Snowflake GES-C01 Dumps - Obtain Brilliant Result (2026)

BONUS!!! Download part of TroytecDumps GES-C01 dumps for free: https://drive.google.com/open?id=1y4bXvL4Nro5ujw-2FDX5JeVKohhXHHCq

While making revisions and modifications to the Snowflake GES-C01 practice exam, our team takes reports from over 90,000 professionals worldwide to make the Snowflake GES-C01 Exam Questions foolproof. To make you capable of preparing for the GES-C01 exam smoothly, we provide actual Snowflake GES-C01 exam dumps.

Our APP online version of GES-C01 exam questions has the advantage of supporting all electronic equipment. You just need to download the online version of our GES-C01 preparation dumps, and you can use our GES-C01 study quiz by any electronic equipment. We can promise that the online version will not let you down. We believe that you will benefit a lot from it if you buy our GES-C01 training materials.

>> GES-C01 Certificate Exam <<

Latest SnowPro® Specialty: Gen AI Certification Exam real exams, GES-C01 vce dumps

As for Snowflake GES-C01 exam, it is the most difficult to pass. But, as long as you believe in TroytecDumps, everything is ok. TroytecDumps Snowflake GES-C01 exam simulations contain the most accurate questions and answers. If you don't believe our Snowflake GES-C01 certification training, you can go to our TroytecDumps. You can find pdf real questions and answers and download it. And the purchase rate is unbelievably high every day. By choosing it, pass rate is 100%. Hurry up! Don't hesitate to add our Snowflake GES-C01 Dumps Torrent to your shopping cart.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q322-Q327):

NEW QUESTION # 322
A data engineering team is designing a scalable data pipeline in Snowflake that involves processing large text inputs with Cortex AI LLM functions. They want to ensure cost efficiency and prevent queries from failing due to exceeding LLM context window limits. They plan to use SNOWFLAKE. CORTEX. COUNT_TOKENS for pre-validation. Which of the following statements are TRUE about the role and cost of COUNT_TOKENS in this scenario? (Select all that apply)

☐ Using
`COUNT_TOKENS`

to estimate tokens for an embedding model like
`snowflake-arctic-embed-m-v1.5`

will help ensure the input text does not exceed its 512-token context window, thus preventing truncation or unexpected behavior.

☐ The
`COUNT_TOKENS`

function directly contributes to the overall token-based billing of the LLM inference call it precedes.

☐ For optimal retrieval quality in RAG scenarios,
`COUNT_TOKENS`

can help facilitate splitting text into smaller chunks, ideally no more than 512 tokens, even for models with larger context windows.

☐ The compute cost for running
`COUNT_TOKENS`

is entirely independent of the length of the input text, as it only reflects the cost of function invocation.

☐ The
`COUNT_TOKENS`

function is universally available in all Snowflake regions and supports token counting for any Cortex model, irrespective of that model's specific regional availability for other inference functions.

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

Answer: C,D,E

Explanation:
Option A is correct. Embedding models like
`snowflake-arctic-embed-m-v1.5`
have a fixed context window of 512 tokens. Using COUNT_TOKENS allows pre-checking if text fits within this limit, preventing truncation that can occur when the input exceeds the context window. Option B is incorrect because COUNT_TOKENS incurs only compute cost to run the function and does not incur additional token-based costs that would add to the billing of subsequent LLM inference calls. Option C is correct. For best search results, Snowflake recommends splitting text into chunks of no more than 512 tokens. This practice generally leads to higher retrieval and downstream LLM response quality in RAG scenarios, and COUNT_TOKENS is a valuable tool for managing these chunk sizes. Option D is incorrect. While COUNT_TOKENS incurs compute cost (not token-based cost), the amount of compute would generally scale with the size of the input text it processes, making the cost not entirely independent of input length, although it's not billed on a per-token basis for its own operation. Option E is correct. The COUNT_TOKENS function is available in all regions for any model, though the models themselves may have specific regional availabilities for other functions.

NEW QUESTION # 323
An ML engineer has developed a custom PyCaret classification model and wants to deploy it to Snowpark Container Services (SPCS) for inference using the Snowflake Model Registry. The model requires specific versions of pycaret' , 'scipy', and 'joblib'. The engineer also wants to make the service accessible via an HTTP endpoint. Which of the following Model Registry and service creation steps are 'most appropriate' for the ML engineer? (Select all that apply.)

- A.

  Log the model using `reg.log_model`, specifying `conda_dependencies` which would be resolved from the **Snowflake Anaconda channel for warehouse deployment**, and setting `target_platforms` to `["SNOWPARK_CONTAINER_SERVICES"]`.

- B.

  Create the service using `mv.create_service`, ensuring `ingress_enabled = True` and specifying `gpu_requests = "1"` if the compute pool is GPU-enabled.

- C.

```
Ensure the `ModelContext` specifies the `model_file` attribute pointing to the serialized PyCaret model, for example:
pycaret_model_context = custom_model.ModelContext(
    model_file='pycaret_best_model.pkl'
);
```

- D.

  [illegible]

- E. Opt for warehouse deployment instead of SPCS, as PyCaret is not natively supported by Snowflake and managing its dependencies in SPCS would be overly complex compared to a warehouse.

Answer: B,C,D

Explanation:
Option A is correct. When bringing an unsupported model type, such as PyCaret, you must define a 'ModelContext' that refers to the serialized model file (e.g., a pickled file). Option B is incorrect. For models deployed to Snowpark Container Services, 'conda_dependencies' are, by default, obtained from 'conda-forge' , not the Snowflake Anaconda channel, which is used for warehouse deployments. Therefore, relying on the Snowflake Anaconda channel for SPCS deployment is incorrect. Option C is correct. While 'conda_dependencies' can be used for SPCS (resolved from 'conda-forge'), 'pip_requirementS are often a more direct and reliable way to specify dependencies for custom or less common third-party Python packages, ensuring they are pulled directly from PyPI if not available in 'conda-forge' . The PyCaret example in the sources, while using 'conda_dependencies' , represents a specific case, and for broader 'custom third-party packages', pip is a strong choice. Option D is correct. To make the deployed service accessible via an HTTP endpoint, must be set to 'True'. Additionally, 'gpu_requests = (or the appropriate number of GPUs) is essential when deploying a model to a GPU compute pool to ensure it leverages the GPU resources for inference. Option E is incorrect. Snowpark Container Services is specifically designed to ease the restrictions of warehouse deployment, allowing for the use of any packages (including PyPI) and enabling large models to run on distributed clusters of GPUs, which is ideal for this scenario.

NEW QUESTION # 324
A Data Application Developer is building a Streamlit chat application powered by Snowflake Cortex Analyst. Users frequently ask questions involving specific product names, such as "What was the total sales of 'Luxury Coffee Beans' last quarter?". The semantic model has a product_name dimension with high cardinality. The developer wants to ensure Cortex Analyst accurately identifies these specific product literals in user queries. Given this scenario, which of the following approaches should the developer consider to optimize literal search capabilities and enhance Cortex Analyst responses?

- Configure the `product_name` dimension in the semantic model to include a `sample_values` array with a comprehensive list of all possible product names.
- Create a Cortex Search Service on the underlying `product_name` column and specify this service in the `cortex_search_service` field of the `product_name` dimension in the semantic model.
- Employ the `AI_COMPLETE` function within the chat application logic to pre-process user queries and extract product names before sending them to Cortex Analyst.
- Add a `verified_query` entry to the semantic model for each specific product name, detailing its sales query.
- Utilise the `AI_PARSE_DOCUMENT` function to extract product names from a reference document containing all product listings and feed them into the semantic model.

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

Answer: B

Explanation:
To improve literal search capabilities for Cortex Analyst, especially with high-cardinality dimensions like product names, integrating with Cortex Search Services is the recommended approach. Cortex Search provides low-latency, high-quality "fuzzy" search over text data, enabling semantic search to find literal values for Cortex Analyst's SQL queries. This integration is supported by specifying the Cortex Search Service in the field of the dimension definition within the semantic model. Option A is not ideal because cortex search service sample _ values are recommended for low-cardinality dimensions (e.g., 1-10 distinct values). A high-cardinality dimension like product names would make this unmanageable and less effective. Option C is incorrect; AI_COMPLETE is a general LLM completion function and not designed for pre-processing queries to extract structured entities for Cortex Analyst's text-to-SQL functionality. Option D is impractical and unscalable for high-cardinality data, as it would require creating a vast number of entries. Option E, while can extract information from documents, verified_query AI_PARSE_DOCUMENT it is not the designated or most efficient method to provide literal values for semantic matching within Cortex Analyst's dimensions; Cortex Search Services are specifically built for this purpose.

NEW QUESTION # 325
A global marketing team uses Snowflake to manage customer feedback in various languages. They need to translate customer reviews from German ("de") into English ("en") for analysis. The reviews are stored in a table named 'CUSTOMER REVIEWS' in a column called 'REVIEW TEXT'. Which of the following SQL statements correctly applies the 'SNOWFLAKE.CORTEX.TRANSLATE function and what is the expected return type for the translated text?

- A. The query

```
SELECT SNOWFLAKE.CORTEX.TRANSLATE(review_text, 'English') FROM CUSTOMER_REVIEWS;
```

- B. The query

```
SELECT SNOWFLAKE.CORTEX.TRANSLATE(review_text, 'de', 'en') FROM CUSTOMER_REVIEWS;
```

- C. The query

[illegible]

- D. The query

```
SELECT SNOWFLAKE.CORTEX.TRANSLATE(review_text, 'en') FROM CUSTOMER_REVIEWS;
```

- E. The query

```
SELECT SNOWFLAKE.CORTEX.TRANSLATE(review_text, 'de', 'en', 'high_accuracy') FROM CUSTOMER_REVIEWS;
```

Answer: B

Explanation:
Option B is correct. The 'SNOWFLAKE.CORTEX.TRANSLATE function takes three arguments: the text to be translated, the source language, and the target language. It returns a STRING value containing the translated text. Option A includes an unsupported 'high_accuracy' option and claims an incorrect return type. Option C uses an incorrect syntax and claims an incorrect return type. Option D uses an incorrect number of arguments. Option E claims an incorrect return type.

NEW QUESTION # 326
A data engineering manager needs to audit Cortex LLM function costs to identify specific SQL queries that are unexpectedly high in token consumption for the 'llama3.1-8b' model. They require granular analysis of prompt, completion, and guardrail token usage for these queries. Which of the following Snowflake methods or views would provide the necessary insights?

☐
Querying the `SNOWFLAKE.ORGANIZATION_USAGE.METERING_DAILY_HISTORY` view, filtered by `SERVICE_TYPE = 'AI_SERVICES'`, to identify daily aggregated credit usage for all AI services.

☐
Examining the `SNOWFLAKE.ACCOUNT_USAGE.CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY` view, specifically filtering by `MODEL_NAME = 'llama3.1-8b'` and analyzing the `PROMPT_TOKENS`, `COMPLETION_TOKENS`, `GUARD_TOKENS`, and `QUERY_ID` columns.

☐
Utilizing the `SNOWFLAKE.CORTEX.COUNT_TOKENS` function to estimate token counts for input prompts before any LLM function execution, comparing these estimates to actual spend.

☐
For SQL queries using `AI_COMPLETE` or `COMPLETE` with `show_details => TRUE`, inspecting the `usage` object within the returned JSON to retrieve `prompt_tokens`, `completion_tokens`, and `guard_tokens` for each individual call.

☐
Analyzing the `SNOWFLAKE.ACCOUNT_USAGE.CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY` view, filtering by the specific Document AI model builds associated with the 'llama3.1-8b' model.

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

Answer: A,D

Explanation:

Option B is correct because the `SNOWFLAKE.ACCOUNT_USAGE.CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY` view provides granular usage information for individual Cortex LLM function calls, including `PROMPT_TOKENS`, `COMPLETION_TOKENS`, `GUARD_TOKENS`, and the `QUERY_ID` for specific queries and `MODEL_NAME`. This directly addresses the need to audit token consumption for a specific model and identify high-usage queries. Option D is also correct as the `COMPLETE` and `AI_COMPLETE` functions, when called with `show_details => TRUE` (available via the `options` argument for `COMPLETE`), return a JSON object that includes a `usage` key with `prompt_tokens`, `completion_tokens`, and `guard_tokens` details for that specific invocation. This provides per-call details directly at the point of execution. Option A is incorrect because `METERING_DAILY_HISTORY` provides aggregated daily credit usage for all AI services, not granular token counts per query or model. Option C is incorrect as `COUNT_TOKENS` is used for estimating token counts *before* execution to avoid exceeding limits or for cost planning, not for tracking *actual* historical usage. Option E is incorrect because `CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY` tracks Document AI processing functions like `!PREDICT` and `PARSE_DOCUMENT`, and aggregates pages processed and credits used, not granular token counts for general LLM functions like `COMPLETE`. The `llama3.1-8b` model is an LLM available for `AI_COMPLETE` / `COMPLETE`.

NEW QUESTION # 327
......

After cracking the SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam you will receive the credential badge. It will pave your way toward well-paying jobs or promotions in any reputed tech company. At TroytecDumps have customizable SnowPro® Specialty: Gen AI Certification Exam (GES-C01) practice exams for the students to review and improve their preparation. The SnowPro® Specialty: Gen AI Certification Exam (GES-C01) practice test material product of TroytecDumps are created by experts with the dedication to help customers crack the SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam on the first attempt.

New GES-C01 Exam Online: https://www.troytecdumps.com/GES-C01-troytec-exam-dumps.html

Our GES-C01 exam questions are the Best choice in terms of time and money, Snowflake GES-C01 Certificate Exam We also offer you a 100% money-back guarantee, in the very rare case of failure or unsatisfactory results, Latest TroytecDumps GES-C01 braindumps will definitely fascinate you with the select number of important questions and answers, Snowflake GES-C01 Certificate Exam There are many certificates for you to get but which kind of certificate is most authorized, efficient and useful?

Check Permissions on Files and List Items, in order for clients GES-C01 to enroll using I, you will need to set permissions accordingly, as by default, only domain admins can access the console.

Our GES-C01 Exam Questions are the Best choice in terms of time and money, We also offer you a 100% money-back guarantee, in the very rare case of failure or unsatisfactory results.

Snowflake GES-C01 Preparation Materials Available In Different Formats

Latest TroytecDumps GES-C01 braindumps will definitely fascinate you with the select number of important questions and answers, There are many certificates for you GES-C01 Certificate Exam to get but which kind of certificate is most authorized, efficient and useful?

The rest of the time you can do anything GES-C01 Certificate Exam you want to do to,which can fully reduce your review pressure.