

Question Snowflake GES-C01 Explanations | GES-C01 Reliable Exam Papers



BTW, DOWNLOAD part of Real4test GES-C01 dumps from Cloud Storage: <https://drive.google.com/open?id=1xFg9ub81SDY6kLOGNukAKTAIG74dAvSk>

With the rapid market development, there are more and more companies and websites to sell GES-C01 guide torrent for learners to help them prepare for exam. If you have known before, it is not hard to find that the study materials of our company are very popular with candidates, no matter students or businessman. Welcome your purchase for our GES-C01 Exam Torrent. As is an old saying goes: Client is god! Service is first! GES-C01 Guide Braindumps can simulate limited-timed examination and online error correcting, and have 24/7 Service Online, GES-C01 Exam Torrent is the best and wisest choice for you to prepare your test.

Every day of our daily life seems to be the same rhythm, work to eat and sleep, and all the daily arrangements, the exam does not go through every day, especially for the key GES-C01 qualification test ready to be more common. In preparing the GES-C01 qualification examination, the GES-C01 study materials will provide users with the most important practice materials. Users can evaluate our products by downloading free demo templates prior to formal purchase.

>> Question Snowflake GES-C01 Explanations <<

Snowflake GES-C01 Reliable Exam Papers - GES-C01 Valid Exam Book

Real4test also presents desktop-based Snowflake GES-C01 practice test software which is usable without any internet connection after installation and only required license verification. SnowPro® Specialty: Gen AI Certification Exam (GES-C01) practice test software is very helpful for all those who desire to practice in an actual SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam-like environment. SnowPro® Specialty: Gen AI Certification Exam (GES-C01) practice test software contains many Snowflake GES-C01 practice exam designs just like the real SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q255-Q260):

NEW QUESTION # 255

A Streamlit application developer wants to use AI_COMPLETE (the latest version of COMPLETE (SNOWFLAKE. CORTEX)) to process customer feedback. The goal is to extract structured information, such as the customer's sentiment, product mentioned, and any specific issues, into a predictable JSON format for immediate database ingestion. Which configuration of the AI COMPLETE function call is essential for achieving this structured output requirement?

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

Answer: D

Explanation:

'AI_COMPLETE Structured OutputS (and its predecessor 'COMPLETE Structured Outputs') specifically allows supplying a JSON schema as the 'response_format' argument to ensure completion responses follow a predefined structure. This significantly reduces the need for post-processing and enables seamless integration with systems requiring deterministic responses. The JSON schema object defines the structure, data types, and constraints, including required fields. While prompting the model to 'Respond in JSON' can improve accuracy for complex tasks, the 'response_format' argument is the direct mechanism for enforcing the schema. Option A is a form of prompt engineering, which can help but does not guarantee strict adherence as 'response_format' does. Option B controls randomness and length, not output structure. Option D is less efficient for extracting multiple related fields compared to a single structured output call. Option E's 'guardrails' are for filtering unsafe or harmful content, not for enforcing output format.

NEW QUESTION # 256

A Gen AI Specialist is developing a conversational analytics application using Cortex Analyst, aiming to provide a seamless multi-turn conversation experience for business users querying structured data. The team observes that follow-up questions are sometimes misinterpreted, especially when the conversation history is long. Which of the following statements accurately describe how Cortex Analyst handles multi-turn conversations and key considerations for optimizing this functionality?

- A. Multi-turn conversation in Cortex Analyst is primarily handled by the CORTEX_ANALYST_MODEL AZURE_OPENAI parameter, which, when enabled, allows Azure OpenAI models to manage conversational context more effectively.
- B. Cortex Analyst incorporates an additional LLM summarization agent before its original workflow to rewrite current-turn questions based on conversation history, with Llama 3.1 70B being a recommended model for this task due to its performance in evaluating summarization quality.
- C. Cortex Analyst simply passes the entire conversation history to all subsequent LLM calls, and optimizing this requires manually truncating the array in messages the REST API request.
- D. When a user shifts intent frequently in a multi-turn conversation, Cortex Analyst automatically resets the conversation history to prevent misinterpretations and improve accuracy.
- E. To address misinterpretation in long conversations, the max_tokens parameter for the Cortex Analyst REST API should be significantly increased to ensure the LLM receives the complete historical context without truncation.

Answer: B

Explanation:

Cortex Analyst supports multi-turn conversations for data-related questions by incorporating an additional LLM summarization agent. This agent processes the conversation history and rewrites the current-turn question to include relevant context from previous turns, thereby providing a more coherent and accurate query for subsequent processing. This approach avoids passing arbitrarily long conversation histories to every LLM agent, which would lead to longer inference times and non-determinism. 'Llama 3.1 70B' has been identified as a sufficient model for this summarization task, achieving high accuracy in rewriting questions. Option A is incorrect because Cortex Analyst specifically uses a summarization agent to avoid simply passing the entire, potentially long, conversation history to all LLM calls. Option C is incorrect. The parameter controls the option to use Azure OpenAI models with Cortex Analyst (a legacy path that is discouraged), but it does not describe the mechanism for handling multi-turn conversational context. Option D is incorrect. While 'max_tokens' influences the length of LLM outputs, the strategy for handling long conversation history in Cortex Analyst is to use a summarization agent to create a concise, relevant context, not simply to increase the token limit to send all historical data. Increasing 'max_tokens' for entire conversation histories would lead to higher costs and potentially longer latencies. Option E is incorrect. The documentation suggests that if a conversation becomes too long or the user's intent shifts frequently, users *might need to reset* the conversation, but it does not state that Cortex Analyst *automatically* performs this reset.

NEW QUESTION # 257

Which of the following SQL snippets, when executed against a single invoice file like "invoice001.pdf", correctly extracts and transforms the desired data, assuming 'json_content' holds the raw Document AI output?

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

Answer: E

Explanation:

Option B correctly uses a Common Table Expression (CTE) to retrieve the raw JSON output from (which is a Document AI method for extracting information from documents in a stage), leveraging to access the document. It then accesses the 'invoice_number' and 'vendor_name' using .value' syntax, appropriate for values returned as an array containing a single object with a 'value' field, as shown in Document AI output examples. The 'LATERAL FLATTEN' clause is correctly applied to expand the array of line items, and 'ARRAY_AGG' combined with 'ARRAY_TO_STRING' converts these items into a comma-separated string. Finally, it groups by the single-value extracted fields.

Option A attempts to flatten the result multiple times or in an incorrect way within the SELECT statement without a proper FROM clause for the flattened data, leading to inefficient or incorrect aggregation. Option C directly references a staged file path (@invoice_docs_stage/invoice001.pdf) without the necessary GET PRESIGNED URL' function, which is required when calling '!PREDICT with a file from a stage. It also incorrectly assumes direct .value' access for array-wrapped single values and does not correctly transform the 'invoice_items' array into a string. Option D's subquery for 'ARRAY AGG' is syntactically problematic for direct column access from the outer query without explicit 'LATERAL FLATTEN' at the top level. Option E only extracts the 'ocrScore' from the document metadata and does not perform the requested data transformations.

NEW QUESTION # 258

An administrator has configured the 'CORTEX_MODELS_ALLOWLIST' parameter to only permit the 'mistral-large?' model at the account level. A user with the 'PUBLIC' role, which has been granted 'SNOWFLAKE.CORTEX USER' and 'SNOWFLAKE "CORTEX- MODEL-ROLE-LLAMA3.1-70B"' , attempts to execute several 'AI COMPLETE queries. Which of the following queries will successfully execute?

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

Answer: B,C

Explanation:

Option A is correct. The query directly references 'MISTRAL-LARGE2', which is explicitly in the account-level , so it will succeed. Option B is correct. When a model name is provided as a string argument, Cortex first treats it as an identifier for a schema-level model object. If found, RBAC is applied. The user's role has 'SNOWFLAKE."CORTEX-MODEL-ROLE-LLAMA3.1-70B"' granted, which provides access to the 'LLAMA3.1-70B' model object in 'SNOWFLAKE.MODELS' , regardless of the 'CORTEX MODELS_ALLOWLIST' setting for plain model names. Option C is incorrect because 'llama3.1-70b' as a plain model name is not in the 'CORTEX_MODELS_ALLOWLIST'. Although the user has access to the model object, a plain string like 'llama3.1-70b' will be looked up in the allowlist after failing to match a model object by that plain name, and the allowlist only has 'MISTRAL-LARGE2'. Option D is incorrect. 'snowflake-arctic' is neither in the 'CORTEX MODELS ALLOWLIST' nor does the user have a specific application role granting access to a 'snowflake-arctic' model object. Option E is incorrect because 'ALTER ACCOUNT' operations can only be performed by the 'ACCOUNTADMIN' role, not by a 'PUBLIC' user role.

NEW QUESTION # 259

A data scientist is tasked with improving the accuracy of an LLM-powered chatbot that answers user questions based on internal company documents stored in Snowflake. They decide to implement a Retrieval Augmented Generation (RAG) architecture using Snowflake Cortex Search. Which of the following statements correctly describe the features and considerations when leveraging Snowflake Cortex Search for this RAG application?

- A. Cortex Search automatically handles text chunking and embedding generation for the source data, eliminating the need for manual ETL processes for these steps.
- B. To create a Cortex Search Service, one must explicitly specify an embedding model and manually manage its underlying infrastructure, similar to deploying a custom model via Snowpark Container Services.
- C. For optimal search results with Cortex Search, source text should be pre-split into chunks of no more than 512 tokens, even when using models with larger context windows like
D. The
- E. Enabling change tracking on the source table for the Cortex Search Service is optional; the service will still refresh automatically even if change tracking is disabled.

Answer: A,C,D

Explanation:

Option A is correct because Cortex Search is a fully managed service that gets users started with a hybrid (vector and keyword) search engine on text data in minutes, without needing to worry about embedding, infrastructure maintenance, or index refreshes. Option B is incorrect because Cortex Search is a fully managed service; users do not need to manually manage the embedding model infrastructure. A default embedding model is used if not specified. Option C is correct because, for best search results with Cortex Search, Snowflake recommends splitting text into chunks of no more than 512 tokens, as smaller chunks typically lead to higher retrieval and downstream LLM response quality, even with models that have larger context windows. Option D is correct because the 'SNOWFLAKE.CORTEX.SEARCH PREVIEW' function allows users to test the search service to confirm it is populated with data and serving reasonable results for a given query. Option E is incorrect because change tracking is required on the source table for the Cortex Search Service to function correctly and reflect updates to the base data.

NEW QUESTION # 260

.....

Most of the materials on the market do not have a free trial function. Even some of the physical books are sealed up and cannot be read before purchase. As a result, many students have bought materials that are not suitable for them and have wasted a lot of money. But GES-C01 guide torrent will never have similar problems, not only because GES-C01 exam torrent is strictly compiled by experts according to the syllabus, which are fully prepared for professional qualification examinations, but also because GES-C01 Guide Torrent provide you with free trial services. Before you purchase, you can log in to our website and download a free trial question bank to learn about GES-C01 study tool.

GES-C01 Reliable Exam Papers: https://www.real4test.com/GES-C01_real-exam.html

Any time if you want to know something about our products GES-C01: SnowPro® Specialty: Gen AI Certification Exam exam cram, we will serve for you immediately, Snowflake Question GES-C01 Explanations Your ability can be stimulated effectively and appropriately, and you would absorb those knowledge points easily, Snowflake Question GES-C01 Explanations If you have some questions, round-the-clock client support are waiting for you, We always strictly claim for our GES-C01 study materials must be the latest version, to keep our study materials up to date, we constantly review and revise them to be at par with the latest Snowflake syllabus for GES-C01 exam

Such direct manipulation of memory can disappear into the GES-C01 implementation of handles, such as 'vector', The lean startup methodology calls this the build measure learn loop.

Any time if you want to know something about our products GES-C01: SnowPro® Specialty: Gen AI Certification Exam exam cram, we will serve for you immediately, Your ability can be stimulated effectively and appropriately, and you would absorb those knowledge points easily.

Snowflake GES-C01 Exam Questions With Free Updates At 25% Discount

If you have some questions, round-the-clock client support are waiting for you, We always strictly claim for our GES-C01 study materials must be the latest version, to keep our study materials up to date, we constantly review and revise them to be at par with the latest Snowflake syllabus for GES-C01 exam

Real4test has formulated the SnowPro® Specialty: Gen AI Certification Exam (GES-C01) product in three versions.

- www.prepawayete.com GES-C01: The Penetration Tester's Guide Test Engine Copy URL  www.prepawayete.com  open and search for « GES-C01 » to download for free GES-C01 Valid Test Labs
- Quiz The Best Snowflake - GES-C01 - Question SnowPro® Specialty: Gen AI Certification Exam Explanations Open website  www.pdfvce.com  and search for "GES-C01" for free download Visual GES-C01 Cert Test
- Latest GES-C01 Test Pass4sure GES-C01 Flexible Testing Engine GES-C01 Updated Demo Go to website  www.examcollectionpass.com open and search for "GES-C01" to download for free GES-C01 Updated Demo
- Free PDF Quiz 2026 GES-C01: SnowPro® Specialty: Gen AI Certification Exam Updated Question Explanations Search for  GES-C01 and easily obtain a free download on www.pdfvce.com GES-C01 Latest Dumps Pdf
- GES-C01 Latest Dumps Pdf Latest GES-C01 Test Online Practice GES-C01 Exams Easily obtain free download of  GES-C01  by searching on [www.prepawaypdf.com] Practice GES-C01 Exams
- New GES-C01 Test Preparation New GES-C01 Exam Practice Latest GES-C01 Test Online Search for  GES-C01 and easily obtain a free download on [www.pdfvce.com] Latest GES-C01 Test Pass4sure
- Free PDF Snowflake - GES-C01 - Fantastic Question SnowPro® Specialty: Gen AI Certification Exam Explanations

Copy URL 「 www.prepawaypdf.com 」 open and search for (GES-C01) to download for free Test GES-C01 Simulator

BTW, DOWNLOAD part of Real4test GES-C01 dumps from Cloud Storage: <https://drive.google.com/open?id=1xFg9ub81SDY6kLOGNukAKTAIG74dAvSk>