

Simulation GES-C01 Questions & Exam GES-C01 Objectives

Snowflake GES-C01 Exam

SnowPro® Specialty: Gen AI Certification Exam

<https://www.passquestion.com/ges-c01.html>



Pass GES-C01 Exam with PassQuestion GES-C01 questions and answers in the first attempt.

<https://www.passquestion.com/>

1 / 23

BTW, DOWNLOAD part of BraindumpsPass GES-C01 dumps from Cloud Storage: <https://drive.google.com/open?id=173Z36u6F4Kmrc4i4KG-0Lu4CVZyvd47d>

If you are not aware of your problem, please take a good look at the friends around you! Now getting an international GES-C01 certificate has become a trend. If you do not hurry to seize the opportunity, you will be far behind others! Now the time cost is so high, choosing GES-C01 Exam Prep will be your most efficient choice. You can pass the GES-C01 exam in the shortest possible time to improve your strength.

Are you still feeling distressed for expensive learning materials? Are you still struggling with complicated and difficult explanations in textbooks? Do you still hesitate in numerous tutorial materials? GES-C01 study guide can help you to solve all these questions. GES-C01 certification training is compiled by many experts over many years according to the examination outline of the calendar year and industry trends. With GES-C01 Study Guide, you only need to spend 20 to 30 hours practicing to take the exam. In addition, GES-C01 certification training has a dedicated expert who updates all data content on a daily basis and sends the updated content to the customer at the first time. Therefore, using GES-C01 guide torrent, you don't need to worry about missing any exam focus.

>> Simulation GES-C01 Questions <<

Exam GES-C01 Objectives - Reliable GES-C01 Exam Online

It points to the exam heart to solve your difficulty. So high quality materials can help you to pass your exam effectively, make you feel easy, to achieve your goal. With the GES-C01 test guide use feedback, it has 98%-100% pass rate. That's the truth from our customers. And it is easy to use for you only with 20 hours' to 30 hours' practice. After using the GES-C01 Test Guide, you will have the almost 100% assurance to take part in an examination. With high quality materials and practices, you will get easier to pass the exam.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q39-Q44):

NEW QUESTION # 39

A data application developer is tasked with building a multi-turn conversational AI application using Streamlit in Snowflake (SiS) that leverages the COMPLETE (SNOWFLAKE.CORTEX) LLM function. To ensure the conversation flows naturally and the LLM maintains context from previous interactions, which of the following is the most appropriate method for handling and passing the conversation history?

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

Answer: E

Explanation:

To provide a stateful, conversational experience with the 'COMPLETE (SNOWFLAKE.CORTEX)' function (or its latest version, 'AI_COMPLETE'), all previous user prompts and model responses must be explicitly passed as part of the argument. This argument expects an array of objects, where each object represents a turn and contains a 'role' ('system', 'user', or 'assistant') and a 'content' key, presented in chronological order. In Streamlit, 'st.session_state' is the standard and recommended mechanism for storing and managing data across reruns of the application, making it ideal for maintaining chat history, by initializing 'st.session_state.messages = []' and appending messages to it. Option A is incorrect because 'COMPLETE' does not inherently manage history from external tables. Option B is incorrect as 'COMPLETE' does not retain state between calls; history must be explicitly managed. Option D is a less effective form of prompt engineering compared to passing structured history, as it loses the semantic role distinction and can be less accurate for LLMs. Option E describes a non-existent parameter for the 'COMPLETE' function.

NEW QUESTION # 40

A data engineer is developing an AI-infused data pipeline in Snowflake Notebooks to analyze Federal Reserve Meeting Minutes and official Statements, which are initially in PDF format. The goal is to determine the FED's stance on interest rates (hawkish, dovish, or neutral) and the reasoning for each ingested PDF using an LLM. The pipeline needs to automate data ingestion, text extraction, LLM inference, and store the results in a Snowflake table. Which sequence of operations and Snowflake features is most appropriate for building this pipeline within Snowflake?

- A. Ingest PDF documents into a directory table. Use 'Document AI' (PREDICT) to extract specific entities and tables from the PDFs into structured JSON. Then, create a 'STREAM' on the stage and a 'TASK' to continuously process new documents, extracting information and potentially performing additional sentiment analysis with another LLM.
- B. Load unstructured PDF files into an internal stage. Use a stored procedure to download new PDFs from the FOMC website. Leverage Snowpark Container Services to deploy a fine-tuned open-source LLM (e.g., Llama 2) for text extraction and sentiment analysis, and orchestrate the pipeline with 'Dynamic Tables' for continuous updates.
- C. Scrape PDF data from an external website, load unstructured PDF files to an internal stage, then use a 'UDF' to parse raw text from PDFs and a separate UDF ('GENERATE_PROMPT') to encapsulate a custom prompt. Finally, use a 'TASK' to automate the process, calling Snowflake's function with the custom prompt at the point of ingestion to generate the sentiment signal and reasoning.
- D. Directly ingest PDF documents into a 'VARIANT' column in a Snowflake table. Then, use the SQL function in 'OCR' mode to extract text and layout. The extracted text is then passed to 'SNOWFLAKE.CORTEX.CLASSIFY TEXT' to determine the sentiment, and the results are stored in a new table.
- E. Scrape data from an external website directly into a Snowflake table using an 'EXTERNAL FUNCTION'. Then, apply 'SNOWFLAKE.CORTEX.EXTRACT ANSWER' with a question like 'What is the FED's stance?' and 'SNOWFLAKE.CORTEX.SUMMARIZE' for reasoning to enrich the table. Automate this using 'STREAMS' and 'TASKS'.

Answer: A,C

Explanation:

Option A is correct. This option directly aligns with the 'AI-Infused Pipelines with Snowflake Cortex' blog post. It describes scraping data, loading to an internal stage, using a stored procedure to download new PDFs, and 'UDFS for parsing text and generating prompts. It explicitly mentions using 'Snowflake's TRY_COMPLETE function for LLM inference with a custom prompt at ingestion, all automated with 'Streams' and 'Tasks'. Option B is incorrect. PARSE DOCUMENT in 'OCR mode extracts text but does not preserve layout and primarily focuses on text and layout extraction, not directly sentiment analysis or complex reasoning. While 'CLASSIFY_TEXT (or can classify text, 'AI PARSE DOCUMENT is for extracting text and layout from documents, not direct ingestion into a 'VARIANT' column for text extraction via 'OCR mode. is a Cortex AI SQL function, typically used on files stored in stages. Option C is incorrect. While Snowpark Container Services can host LLMs and Dynamic Tables can automate updates, using an external LLM for text extraction when 'Document AI' or are available for native PDF processing is not the most direct approach. Additionally, 'Dynamic Tables' do not support incremental refresh with 'COMPLETE', and the prompt generation and sentiment analysis would still need to be explicitly defined. Option D is incorrect. Directly scraping into a Snowflake table with an 'EXTERNAL FUNCTION' isn't covered as the primary ingestion method for PDFs in the context of LLMs for sentiment. 'EXTRACT ANSWER and 'SUMMARIZE are task-specific Cortex functions, but the core task is a multi-step pipeline for PDF content analysis, which would be better served by a robust document processing solution. The focus of the pipeline example is on PDFs, not general website scraping directly to a table for immediate LLM application. Also, 'EXTRACT_ANSWER is for extracting a specific answer, while the sentiment is a classification. Option E is correct. This option uses 'Document AI', which is specifically designed to extract structured information (entities, tables) from unstructured documents like PDFs using 'Arctic-TILT'. It explicitly mentions creating a 'STREAM' on a stage and a 'TASK' for continuous processing, which is a standard pattern for Document AI pipelines. This approach directly handles the PDF extraction and structuring into JSON, which can then be further processed for sentiment analysis or reasoning if needed. The output JSON from '!PREDICT includes various extracted fields and can be parsed.

NEW QUESTION # 41

A financial analytics team is developing an application to extract specific, structured financial data (e.g., company name, revenue, profit margin) from various news articles using Snowflake Cortex LLM functions. They require the output to strictly conform to a predefined JSON schema and want to ensure robust error handling. Which of the following statements are crucial considerations for achieving this goal?

- A. To guarantee that critical fields like 'company name' and 'revenue' are always extracted, these properties must be explicitly listed within the 'required' array of the JSON schema provided to AI_COMPLETE.
- B. The complexity of the JSON schema provided to AI_COMPLETE has no impact on compute costs, as only the input text and generated content tokens are billed.
- C. The AI_COMPLETE function should be used with the response_format argument, supplying a JSON schema object that defines the required structure, data types, and constraints for the output.
- D. For enhanced reliability in production pipelines, the team should wrap their AI_COMPLETE calls within TRY_COMPLETE, as it returns a structured error object if the model fails to adhere to the schema, allowing for detailed debugging.
- E. Setting the temperature option to 0 in the AI_COMPLETE call is essential for obtaining the most consistent and accurate structured JSON outputs, regardless of task complexity or model used.

Answer: A,C,E

Explanation:

Option A is correct. AI_COMPLETE Structured Outputs allows specifying a JSON schema via the argument to ensure response_format responses follow a defined structure, data types, and constraints. Option B is correct. Using the field in the JSON schema ensures that required specified properties are extracted, or an error is raised by making extraction of critical information reliable. Option C is incorrect. COMPLETE, performs the same operation as COMPLETE (or AI_COMPLETE) but returns instead of raising an error when the operation cannot be TRY COMPLETE NULL performed. It does not return a structured error object for detailed debugging, but rather handles the error by returning allowing a pipeline to NULL, continue. Option D is correct. For the most consistent results and to optimize JSON adherence accuracy, it is recommended to set the temperature option to 0 when calling COMPLETE (or AI_COMPLETE). Option E is incorrect. The number of tokens processed (and billed) increases with schema complexity. A larger and more complex supplied schema generally consumes more input and output tokens, leading to higher compute costs.

NEW QUESTION # 42

A financial services company uses Snowflake Cortex's AI_COMPLETE for sentiment analysis on customer call transcripts, which contain personally identifiable information (PII). They also fine-tune a llama3.1-70b model with proprietary financial data. Which of

the following statements accurately describe Snowflake's Gen AI principles regarding data privacy, model usage, and governance in this scenario?

- A. The fine-tuned llama3.1-70b model, including the proprietary training data used, is exclusively owned by the financial institution and is not shared with other Snowflake customers.
- B. When using Cortex Analyst with Snowflake-hosted LLMs, metadata and prompts are transmitted outside Snowflake's governance boundary for processing, incurring additional cross-cloud data transfer costs.
- C. Enabling Cortex Guard for AI_COMPLETE automatically anonymizes all PII within the prompt before it reaches the LLM, regardless of the model chosen, to ensure privacy.
- D. Customer Data (inputs and outputs) for AI_COMPLETE, including PII, are guaranteed not to be available to other customers or used to train models made available to others.
- E. Snowflake's metadata fields, such as table and column names, should not contain personal, sensitive, or export-controlled data when using Snowflake AI services, to maintain data governance.

Answer: A,D,E

Explanation:

Option A is correct because Snowflake's trust and safety principles explicitly state that Usage and Customer Data (inputs and outputs of Snowflake AI Features) are NOT available to other customers and are NOT used to train, re-train, or fine-tune Models made available to others. Option B is correct because fine-tuned models built using your data can only be used by you, ensuring exclusivity. Option C is correct as customers should ensure no personal, sensitive, or regulated data is entered as metadata when using the Snowflake service. Option D is incorrect because Cortex Guard filters unsafe and harmful responses from the LLM; it does not automatically anonymize PII in the prompt before it reaches the LLM. Option E is incorrect because when Cortex Analyst is powered by Snowflake-hosted LLMs, data (including metadata or prompts) stays within Snowflake's governance boundary.

NEW QUESTION # 43

An organization enforces strict LLM access control. The has configured 'CORTEX_MODELS_ALLOWLIST='mistral-large2' and executed 'CALL SNOWFLAKE.MODELS.CORTEX_BASE_MODELS_REFRESH()'. A developer, whose role DEV_ROLE has been granted 'SNOWFLAKE.CORTEX_USER' and 'SNOWFLAKE.'CORTEX-MODEL-ROLE-CLAUDE-3-5- SONNET' , attempts to make a REST API call to 'api/v2/cortex/inference:complete' using 'claude-3-5-sonnet' as the model name in the request body. Which of the following statements are true regarding this scenario?

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

Answer: A,E

Explanation:

Option A is incorrect. When the model is specified as a plain string "claude-3-5-sonnet" in the REST API request, Cortex first attempts to match it as a schema-level model object by that plain name. If this fails (as is typical for built-in models which often use quoted identifiers like "CLAUDE-3-5-SONNET"), it then checks the plain name against the 'CORTEX_MODELS_ALLOWLIST. Since 'claude-3-5- sonnet' is not in the allowlist, the call will be denied despite the user's role having the application role for the model object 'SNOWFLAKE.MODELS."CLAUDE-3-5-SONNET". option B is correct because the parameter explicitly limits which models can be used by the 'Cortex LLM REST API'. Option C is correct. To use 'claude-3-5-sonnet' via the REST API, it either needs to be added to the account-level as a plain name, or the request must explicitly use the fully-qualified object identifier 'SNOWFLAKE.MODELS."CLAUDE-3-5-SONNET"' which would then be subject to RBAC on that model object. Option D is incorrect. An HTTP '403 Not Authorized' error typically indicates that the account is not enabled for the REST API or the calling user's default role lacks the 'SNOWFLAKE.CORTEX_USER database role. For 'AI_COMPLETE (which 'COMPLETE' maps to), an error related to an unallowed model would contain information about how to modify the allowlist, implying a different error message than a generic '403'. Option E is incorrect. The 'ENABLE CORTEX_ANALYST_MODEL_AZURE_OPENAI parameter is specifically for Cortex Analyst and determines if it can use Azure OpenAI models, but it is incompatible with model-level RBAC for Cortex Analyst. It does not affect generic 'COMPLETE (SNOWFLAKE.CORTEX)' REST API calls.

NEW QUESTION # 44

.....

Nowadays, online shopping has been greatly developed, but because of the fear of some uncontrollable problems after payment, there are still many people don't trust to buy things online, especially electronic products. But you don't have to worry about this when buying our GES-C01 Study Materials. Not only will we fully consider for customers before and during the purchase, but we will also provide you with warm and thoughtful service after payment. We have a special technical customer service staff to solve all kinds of consumers' problems.

Exam GES-C01 Objectives: <https://www.braindumps.com/Snowflake/GES-C01-practice-exam-dumps.html>

Snowflake Simulation GES-C01 Questions We firmly believe that you cannot be an exception, Snowflake Simulation GES-C01 Questions It is beneficial for those applicants who are busy in daily routines, Snowflake Simulation GES-C01 Questions Professionals be professionals, The content of our GES-C01 practice braindumps is chosen so carefully that all the questions for the exam are contained, Snowflake GES-C01 dumps come in a downloadable PDF format that you can print out and prepare at your own pace.

Making Innovation Work: How to Manage It, Measure It, and Profit Simulation GES-C01 Questions from It, Updated Edition, Management told the team to drop what they were doing and focus on resolving the crisis.

We firmly believe that you cannot be an exception, GES-C01 It is beneficial for those applicants who are busy in daily routines, Professionals be professionals, The content of our GES-C01 practice braindumps is chosen so carefully that all the questions for the exam are contained.

Study Through Online Snowflake GES-C01 Practice Test

Snowflake GES-C01 dumps come in a downloadable PDF format that you can print out and prepare at your own pace.

- Latest GES-C01 Exam Book Clear GES-C01 Exam Test GES-C01 Practice Search for { GES-C01 } and download it for free on www.examdiscuss.com website Certification GES-C01 Exam Cost
- GES-C01 Latest Exam Registration Certification GES-C01 Exam Cost Guaranteed GES-C01 Questions Answers Simply search for www.pdfvce.com GES-C01 for free download on (www.pdfvce.com) GES-C01 Test Simulator Online
- Pass Guaranteed Quiz Snowflake - GES-C01 - SnowPro® Specialty: Gen AI Certification Exam Authoritative Simulation Questions The page for free download of “ GES-C01 ” on { www.examdiscuss.com } will open immediately GES-C01 Valid Dumps Book
- GES-C01 Latest Exam Registration Clear GES-C01 Exam GES-C01 Test Simulator Online Search for GES-C01 and download exam materials for free through www.pdfvce.com Dumps GES-C01 Vce
- GES-C01 Accurate Prep Material Latest GES-C01 Exam Book Reliable GES-C01 Test Labs Easily obtain GES-C01 for free download through www.examcollectionpass.com Training GES-C01 Kit
- Test GES-C01 Practice GES-C01 Valid Dumps Book Clear GES-C01 Exam Enter (www.pdfvce.com) and search for [GES-C01] to download for free Test GES-C01 Practice
- GES-C01 Latest Exam Registration GES-C01 Latest Exam Registration Real GES-C01 Exam Answers Search for GES-C01 and obtain a free download on (www.easy4engine.com) Real GES-C01 Exam Answers
- Windows-based Snowflake GES-C01 Practice Exam Software Search for GES-C01 and obtain a free download on www.pdfvce.com Test GES-C01 Practice
- 2026 Snowflake GES-C01: SnowPro® Specialty: Gen AI Certification Exam Perfect Simulation Questions www.prepawayexam.com is best website to obtain GES-C01 for free download Learning GES-C01 Materials
- Intereactive GES-C01 Testing Engine GES-C01 PDF Download GES-C01 Intereactive Testing Engine Download GES-C01 for free by simply entering www.pdfvce.com website Real GES-C01 Exam Answers
- Three Easy-to-Use Formats of www.troytecdumps.com GES-C01 Exam Copy URL www.troytecdumps.com open and search for GES-C01 to download for free Real GES-C01 Exam Answers
- ameiivkg640274.governor-wiki.com, wjhsd.instructure.com, ok-social.com, brontewhqj359096.tkzblog.com, haimaygvw678211.wikicarrier.com, bookmarkfox.com, denissnay920387.thebindingwiki.com, iantzxr119447.angelinsblog.com, tedgkyx172821.onzeblog.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, Disposable vapes

What's more, part of that BraindumpsPass GES-C01 dumps now are free: <https://drive.google.com/open?id=173Z36u6F4Kmrc4i4KG-0Lu4CVZyvd47d>