

Pass Guaranteed Snowflake - GES-C01–Efficient New Cram Materials



DOWNLOAD the newest RealValidExam GES-C01 PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1ØfA9Rqd_FJ1OOuwuK7j04EYNfh93rw8

If you try to free download the demos on the website, and you will be amazed by our excellent GES-C01 preparation engine. We can absolutely guarantee that even if the first time to take the exam, candidates can pass smoothly. You can find the latest version of GES-C01 Practice Guide in our website and you can practice GES-C01 study materials in advance correctly and assuredly. The following passages are their advantages for your information

With over a decade's business experience, our GES-C01 test torrent attached great importance to customers' purchasing rights all along. There is no need to worry about virus on buying electronic products. For we make endless efforts to assess and evaluate our GES-C01 exam prep' reliability for a long time and put forward a guaranteed purchasing scheme, we have created an absolutely safe environment and our GES-C01 Exam Question are free of virus attack. Given that there is any trouble with you, please do not hesitate to leave us a message or send us an email; we sincere hope that our GES-C01 test torrent can live up to your expectation.

>> New GES-C01 Cram Materials <<

Free Snowflake GES-C01 Pdf Guide | GES-C01 Exam Syllabus

All contents of the GES-C01 exam questions are masterpieces from experts who imparted essence of the exam into our GES-C01 study prep. So our high quality and high efficiency GES-C01 practice materials conciliate wide acceptance around the world. By incubating all useful content GES-C01 training engine get passing rate from former exam candidates of 98 which evince our accuracy rate and proficiency.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q303-Q308):

NEW QUESTION # 303

A business user frequently asks Cortex Analyst questions that require filtering on specific product names, such as "What were the sales for 'iced tea' last month?" The 'product' dimension has many distinct values (high cardinality), and Cortex Analyst sometimes struggles to accurately identify the exact literal product name, leading to less precise SQL queries. The Gen AI Specialist wants to enhance Cortex Analyst's ability to find these literal values for the 'product' dimension. To improve Cortex Analyst's literal search capability for the high-cardinality 'product' dimension, which of the following is the most appropriate and recommended approach to configure in the semantic model?

- Increase the number of `sample_values` for the 'product' dimension in the semantic model to several hundred, covering all known product names.
- Create a Cortex Search Service on the underlying 'product' column and specify this service in the `cortex_search_service` field of the 'product' dimension in the semantic model.
- Embed the entire list of product names directly into the `description` field of the 'product' dimension.
- Implement a custom UDF to perform fuzzy string matching on product names within the SQL generated by Cortex Analyst.
- Rely solely on the LLM's inherent ability to infer product names from the user's natural language question without additional configuration.

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

Answer: E

Explanation:

Cortex Analyst offers solutions to improve literal usage, including semantic search over sample values in the semantic model and semantic search using Cortex Search Services. For dimensions with high cardinality (many distinct values), creating a Cortex Search Service on the underlying column and specifying it in the field of the dimension within the semantic model is the recommended approach. This allows for high-quality "fuzzy" search to find literal values needed for Cortex Analyst's SQL queries. Option A is less effective for high-cardinality dimensions because only a fixed-size set of sample values is presented to the LLM, regardless of how many are provided. Option C is not the intended use for the 'description' field and could exceed context window limits. Option D, while a possible technical solution, bypasses the integrated and optimized Cortex Search functionality designed for this purpose. Option E is explicitly contradicted by the scenario, which indicates the LLM struggles, and the available solutions are designed to address this limitation.

NEW QUESTION # 304

A financial institution needs to process thousands of incoming PDF loan application forms daily, extracting applicant names, loan amounts, and submission dates, and loading them into a Snowflake table. They aim for continuous processing with minimal manual intervention. Which of the following statements correctly describe how Document AI can be used in an automated SQL pipeline for this purpose?

- A. The SNOWFLAKE .DOCUMENT_INTELLIGENCE_CREATOR database role alone is sufficient for defining the model build and configuring the processing pipeline, without needing additional CREATE MODEL privileges on the schema.
- B. Document AI's PREDICT method natively supports all PDF files up to 500 MB and 500 pages, allowing for large-scale, single-query processing without requiring users to split documents into smaller chunks.
- C. The pipeline can leverage the `<model build name> ! PREDICT` method within a CREATE TASK statement to automatically process new PDFs as they arrive in an internal or external stage, once the Document AI model build is published.
- D. The extracted information, including confidence scores and values, is returned as a JSON object, which can then be parsed into separate columns in a Snowflake table using SQL functions like LATERAL FLATTEN.
- E. To ensure continuous data ingestion and processing, a STREAM can be created on the stage to detect new PDF documents, triggering the TASK for extraction and subsequent loading into a Snowflake table.

Answer: C,D,E

Explanation:

Option A is correct because DocumentAI supports creating automated pipelines with tasks that call the method to extract information from documents in a stage. Option B is correct as streams are used to detect new data (e.g., PDFs) in a stage, and tasks can be set up to execute when new data is available in the stream, enabling continuous processing. Option E is correct because the 'PREDICT method returns its results as a JSON object, which typically contains 'score' and 'value' fields for extracted entities, and this JSON output can be parsed into separate columns using 'LATERAL FLATTEN'. Option C is incorrect as, in addition to the 'SNOWFLAKE.DOCUMENT_INTELLIGENCE_CREATOR database role, the role used must also have 'CREATE

SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE and 'CREATE MODEL' privileges on the schema where the model build is located. Option D is incorrect because DocumentAI has specific limitations on document size (max 50 MB) and page count (max 125 pages per document), and also limits processing to a maximum of 1000 documents in one query.

NEW QUESTION # 305

A Snowflake administrator is tasked with monitoring the efficiency and cost-effectiveness of their Cortex Analyst deployments. They need to identify if certain semantic models are generating a high volume of failed or expensive queries. Which of the following approaches or statements are crucial for effectively monitoring and identifying issues with Cortex Analyst usage and associated costs?

Querying the `SNOWFLAKE.LOCAL.CORTEX_ANALYST_REQUESTS` function to access detailed logs, including generated SQL and error/warning messages, for specific semantic models.

Analyzing the `CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY` view to track pages processed and credits used, as it aggregates all AI Services activity including Cortex Analyst.

Inspecting the `METERING_DAILY_HISTORY` view with a `SERVICE_TYPE` filter of 'AI_SERVICES' to get a daily aggregate of all AI service credit consumption, including Cortex Analyst.

Reviewing the 'prompt_tokens' and 'completion_tokens' columns in the query logs to directly assess the token count-based cost incurred by each Cortex Analyst message.

Understanding that there is typically a 1-2 minute lag between a Cortex Analyst request being made and its logs becoming visible in the event table.

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

Answer: A,D,E

Explanation:

Option A is correct because Cortex Analyst logs requests to an event table, and the function can be used to query these logs, which include generated SQL and errors/warnings, helping identify issues. Option C is correct as the 'METERING_DAILY_HISTORY' view tracks daily credit usage for services, including 'AI_SERVICES', which encompasses Cortex Analyst. Option E is correct as there is a reported 1-2 minute lag for Cortex Analyst requests to become visible in the event table logs. Option B is incorrect because 'CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY' specifically displays Document AI processing function activity, not all AI Services or Cortex Analyst. Option D is incorrect because Cortex Analyst costs are based on messages, not tokens, so 'prompt_tokens' and 'completion_tokens' would not be relevant for direct cost assessment in this context.

NEW QUESTION # 306

A Gen AI developer has a Document AI pipeline that uses a query with 'GET PRESIGNED URL' to process multi-page PDF documents. Despite the internal stage being correctly set up with 'SNOWFLAKE SSE' encryption and the model build being published, they observe inconsistent results. Some documents result in a Received HTTP 403 response for presigned URL. URL may be expired.

error, while other documents (containing complex diagrams and dense text in an unsupported language like Korean) are processed, but the extracted information is often incomplete or inaccurate.

Which two factors are most likely contributing to these observed issues?

- A. The 'PREDICT' method is being called with an outdated model build version instead of the latest one, leading to performance degradation.
- B. The role lacks the 'EXECUTE TASK' privilege, preventing the scheduled pipeline tasks from running.
- C. The Document AI model is returning answers longer than its limit of 512 tokens for entity extraction or 2048 tokens for table extraction.
- D. The documents are in an unsupported language or exceed the maximum page length of 125 pages, causing extraction failures or inaccuracies.
- E. The default expiration time for the 'GET PRESIGNED URL' function is causing some URLs to expire before the Document AI model can process them.

Answer: D,E

Explanation:

The error 'Received HTTP 403 response for presigned URL. URL may be expired.' directly indicates that the function's default expiration time is causing some documents to be inaccessible by the Document AI model. This is a common issue when processing pipelines encounter delays. Additionally, the observation of incomplete or inaccurate extraction for documents with 'dense text in an unsupported language like Korean' directly points to language limitations. Document AI explicitly lists supported languages (English, Spanish, French, German, Portuguese, Italian, and Polish) and states that results for other languages might not be satisfactory. While the question mentions 'multi-page PDF documents' without explicitly stating they exceed page limits, the mention of 'complex diagrams and dense text' can also imply potential issues if page length (max 125 pages) is exceeded or other document requirements are not met. Thus, option D comprehensively covers these content-related issues. Option A (outdated model version) is unlikely to cause these specific errors, as the latest model is used by default if not specified. Option C (missing 'EXECUTE TASK' privilege) would prevent task execution entirely, not cause intermittent URL issues or content-specific extraction problems. Option E (answers exceeding token limits) would be reflected in truncated output, not necessarily 'incomplete or inaccurate' extraction in the sense of failing to identify information in the first place.

NEW QUESTION # 307

A data science team is using SNOWFLAKE.CORTEX.CLASSIFY_TEXT to categorize product reviews into detailed segments like 'Bug Report - Critical', 'Feature Request - UI/UX', 'General Praise', or 'Query - Billing Issue'. For highly nuanced reviews, they find the initial classifications lack precision, and they are also concerned about the associated compute costs for processing large volumes of data. Which strategies should they employ to optimize classification accuracy and manage costs effectively with this function?

- A. If classifying thousands of reviews, they can significantly reduce overall compute costs by setting the temperature option to 0.0 within CLASSIFY_TEXT to ensure deterministic and cheaper inference.
- B. To reduce input token costs for classifications, the input text should be pre-processed to remove common stop words and punctuation, as these characters are counted as billable tokens without contributing to classification accuracy.
- C. To improve accuracy for ambiguous classifications, they should augment the list_of_categories with explicit description and examples for each category, understanding that these additions will increase input token costs for each record processed.

```
SELECT SNOWFLAKE.CORTEX.CLASSIFY_TEXT(
  'The app freezes after login, making it unusable.',
  [
    { 'label': 'Bug Report - Critical', 'description': 'Software defect causing severe issues with application functionality.', 'examples': ['App crashes frequently', 'Login button unresponsive'] },
    { 'label': 'Feature Request - UI/UX', 'description': 'Suggestion for user interface or experience improvement.' }
  ]
);
```

- D. CLASSIFY_TEXT labels, descriptions, and examples are counted as input tokens only once per function call, regardless of the number of records processed in a batch, to optimize cost efficiency.
- E. For complex scenarios where the relationship between review text and categories is not straightforward, including a concise task_description (e.g., 'Classify the product review focusing on technical support relevance') in the options argument is recommended to guide the model.

```
SELECT SNOWFLAKE.CORTEX.CLASSIFY_TEXT(
  'The new update introduced a confusing navigation bar that is hard to use.',
  ['Bug Report', 'Feature Request'],
  { 'task_description': 'Classify the feedback regarding app usability issues.' }
);
```

Answer: C,E

Explanation:

Option A is correct because adding label descriptions and examples can improve classification accuracy, especially when category definitions are ambiguous. The source explicitly states that each label, description, and example counts as input tokens for each record processed by a 'CLASSIFY_TEXT' function call, incurring costs accordingly. Option B is correct because adding a clear 'task_description' can improve accuracy when the relationship between the input text and categories is ambiguous or nuanced. Option C is incorrect; while token counts contribute to cost, the sources do not recommend removing stop words and punctuation for cost reduction or as a general best practice for SCLASSIFY TEXT. The focus is on using plain English input. Option D is incorrect because the 'temperature' option is available for 'COMPLETE' and functions to control output randomness, but it is not listed as an option for 'CLASSIFY_TEXT' in its syntax. Furthermore, while a lower temperature can make results more deterministic, the source does not link it to 'cheaper' inference cost for these task-specific functions, but rather to consistency for 'COMPLETE'. Option E is incorrect because 'AI_CLASSIFY' labels, descriptions, and examples are indeed counted as input tokens for 'each record processed', not just once per call, as clearly stated in the cost considerations.

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, 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, allprotrainings.com,
www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, 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

DOWNLOAD the newest RealValidExam GES-C01 PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=19fA9Rqd_FJ1OOuwuK7j04EYNfH93rw8