

PassTIP 에서 Snowflake GES-C01 덤프를 구매하시면 일년무료 업데이트서비스를 받을수 있습니다. 일년무료 업데이트서비스란 구매일로부터 1년동안 구매한 덤프가 업데이트될때마다 구매시 사용한 메일주소로 가장 최신버전을 보내드리는것을 의미합니다. Snowflake GES-C01 덤프에는 가장 최신시험문제의 기출문제가 포함되어있어 높은 적중율을 자랑하고 있습니다.

최신 Snowflake Certification GES-C01 무료샘플문제 (Q299-Q304):

질문 # 299

A data engineer is designing an automated pipeline to process customer feedback comments from a 'new_customer_reviews' table, which includes a 'review_text' column. The pipeline needs to classify each comment into one of three predefined categories: 'positive', 'negative', or 'neutral', and store the classification label in a new 'sentiment_label' column.

Which of the following statements correctly describe aspects of implementing this data transformation using 'SNOWFLAKE.CORTEX.CLASSIFY_TEXT' in a Snowflake pipeline?

- A. The classification can be achieved by integrating a 'SELECT' statement with
- B. The argument must contain exactly three unique categories for sentiment classification.
- C. The cost for 'CLASSIFY_TEXT' is incurred based on the number of pages processed in the input document.
- D. Both the input string to classify and the are case-sensitive, potentially yielding different results for variations in capitalization.
- E. Including an optional 'task_description' such as

정답: A,D,E

설명:

Option A is correct. 'SNOWFLAKE.CORTEX.CLASSIFY_TEXT' classifies free-form text into categories and returns an 'OBJECT' value (VARIANT) where the 'label' field specifies the category. This can be extracted using '['label']' and seamlessly integrated into 'INSERT' or 'UPDATE' statements within a pipeline task for data transformation.

Option B is correct. Adding a clear 'task_description' to the 'options' argument for 'CLASSIFY_TEXT' can significantly improve classification accuracy. This is particularly useful when the relationship between the input text and the provided categories is ambiguous or nuanced.

Option C is incorrect. incurs compute cost based on the number of tokens processed (both input and output tokens), not on the number of pages in a document. Functions like 'AI_PARSE_DOCUMENT' bill based on pages.

Option D is incorrect. The argument for 'CLASSIFY_TEXT' must contain at least two and at most 100 unique categories. It is not strictly limited to three for any classification task, including sentiment.

Option E is correct. Both the 'input' string to classify and the are case-sensitive, meaning that differences in capitalization for either the input text or the category labels can lead to different classification results.

질문 # 300

A data processing team is using Snowflake Document AI to extract data from incoming supplier invoices. They observe that many documents are failing to process, and successful extractions are taking longer than expected, leading to increased costs. Upon investigation, they find error messages such as

. Additionally, their 'X-LARGE' virtual warehouse is constantly active, contributing to higher-than-anticipated bills. Which two of the following actions are essential steps to troubleshoot and address the root causes of these processing errors and optimize their Document AI pipeline?

- A. Increase the 'max_tokens' parameter within the '!PREDICT' function options to accommodate longer document responses from the model.
- B. Implement a pre-processing step to split documents exceeding 125 pages or 50 MB into smaller, compliant files before loading to the stage.
- C. Scale down the virtual warehouse to 'X-SMALL' or 'SMALL' size, as larger warehouses do not increase Document AI query processing speed and incur unnecessary costs.
- D. Redefine extraction questions to be more generic and encompassing, reducing the number of distinct questions needed per document.
- E. Configure the internal stage used for storing invoices with 'ENCRYPTION = (TYPE = 'SNOWFLAKE SSE')'.

정답: B,E

설명:

The error messages 'Document has too many pages. Actual: 130. Maximum: 125.' and 'File exceeds maximum size. Actual:

54096026 bytes. Maximum: 50000000 bytes: directly indicate that the documents do not meet Document AI's input requirements, which specify a maximum of 125 pages and 50 MB file size. Therefore, implementing a pre-processing step to split or resize these documents is an essential solution (Option B). The error 'cannot identify image file < io.BytesIO object at Ox...>' is a known issue that occurs when an internal stage used for Document AI is not configured with encryption. Correctly configuring the stage with this encryption type is crucial for resolving this processing error (Option D). Option A, while addressing cost optimization, is not a root cause of the 'processing errors' themselves, although it is a best practice for cost governance as larger warehouses do not increase Document AI query processing speed. Option C is incorrect; best practices for question optimization suggest being specific, not generic. Option E is incorrect as 'max_tokenS' relates to the length of the model's output, not the input document's size or page limits.

질문 # 301

A financial analytics team is using to extract specific financial metrics (e.g., revenue, profit margin) from quarterly reports and requires the output in a strict JSON format for automated ingestion into a data warehouse. They've encountered issues where the LLM sometimes generates malformed JSON or includes extraneous text. Which of the following approaches will help ensure deterministic, schema-compliant JSON outputs and mitigate these 'hallucinations' related to format?

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

정답: A,B,D,E

설명:

Option A is correct because 'AI_COMPLETE Structured Outputs allows supplying a JSON schema via the 'response_format' argument, which defines the structure, data types, and constraints, ensuring responses conform to the schema and reducing post-processing. Option B is correct as it specifically states requirements for OpenAI (GPT) models when using 'response_format', including 'additionalProperties: false' and explicitly listing all properties in the 'required' field. Option C is correct because, for medium-to-complex tasks, adding 'Respond in JSON' to the prompt can further improve JSON adherence accuracy, even when 'response_format' is specified. Option D is correct as setting 'temperature' to 0 is recommended for the most consistent results, which is crucial for deterministic JSON outputs, regardless of task or model. Option E is incorrect because 'CLASSIFY_TEXT' is for categorizing text, not for validating JSON structure, making it an inappropriate tool for this task.

질문 # 302

A Gen AI Specialist needs to extract the 'invoice number' and 'total amount' from a specific invoice PDF, 'invoice_001 .pdf, located in an internal stage named They want to use the default (latest) model build version for a model named 'invoice_processor'. Which SQL query correctly uses the '!PREDICT method to extract the required information, and what key fields would be present in the JSON output for a successful extraction of 'invoice_number' and 'total_amount'?

- A. □
- B. □
- C. □
- D. □
- E. □

정답: D

설명:

Option C is correct. The '!PREDICT' method is used with 'GET_PRESIGNED URL' to access a document from a stage. If the model build version is not specified, the latest available version is used by default. The 'T-ROM DUAL' clause is appropriate for extracting information from a single specified document. The function returns a JSON object, where extracted entities like and 'total amount' are typically represented as arrays of objects, each containing 'score' and 'value' fields, and '_documentMetadata.ocrScore' provides the OCR confidence. Option A uses 'FROM DIRECTORY(@invoices_stage which would process all documents in the stage, not just a single specified file. Option B uses an incorrect string literal for the model version and has incorrect keys for the output fields. Option D uses an incorrect syntax for calling the '!PREDICT method and incorrectly assumes direct column output. Option E uses 'URL()' which is not the correct function for generating a presigned URL, and also describes incorrect output fields.

질문 # 303

A team is developing a critical business intelligence application that leverages Snowflake Cortex Analyst to provide natural language querying capabilities over complex structured data. To minimize operational costs while maintaining high accuracy, which of the following strategies are most effective for optimizing the cost efficiency of the Cortex Analyst service?

- A. Leveraging Cortex Search Services integration within the semantic model to improve literal value matching, thereby reducing the need for Cortex Analyst to perform expensive fuzzy string matching or re-prompt the user.
- B. Configuring a custom instruction with a short, precise task description to reduce the input token count for the LLMs orchestrating SQL generation.
- C. Using a smaller, less capable LLM as the underlying summarization agent for multi-turn conversations to reduce token processing costs, even if it slightly degrades conversational context.
- D. Implementing a comprehensive Verified Query Repository (VQR) to guide Cortex Analyst towards pre-validated SQL queries for common questions, which ensures predictable execution and reduces LLM inference iterations.
- E. Optimizing the semantic model YAML file by reducing the number of logical tables and columns to decrease the metadata processed by Cortex Analyst's LLMs per message.

정답: A,D

설명:

Option B is correct because a Verified Query Repository (VQR) helps Cortex Analyst leverage pre-validated SQL for similar questions, improving accuracy and potentially reducing the number of LLM inference calls or complex reasoning steps required for SQL generation, thus making usage more efficient and reducing cost associated with less optimal LLM calls. Option D is correct because integrating Cortex Search Services improves literal search, helping Cortex Analyst find exact literal values needed for SQL queries more accurately and efficiently, which can reduce ambiguity and the need for multiple LLM iterations or incorrect queries, ultimately leading to more cost-effective message processing. Option A is incorrect: While using a smaller LLM might seem to save cost, Llama 3.1 70B was specifically chosen as the summarization agent for multi-turn conversations in Cortex Analyst due to its higher accuracy in rephrasing questions and avoiding errors, implying that a less capable model would degrade performance and potentially lead to more (and thus more expensive) overall messages to achieve a correct answer. The cost for Cortex Analyst is per message, not per token for this component. Option C is incorrect. While a well-scoped semantic model is recommended for accuracy, the sources do not explicitly state that reducing the number of logical tables and columns 'directly' reduces the per-message cost of Cortex Analyst, which is fixed per message. The impact would be indirect through improved accuracy or reduced processing complexity, but not a direct cost reduction based on metadata size for the fixed per-message billing. Option E is incorrect. Cortex Analyst cost is based on the number of messages, not the token count of prompts. While good prompt engineering (like concise custom instructions) is generally good practice, it does not directly reduce the per-message cost of Cortex Analyst as it would for token-based LLM calls.

질문 # 304

.....

PassTIP의 Snowflake 인증 GES-C01 덤프는 인터넷에서 검색되는 Snowflake 인증 GES-C01 시험 공부 자료 중 가장 출중한 시험 준비 자료입니다. Snowflake 인증 GES-C01 덤프를 공부하면 시험 패스는 물론이고 IT 지식을 더 많이 쌓을 수 있어 일거양득입니다. 자격증을 취득하여 자신 있게 승진하여 연봉 협상하세요.

GES-C01 학습 자료 : <https://www.passtip.net/GES-C01-pass-exam.html>

PassTIP에는 Snowflake GES-C01 인증 시험의 특별한 합습 가이드가 있습니다, GES-C01 덤프에 있는 문제만 이해하고 완벽하게 공부하신다면 SnowPro® Specialty: Gen AI Certification Exam 최신 시험을 한방에 패스하여 자격증을 쉽게 취득할 수 있을 것입니다, PassTIP GES-C01 학습 자료는 시험에서 불합격 성적표를 받으시면 덤프 비용을 환불하는 서비스를 제공해 드려 아무런 걱정 없이 시험에 도전하도록 힘이 되어 드립니다, Snowflake GES-C01 덤프는 시험 패스율이 거의 100%에 달하여 많은 사랑을 받아왔습니다, PassTIP의 Snowflake 인증 GES-C01 덤프를 구매하시면 밝은 미래가 보입니다.

그가 프랑스로 떠나고 한국에 없을 걸 알면서도 퇴근 때마다 누군가를 버릇처럼 찾고 있을 자신이, 가볍게 대답한 예원이 미련 없이 문고리를 잡았다, PassTIP에는 Snowflake GES-C01 인증 시험의 특별한 합습 가이드가 있습니다.

높은 적중율을 자랑하는 GES-C01 높은 통과율 시험 대비 자료 덤프 자료로 SnowPro® Specialty: Gen AI Certification Exam 시험 패스가

GES-C01 덤프에 있는 문제만 이해하고 완벽하게 공부하신다면 SnowPro® Specialty: Gen AI Certification Exam 최신 시험을 한방에 패스하여 자격증을 쉽게 취득할 수 있을 것입니다, PassTIP는 시험에서 불합격 성적표를 받으시면 덤프

