

Instant GES-C01 Discount & Training GES-C01 Tools



DOWNLOAD the newest PrepAwayETE GES-C01 PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1ofpDS6Wo1hzOqjYszY3WvhBBguDcPR-G>

No matter how much you study, it can be difficult to feel confident going into the SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam. However, there are a few things you can do to help ease your anxiety and boost your chances of success. First, make sure you prepare with Real GES-C01 Exam Dumps. If there are any concepts you're unsure of, take the time to take GES-C01 practice exams until you feel comfortable.

Free update for 365 days is available if you buy GES-C01 exam braindumps from us. That is to say, in the following year, you can get the latest information about the GES-C01 exam dumps timely. And the update version will be sent to your email automatically. In addition, the GES-C01 Exam Braindumps are compiled by experienced experts who are quite familiar with the dynamics about the exam center, therefore the quality and accuracy of the GES-C01 exam braindumps can be guaranteed.

>> Instant GES-C01 Discount <<

Training GES-C01 Tools & Pass4sure GES-C01 Exam Prep

Our latest GES-C01 exam torrent is comprehensive, covering all the learning content you need to pass the qualifying GES-C01 exams. Users with qualifying exams can easily access our web site, get their favorite latest GES-C01 study guide, and before downloading the data, users can also make a free demo of our GES-C01 Exam Questions for an accurate choice. Users can easily pass the GES-C01 exam by learning our GES-C01 practice materials, and can learn some new knowledge in this field for you have a brighter future.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q301-Q306):

NEW QUESTION # 301

A Gen AI specialist is preparing to upload a large volume of diverse documents to an internal stage for Document AI processing. The objective is to extract detailed information, including lists of items and potentially classifying document types, and then automate this process. Which of the following statements represent 'best practices or important considerations/limitations' when preparing documents and setting up the Document AI workflow in Snowflake? (Select ALL that apply.)

- A. To improve model training, documents uploaded should represent a real use case, and the dataset should consist of diverse documents in terms of both layout and data.
- B. When defining data values for extraction, especially for nonstandard formats or combinations of values, fine-tuning the model with annotations is generally more effective than relying solely on complex prompt engineering.
- C. If the Document AI model does not find an answer for a specific field, the '!PREDICT method will omit the 'value' key but will still return a 'score' key to indicate confidence that the answer is not present.
- D. Documents with a page count exceeding 125 pages or a file size greater than 50 MB will be processed, but with a potential reduction in extraction accuracy.
- E. For continuous processing of new documents, it is best practice to create a stream on the internal stage and a task to automate the '!PREDICT method execution.

Answer: A,B,C,E

Explanation:

Several best practices and limitations apply to Document AI, particularly during the preparation and training phases: - **Option A: Correct.** To enhance model training, it is crucial that the uploaded documents accurately reflect a real use case and that the dataset exhibits diversity in both layout and content. This includes variation in information, not just all documents containing the same data or presented in the same form. - **Option B: Correct.** If the Document AI model cannot find an answer for a specified field, it will not return a 'value' key, but it will still provide a 'score' key, which indicates its confidence that the answer is absent from the document. This behaviour is shown in the example for 'buyer_name' field in the '!PREDICT' JSON output. - **Option C: Correct.** Creating an automated pipeline for continuous document processing involves setting up a stream on the internal stage and a task to trigger the '!PREDICT' method. This approach is explicitly mentioned as a capability and best practice for creating document processing pipelines. - **Option D: Incorrect.** Documents exceeding 125 pages or 50 MB in size are strict limitations for Document AI and will result in processing errors, not merely reduced accuracy. Examples of errors include 'Document has too many pages. Actual: 150. Maximum: 125.' or 'File exceeds maximum size. Actual: 54096026 bytes. Maximum: 50000000 bytes.'. - **Option E: Correct.** For complex extraction scenarios, such as nonstandard formats, lists, or classification, 'showing' the model what is expected through annotations and fine-tuning is often more effective than 'telling' it via elaborate prompt engineering. This is described as 'Show, don't tell' in the best practices, recommending annotations for combinations of values, arrays, nonstandard formats, normalization, and classification tasks instead of spending too much time on prompt engineering alone.

NEW QUESTION # 302

A data architect is designing a workflow to programmatically extract highly structured data from various text inputs using Snowflake Cortex AI's AI_COPCOMPLETE function through its REST API. They require the output to strictly adhere to a complex JSON schema for downstream processing and need to manage associated costs. Which of the following statements accurately describe aspects of this approach?

which can include a 'required' array for mandatory fields.

- | Snowflake Cortex AI_COMPLETE Structured Outputs incurs additional compute cost specifically for the overhead of validating each generated token against the applied JSON schema, on top of standard token processing costs.
- | For optimal JSON adherence accuracy, especially with complex reasoning tasks, it is always recommended to include strong prompt instructions like 'Respond in strict JSON' in the LLM prompt, in addition to providing the response_format.
- | The response_format argument and its JSON schema can be specified for any model supported by AI_COMPLETE via the REST API, including mistral-large2 and claude-3-5-sonnet.
- | Increasing the complexity or nesting depth of the JSON schema, while ensuring output structure, will generally lead to a reduction in the number of input and output tokens consumed, thereby lowering costs.

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

Answer: A,B,E

Explanation:

Let's review each option: - Option A is correct. To obtain a structured response via the 'COMPLETE' or 'AI_COMPLETE' function through the REST API, the 'response_format' argument is specified within the request body, requiring "type": "json" and a "schema" object. The 'required' array within the JSON schema is critical for defining mandatory fields, and the function will raise an error if any specified required field cannot be extracted. - Option B is incorrect. 'AI_COMPLETE' Structured Outputs incurs compute cost based on the number of tokens processed, but it does not incur additional compute cost for the overhead of verifying each token against the supplied JSON schema. - Option C is correct. While structured outputs generally don't always require explicit prompting for simple tasks, for complex reasoning tasks, adding specific instructions like 'Respond in JSON' or details about the schema to the prompt can significantly improve JSON adherence accuracy. - Option D is correct. All models supported by 'AI_COMPLETE' support structured output. The provided documentation includes REST API examples demonstrating the use of 'response_format' with models like 'mistral-large2' and 'claude-3-5-sonnet'. - Option E is incorrect. The number of tokens processed (and thus billed) generally increases with schema complexity. A larger and more complex schema, especially one with deep nesting, will consume a greater number of input and output tokens.

NEW QUESTION # 303

A financial services company is developing an automated data pipeline in Snowflake to process Federal Reserve Meeting Minutes, which are initially loaded as PDF documents. The pipeline needs to extract specific entities like the FED's stance on interest rates ('hawkish', 'dovish', or 'neutral') and the reasoning behind it, storing these as structured JSON objects within a Snowflake table. The goal is to ensure the output is always a valid JSON object with predefined keys. Which AI_COMPLETE configuration, used within an in-line SQL statement in a task, is most effective for achieving this structured extraction directly in the pipeline?

- Using a simple prompt like 'Extract FED stance on interest rates and the reasoning as JSON from the document: [document_text]' and expecting the LLM to format it correctly without further configuration.
- Setting the `temperature` parameter to 0 and `max_tokens` to a sufficiently large value, along with a prompt to output JSON, to inherently guide the LLM to a structured format.
- Employing the `response_format` argument within `AI_COMPLETE` with a JSON schema that explicitly defines the 'stance' (enum: 'hawkish', 'dovish', 'neutral') and 'reasoning' (string) fields, ensuring strict adherence to the output structure.
- Utilizing multiple calls to `SNOWFLAKE.COREX.EXTRACT_ANSWER()` or `AI_EXTRACT()`, one for the stance and another for the reasoning, then manually constructing the JSON in a subsequent SQL step.
- Relying on default `AI_COMPLETE` behavior, as Snowflake Cortex automatically detects the need for JSON output when entity extraction with specific fields is implied by the prompt.

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

Answer: B

Explanation:

To ensure that LLM responses adhere to a predefined JSON structure, the 'AI_COMPLETE' function's 'response_format' argument, which accepts a JSON schema, is the most effective and direct method. This mechanism enforces the structure, data types, and required fields, significantly reducing the need for post-processing and ensuring deterministic, high-quality output. The AI-Infused Data Pipelines with Snowflake Cortex blog highlights asking the LLM to create a JSON object for maximizing utility. While setting 'temperature' to 0 can improve consistency, it does not enforce a specific schema. Prompt engineering (Option A) can help but does not guarantee strict adherence. Using multiple extraction calls (Option D) is less efficient and robust for extracting multiple related fields than a single 'AI_COMPLETE' call with a structured output schema. Snowflake Cortex does not automatically infer and enforce a JSON schema without explicit configuration (Option E).

NEW QUESTION # 304

A marketing analyst wants to quickly gauge the overall sentiment of customer feedback stored in a Snowflake table called `CUSTOMER_FEEDBACK`, which has a column `FEEDBACK_TEXT`. They decide to use the `SNOWFLAKE.COREX.SENTIMENT` function to process a review. Consider the following SQL query for a specific review:

Which of the following correctly describes the expected output format and interpretation of the `sentiment_score` for the given input?

- A. The output will be a boolean value (TRUE/FALSE) indicating if the sentiment is positive, and 0.5 represents a neutral sentiment.
- B. The output will be an integer between 0 and 100, where higher values denote more positive sentiment.
- C. **The output will be a floating-point number between -1 and 1 (inclusive), where a value of 1 indicates strong positive sentiment and -1 indicates strong negative sentiment.**
- D. The output will be a string like 'Positive' or 'Negative', and a score close to 1 indicates strong positive sentiment.
- E. The output will be a JSON object containing a 'label' field, and values around 0 indicate a neutral sentiment.

Answer: C

Explanation:

Option B is correct. The function returns a floating-point number in the range of -1 to 1 (inclusive). A SNOWFLAKE.COREX.SENTIMENT score of 1 signifies strong positive sentiment, -1 signifies strong negative sentiment, and values around 0 indicate neutral sentiment. The function does not return a string label, a JSON object, a boolean, or an integer percentage.

NEW QUESTION # 305

An enterprise is deploying a new RAG application using Snowflake Cortex Search on a large dataset of customer support tickets. The operations team is concerned about managing compute costs and ensuring efficient index refreshes for the Cortex Search Service, which needs to be updated hourly. Which of the following considerations and configurations are relevant for optimizing cost and performance of the Cortex Search Service in this scenario?

- A. **CHANGE_TRACKING**
- B. For optimal performance and cost efficiency, Snowflake recommends using a dedicated warehouse of size no larger than **MEDIUM** for each Cortex Search Service.

- C. The CREATE CORTEX SEARCH SERVICE
- D. The primary cost driver for Cortex Search is the number of search queries executed against the service, with the volume of indexed data (GB/month) having a minimal impact on overall billing.
- E. For embedding text, selecting a model like

 snowflake-arctic-embed-m-v1.5
 (0.03 credits/million tokens) over voyage-multilingual-2
 (0.07 credits/million tokens) could significantly reduce EMBED_TEXT_TOKENS

Answer: A,B,C,E

Explanation:

Option A is correct because a Cortex Search Service requires a virtual warehouse to refresh the service, which runs queries against base objects when they are initialized and refreshed, incurring compute costs. Option B is correct because the cost of embedding models varies. For example, 'snowflake-arctic-embed-m-v1.5' costs 0.03 credits per million tokens, while 'voyage-multilingual-2' costs 0.07 credits per million tokens. Choosing a more cost-effective model like 'snowflake-arctic-embed-m-v1.5' for English-only data can reduce token costs. Option C is correct because Snowflake recommends using a dedicated warehouse of size no larger than MEDIUM for each Cortex Search Service to achieve optimal performance. Option D is correct because change tracking is required for the Cortex Search Service to be able to detect and process updates to the base table, enabling incremental refreshes that are more efficient than full re-indexing. Option E is incorrect because Cortex Search Services incur costs based on virtual warehouse compute for refreshes, 'EMBED_TEXT_TOKENS' cost per input token, and a charge of 6.3 Credits per GB/mo of indexed data. The volume of indexed data has a significant impact, not minimal.

NEW QUESTION # 306

.....

It is apparent that a majority of people who are preparing for the GES-C01 exam would unavoidably feel nervous as the exam approaching. If you are still worried about the coming exam, since you have clicked into this website, you can just take it easy now, I can assure you that our company will present the antidote for you--our GES-C01 Learning Materials. As the most popular study materials in the market, our GES-C01 practice guide can give you 100% pass guarantee. You will feel grateful if you choose our GES-C01 training questions.

Training GES-C01 Tools: <https://www.prepawayete.com/Snowflake/GES-C01-practice-exam-dumps.html>

Not only can our GES-C01 guide torrent offer the best service for every user with the totally free update for a year, which gets rid of the worries whether customers can acquire the latest learning material, but also almost 100% passing rate is guaranteed, our GES-C01 exam guide is dedicated to helping everyone achieve the success of certification, Especially if you do not choose the correct study materials and find a suitable way, it will be more difficult for you to pass the exam and get the GES-C01 related certification.

Assigns Information Stewards to govern information assets throughout their lifecycle, Instant GES-C01 Discount The Name Format pop-up menu provides renaming options that allow you to include sequential numbers, version names, date and time stamps, and counters.

Perfect Instant GES-C01 Discount | Amazing Pass Rate For GES-C01 Exam | High Pass-Rate GES-C01: SnowPro® Specialty: Gen AI Certification Exam

Not only can our GES-C01 guide torrent offer the best service for every user with the totally free update for a year, which gets rid of the worries whether customers can acquire the latest learning material, but also almost 100% passing rate is guaranteed, our GES-C01 Exam Guide is dedicated to helping everyone achieve the success of certification.

Especially if you do not choose the correct study materials and find a suitable way, it will be more difficult for you to pass the exam and get the GES-C01 related certification.

In fact, the overload of learning seems not to be a good GES-C01 method, once you are weary of such a studying mode, it's difficult for you to regain interests and energy.

Expertly curated PDF question answers will help you pass the exam on the first attempt. The GES-C01 actual exam guide can provide them with efficient and convenient learning platform Latest GES-C01 Dumps Files so that they can get the certification as soon as possible in the shortest possible time.

P.S. Free & New GES-C01 dumps are available on Google Drive shared by PrepAwayETE: <https://drive.google.com/open?id=1ofpDS6W01hzOqjYszY3WvhBBguDcPR-G>