GES-C01 New Exam Materials & GES-C01 Most Reliable Questions

- ☐ For `RAG_Config_A`, instrument the `generate_answer` function with @instrument(span_type=SpanAttributes.SpanType.GENERATION). For `RAG_Config_B`, instrument its equivalent `generate_answer` function similarly, and ensure both are registered as part of distinct `TruApp` versions or runs for comparison.
- ☐ Enable cross-region inference using the CORTEX_ENABLED_CROSS_REGION parameter to ensure both `llama3.1-8b` and `mistral-7b` models are available, as this directly enables the comparison feature within AI Observability.
- ☐ Instrument the context retrieval component in both configurations with @instrument(span_type=SpanAttributes.SpanType.RETRIEVAL) to allow for the calculation of `context_relevance` metrics for each, which can then be used in comparative evaluations.
- ☐ Create separate runs (using add_run()) with distinct run_name or label in the RunConfig) for each RAG configuration, specifying the respective LLM as llm_judge_name, and explicitly list `answer_relevance` and `groundedness` in the metrics parameter when calling compute_metrics().
- ☐ Focus solely on `prompt_tokens` and `completion_tokens` via the `CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY` view, as these metrics provide the most direct comparison of LLM performance for RAG applications.

DOWNLOAD the newest TestSimulate GES-C01 PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1dn2rA8kdSbGjc0TOPJFSv4cNSWBetgoY

Don't waste time, buy the latest GES-C01 pdf questions and practice tests from TestSimulate and get successful. You can free download the demo of any format of Snowflake GES-C01 test questions before purchase. You can claim a refund if you don't pass the Snowflake GES-C01 Certification Exam after using these actual Snowflake GES-C01 exam dumps.

You must have felt the changes in the labor market. Today's businesses require us to have more skills and require us to do more in the shortest possible time. We are really burdened with too much pressure. GES-C01 simulating exam may give us some help. With our GES-C01 Study Materials, we can get the GES-C01 certificate in the shortest possible time. And our pass rate is high as 98% to 100% which is unbeatable in the market.

>> GES-C01 New Exam Materials <<

Get Snowflake GES-C01 Dumps For Quick Preparation [2025]

We aim to leave no misgivings to our customers so that they are able to devote themselves fully to their studies on GES-C01 guide materials and they will find no distraction from us. I suggest that you strike while the iron is hot since time waits for no one. With our GES-C01 Exam Questions, you will be bound to pass the exam with the least time and effort for its high quality. With our GES-C01 study guide for 20 to 30 hours, you will be ready to take part in the exam and pass it with ease.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q216-Q221):

NEW QUESTION # 216
A new data analyst is trying to incorporate sentiment analysis using SNOWFLAKE. CORTEX. SENTIMENT within a Snowflake data pipeline that uses dynamic tables. They execute the following SQL to create a dynamic table for daily sentiment aggregation:

```
REATE OR REPLACE DYNAMIC TABLE daily_sentiments
ARGET_LAG = '1 hour'
AREHOUSE = my_analytics_wh
S
ELECT
eview_date,
NOWFLAKE.CORTEX.SENTIMENT(review_content) AS review_sentimen
ROM
roduct_reviews;
```


However, this operation fails. Which of the following is the most direct reason for the failure of this specific setup?

- A. The warehouse my_analytics_wh is likely not a Snowpark-optimized warehouse, which is a requirement for Cortex functions within dynamic tables.
- B. The CORTEX_USER database role was not granted to the analyst's role, preventing the execution of Cortex functions.
- C. SNOWFLAKE. CORTEX. SENTIMENT and other Snowflake Cortex functions are currently incompatible with dynamic tables.
- D. The review_content column, if containing non-English text, would cause the SENTIMENT function to fail outright rather

than produce inaccurate results.
- E. The TARGET_LAG for dynamic tables must be explicitly set to '1 day' or longer when integrating with Cortex functions.

Answer: C

Explanation:
Option B is correct. Snowflake Cortex functions, including SNOWFLAKE .CORTEX. SENTIMENT, do not support dynamic tables. This is a fundamental limitation that would cause the CREATE DYNAMIC TABLE statement to fail when trying to incorporate a Cortex function. While the 'CORTEX_USER role is indeed required for calling Cortex AI functions, the direct failure in this scenario is due to the incompatibility with dynamic tables. Option C is incorrect as there's no specified TARGET_LAG' limitation. Option D is incorrect; Snowflake recommends using a smaller warehouse (no larger than MEDIUM) for Cortex functions, but a Snowpark-optimized warehouse is not a strict requirement, and larger warehouses do not increase performance. Option E is incorrect because 'SENTIMENT is designed for English-language input text, and non-English text would likely lead to unexpected or inaccurate results, not a direct failure of the function call itself.

NEW QUESTION # 217
A Gen AI Specialist is tasked with enhancing a Cortex Analyst semantic model to improve the accuracy of literal string searches for product names within user queries. The product names are stored in a high-cardinality PRODUCT_NAME column in the underlying PRODUCT table. The current semantic model already defines a dimension for PRODUCT_NAME. Which of the following configurations and considerations are essential for integrating Cortex Search with Cortex Analyst to achieve this goal?

- Define a `sample_values` array within the `PRODUCT_NAME` dimension in the semantic model YAML, providing a comprehensive list of all possible product names to enable semantic similarity search by Cortex Analyst's internal mechanisms.
- Create a Cortex Search Service on the `PRODUCT_NAME` column of the underlying `PRODUCT` table and configure the `cortex_search_service` field within the `PRODUCT_NAME` dimension in the semantic model YAML to reference this service.
- Ensure the `PRODUCT_NAME` dimension's `data_type` is set to `VARIANT` to allow for flexible matching of various product name formats.
- Increase the `max_tokens` parameter for the Cortex Analyst REST API calls to accommodate longer product name literals in the input prompt.
- Specify `use_as_onboarding_question: true` for relevant product-related verified queries in the Verified Query Repository to pre-load common product searches.

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

Answer: B

Explanation:
For dimensions with high-cardinality, Cortex Analyst recommends integrating with Cortex Search to improve literal string searches. A Cortex Search Service can be created on the relevant column (e.g., 'PRODUCT NAME') to perform a semantic search for literal values. The semantic model's dimension should then include the configuration, referencing the created service. Option A is incorrect because 'sample_values' are recommended for dimensions with low-cardinality (approximately 1-10 distinct values) to avoid exceeding the LLM's context window. For high-cardinality data, Cortex Search is the appropriate solution. Option C is incorrect because 'VARIANT and other complex data types are currently not supported for dimensions in Cortex Analyst semantic models. Option D is incorrect. While 'max_tokens' can be controlled for 'COMPLETE functions, Cortex Analyst's primary mechanism for literal search improvement is through semantic search over sample values or Cortex Search Services, not solely by increasing token limits. Option E is incorrect. The 'use_as_onboarding_question' flag is used for Verified Queries to explicitly suggest questions to users as a starting point, not to improve the accuracy of literal string matching within queries.

NEW QUESTION # 218
A data scientist is leveraging various Snowflake Cortex LLM functions to process extensive text data for an application. To effectively manage their budget, they need a clear understanding of how costs are incurred for each specific function. Which of the following statements accurately describe how costs are calculated for Snowflake Cortex LLM functions, with a particular focus on token usage?

The SNOWFLAKE.CORTEX.EMBED_TEXT_768 function incurs compute costs based on both the input and output tokens processed.

For the SNOWFLAKE.CORTEX.EXTRACT_ANSWER function, billable tokens are determined by the sum of tokens present in the source_document and question fields.

When utilizing AI_COMPLETE with structured outputs, an additional compute cost is explicitly charged for the overhead of validating each generated token against the provided JSON schema.

The SNOWFLAKE.CORTEX.CLASSIFY_TEXT function includes the token count from category descriptions and examples as part of the input tokens for each record processed, thereby increasing the overall cost.

The AI_PARSE_DOCUMENT function is billed based on the total number of individual documents processed, regardless of the number of pages within each document.

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

Answer: B,D

Explanation:
Option B is correct because for the 'EXTRACT_ANSWER function, the number of billable tokens is the sum of the tokens in the Trom_text' (source_document) and 'question' fields. Option D is correct as for 'CLASSIFY TEXT (or labels, descriptions, and examples provided in the categories are counted as input tokens for each record processed, which directly increases the cost. Option A is incorrect because 'EMBED TEXT 768' and 'EMBED TEXT 1024' functions only count 'input tokens' towards the billable total, not both input and output tokens. Option C is incorrect because Cortex Structured Outputs does not incur additional compute cost for the overhead of verifying tokens against the supplied JSON schema, although schema complexity can increase total token consumption. Option E is incorrect because (and 'SNOWFLAKE.CORTEX.PARSE_DOCUMENT) billing is based on the 'number of document pages processed' (e.g., 3.33 Credits per 1 ,000 pages for Layout mode), not just the number of documents. For paged formats (PDF, DOCX), each page is billed as a page; for image files, each image is a page; for HTML/TXT, each 3,000 characters is a page.

NEW QUESTION # 219
A data engineering team needs to establish an automated pipeline in Snowflake to continuously extract 'contract_id' and effective_date' from new PDF contract documents uploaded to an internal stage named They have a pre-trained Document AI model named 'contract_processor'. Which of the following sets of SQL commands correctly configures the necessary Snowflake objects for this automated processing pipeline, including handling file access and initial data loading?

- A.

```
CREATE STAGE @contract_documents_stage;
CREATE STREAM contract_stream ON TABLE contract_processor_raw_table;
CREATE TASK process_contracts
WAREHOUSE = doc_ai_wh
SCHEDULE = '15 MINUTE'
AS
INSERT INTO processed_contracts_table (file_name, contract_id, effective_date)
SELECT
    file_path AS file_name,
    contract_processor!PREDICT(file_content, 'latest'):contract_id.value::STRING AS contract_id,
    contract_processor!PREDICT(file_content, 'latest'):effective_date.value::STRING AS effective_date
FROM
    (SELECT file_path, GET_PRESIGNED_URL('@contract_documents_stage', file_path) AS file_content FROM contract_processor_raw_table WHERE METADATA$ACTION = 'INSERT');
ALTER TASK process_contracts RESUME;
```

- B.

```
CREATE INTERNAL STAGE contract_documents_stage ENCRYPTION=(TYPE='SNOWFLAKE_SSE');
CREATE DYNAMIC TABLE processed_contracts_dt
TARGET_LAG = '15 minutes'
WAREHOUSE = doc_ai_wh
AS
SELECT
    RELATIVE_PATH AS file_name,
    contract_processor!PREDICT(GET_PRESIGNED_URL('@contract_documents_stage', RELATIVE_PATH)):contract_id.value::STRING AS contract_id,
    contract_processor!PREDICT(GET_PRESIGNED_URL('@contract_documents_stage', RELATIVE_PATH)):effective_date.value::STRING AS effective_date
FROM
    DIRECTORY(@contract_documents_stage);
```

- C.

```
CREATE STAGE @contract_documents_stage DIRECTORY = (ENABLE = TRUE);
CREATE PIPE contract_pipe
AS
COPY INTO processed_contracts_table
FROM (SELECT contract_processor!PREDICT(metadata$filename, 1) FROM @contract_documents_stage);
ALTER PIPE contract_pipe REFRESH;
```

- D.

```
CREATE STAGE @contract_documents_stage DIRECTORY = (ENABLE = TRUE) ENCRYPTION=(TYPE='SNOWFLAKE_SSE');
CREATE STREAM contract_stream ON STAGE @contract_documents_stage;
CREATE TASK process_contracts
WAREHOUSE = doc_ai_wh
SCHEDULE = '15 MINUTE'
WHEN SYSTEM$STREAM$HAS_DATA('contract_stream')
AS
INSERT INTO processed_contracts_table (file_name, contract_id, effective_date)
SELECT
    s.RELATIVE_PATH AS file_name,
    p.json_content:contract_id.value::STRING AS contract_id,
    p.json_content:effective_date.value::STRING AS effective_date
FROM
    contract_stream s,
    LATERAL FLATTEN(INPUT => contract_processor!PREDICT(GET_PRESIGNED_URL('@contract_documents_stage', s.RELATIVE_PATH))) p (json_content);
ALTER TASK process_contracts RESUME;
```

- E.

```
CREATE STAGE @contract_documents_stage ENCRYPTION=(TYPE='SNOWFLAKE_SSE');
CREATE STREAM contract_stream ON STAGE @contract_documents_stage;
CREATE TASK process_contracts
WAREHOUSE = doc_ai_wh
SCHEDULE = '15 MINUTE'
WHEN SYSTEM$STREAM_HAS_DATA('contract_stream')
AS
INSERT INTO processed_contracts_table (file_name, contract_id, effective_date)
SELECT
    RELATIVE_PATH AS file_name,
    json_content:contract_id.value::STRING AS contract_id,
    json_content:effective_date.value::STRING AS effective_date
FROM
    contract_stream,
    LATERAL FLATTEN(INPUT => contract_processor!PREDICT(GET_PRESIGNED_URL('@contract_documents_stage', RELATIVE_PATH))) AS json_content;
ALTER TASK process_contracts RESUME;
```

Answer: D

Explanation:

Option C is correct. It correctly sets up the internal stage with `DIRECTORY = (ENABLE = TRUE)` and `ENCRYPTION = (TYPE = 'SNOWFLAKE_SSE')`, which are prerequisites for Document AI. A stream is then correctly created on this stage to track new files. The task is appropriately defined with a warehouse, schedule, and `SYSTEM$STREAM_HAS_DATA` to trigger processing when new data arrives. Inside the task, `GET_PRESIGNED_URL` is used to provide the document path to the `!PREDICT` function, ensuring proper file access from the stage. The `LATERAL FLATTEN` construct is correctly applied to parse the JSON output from `!PREDICT`, which typically returns extracted entities like `contract_id` and `effective_date` as arrays of objects with `value` and `score` keys. Finally, the task is resumed. Option A is missing `DIRECTORY = (ENABLE = TRUE)` for the stage, which is necessary for directory tables used by streams on stages. Option B has multiple errors, including creating a stream on a table instead of a stage, an incorrect `!PREDICT` call, and an unsupported mechanism for passing file content. Option D attempts to use a Dynamic Table, but Snowflake Cortex functions, including Document AI's `!PREDICT`, do not currently support dynamic tables. Option E uses `CREATE PIPE`, which is not the standard Snowflake mechanism for Document AI automated pipelines using streams and tasks, and the `!PREDICT` syntax for file access is incorrect.

NEW QUESTION # 220

An ML engineer is planning a fine-tuning project for a
llama3.1-8b
model to summarize long customer support tickets. They are considering the impact of dataset size and max_epochs on cost and performance, as well as the behavior of the fine-tuned model for inference. Which statements about cost and performance in Snowflake Cortex Fine-tuning are true? (Select all that apply)

- A. The cost for inferencing with a fine-tuned model using the
- B. For optimal cost efficiency, especially with smaller datasets, the
- C. D When fine-tuning a
- D. For large fine-tuning jobs with substantial datasets, particularly when exceeding millions of rows, utilizing Snowpark-optimized warehouses is recommended for improved performance during the training phase.
- E. The compute cost for fine-tuning is primarily determined by multiplying the number of input tokens in the training data by

the number of epochs trained.

Answer: C,D,E

Explanation:
Option A is correct. For the
llama3.1-8b
model, the context window specifically allotted for the
prompt
during fine-tuning is 20,000 tokens, and for the
completion
is 4,000 tokens. Option B is correct. The compute cost incurred for Cortex Fine-tuning is based on the number of tokens used in training, which is calculated as 'number of input tokens number of epochs trained'. Option C is incorrect. While max_epochs can be set to a value from 1 to 10 (inclusive), the default is automatically determined by the system. Setting it to the maximum for 'optimal cost efficiency' is not universally recommended, as a higher number of epochs directly increases the compute cost, and the goal is often to select the smallest model that satisfies the need. Option D is incorrect. When using the COMPLETE function for inference with a fine-tuned model, *both* input and output tokens incur compute cost. Option E is correct. Snowpark-optimized warehouses are recommended for Snowpark workloads with large memory requirements, such as ML training use cases, particularly if the training data has more than 5 million rows. Fine-tuning is an ML training process, so this guidance applies.

NEW QUESTION # 221
......

search for ➡ GES-C01 □□□ for free download □Reliable GES-C01 Exam Online

- Snowflake GES-C01 Online Practice Test □ The page for free download of □ GES-C01 □ on [www.pdfvce.com] will open immediately □Practice GES-C01 Test
- 2025 GES-C01 – 100% Free New Exam Materials | Valid SnowPro® Specialty: Gen AI Certification Exam Most Reliable Questions □ Easily obtain ➥ GES-C01 □ for free download through ➡ www.examsreviews.com □□□ □Test GES-C01 Score Report
- Snowflake GES-C01 Online Practice Test □ Open 《 www.pdfvce.com 》 and search for ➤ GES-C01 □ to download exam materials for free □GES-C01 Valid Practice Materials
- Reliable GES-C01 Exam Questions □ Trustworthy GES-C01 Practice □ Latest GES-C01 Test Blueprint □ Search on ➥ www.dumps4pdf.com □ for □ GES-C01 □ to obtain exam materials for free download □New GES-C01 Braindumps Free
- Practice GES-C01 Test □ New GES-C01 Test Test 🈸 Reliable GES-C01 Exam Questions □ Open website ☀ www.pdfvce.com □☀□ and search for 「 GES-C01 」 for free download □Trustworthy GES-C01 Practice
- Snowflake GES-C01 Online Practice Test □ Search for ▹ GES-C01 ◃ and obtain a free download on □ www.examcollectionpass.com □ □New GES-C01 Braindumps Free
- daotao.wisebusiness.edu.vn, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, onlinecourseshub.com, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, courses.katekoronis.com, ce.snpolytechnic.com, drmsobhy.net, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, Disposable vapes

BTW, DOWNLOAD part of TestSimulate GES-C01 dumps from Cloud Storage: https://drive.google.com/open?id=1dn2rA8kdSbGjc0TOPJFSv4cNSWBetgoY