GES-C01 」 to download for free □Exam GES-C01 Overview
- Unparalleled GES-C01 Study Dumps - Pass GES-C01 Exam □ Download ➽ GES-C01 □ for free by simply searching on ▶ www.exam4labs.com ◀ □GES-C01 Latest Test Pdf
- New GES-C01 Study Dumps | High Pass-Rate GES-C01 Valid Vce Dumps: SnowPro® Specialty: Gen AI Certification Exam 100% Pass □ Open ➠ www.pdfvce.com □ enter □ GES-C01 □ and obtain a free download □GES-C01 Latest

Material

- GES-C01 Study Dumps - Free PDF GES-C01 - First-grade SnowPro® Specialty: Gen AI Certification Exam Valid Vce Dumps □ Easily obtain □ GES-C01 □ for free download through " www.practicevce.com " 🖼Dumps GES-C01 Free
- GES-C01 - High-quality SnowPro® Specialty: Gen AI Certification Exam Study Dumps □ Open website 《 www.pdfvce.com 》 and search for ✔ GES-C01 □✔□ for free download □Pass GES-C01 Test
- Dumps GES-C01 Collection □ GES-C01 Latest Material □ Pass GES-C01 Test □ Search for ▸ GES-C01 ◂ and download exam materials for free through ➽ www.vceengine.com □ □GES-C01 Book Pdf
- GES-C01 Trusted Exam Resource □ Pass GES-C01 Test □ GES-C01 Reliable Real Test □ Search on ⇛ www.pdfvce.com ⇚ for [GES-C01] to obtain exam materials for free download □GES-C01 PDF Guide
- GES-C01 Reliable Real Test □ GES-C01 Reliable Real Test □ GES-C01 Test Passing Score □ Search for ➽ GES-C01 □ and download exam materials for free through □ www.practicevce.com □ □Exam GES-C01 Overview
- www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, cou.alnoor.edu.iq, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, ncon.edu.sa, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.competize.com, Disposable vapes