2025 GES-C01: Fantastic SnowPro® Specialty: Gen AI Certification Exam Dump Collection

The ITexamReview guarantees their customers that if they have prepared with SnowPro® Specialty: Gen AI Certification Exam practice test, they can pass the SnowPro® Specialty: Gen AI Certification Exam (GES-C01) certification easily. If the applicants fail to do it, they can claim their payment back according to the terms and conditions. Many candidates have prepared from the actual Snowflake GES-C01 Practice Questions and rated them as the best to study for the examination and pass it in a single try with the best score.

Our SnowPro® Specialty: Gen AI Certification Exam (GES-C01) prep material also includes web-based and desktop SnowPro® Specialty: Gen AI Certification Exam (GES-C01) practice tests for you to put your skills to the test. Our SnowPro® Specialty: Gen AI Certification Exam (GES-C01) practice exams simulate the real Prepare for your SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam environment, so you can experience the pressure and environment of the actual test before the day arrives. You'll receive detailed feedback on your performance, so you know what areas to focus on and improve.

>> GES-C01 Dump Collection <<

Simulations Snowflake GES-C01 Pdf - Exam GES-C01 Assessment

When you see other people in different industry who feel relaxed with high salary, do you want to try another field? And is the difficulty of learning a new piece of knowledge often deterring you? It doesn't matter, now GES-C01 practice exam offers you a great opportunity to enter a new industry. Our GES-C01 learning material was compiled from the wisdom and sweat of many industry experts. And it is easy to learn and understand our GES-C01 exam questions.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q344-Q349):

NEW QUESTION # 344
A financial institution is deploying a sentiment analysis application that uses Snowflake Cortex 'SENTIMENT' and 'COMPLETE' functions, with different LLMs, for processing customer feedback. They are using AI Observability (Public Preview) to compare the cost- efficiency of using 'mistral-7b' versus 'claude-3-5-sonnet' as LLM judges for evaluation metrics, and also tracking the overall cost of their AI Observability usage. Which statements accurately reflect the cost implications and monitoring tools for this scenario?

- ☐ AI Observability incurs charges for LLM judges (e.g., `mistral-7b`, `claude-3-5-sonnet`) invoked via `COMPLETE (SNOWFLAKE.CORTEX)` calls to compute evaluation metrics, and these charges are based on 'tokens processed'.
- ☐ The `CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY` view should be used to monitor the credit consumption of the LLM judges specifically, as it tracks all Cortex function calls.
- ☐ Comparing models with vastly different context windows (e.g., `mistral-7b` at 32k tokens vs. `claude-3-5-sonnet` at 200k tokens) using AI Observability will not impact the billed 'tokens processed' if the actual prompt sizes are small and similar.
- ☐ In addition to LLM judge costs, warehouse charges are incurred for tasks managing evaluation runs and for queries used to compute evaluation metrics within AI Observability.
- ☐ The `METERING_DAILY_HISTORY` view, filtering by `SERVICE_TYPE ILIKE '%ai_services%'`, can provide an overview of daily credit usage for all AI services, including AI Observability LLM judge activity and associated warehouse costs.

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

Answer: A,B,E

Explanation:
Option A is correct because AI Observability utilizes LLM judges (such as 'mistral-7b' or 'claude-3-5-sonnet') through 'COMPLETE (SNOWFLAKE.CORTEX)' function calls to compute evaluation metrics, and these calls incur charges based on the 'tokens processed'. Option D is correct as, beyond LLM judge costs, AI Observability also incurs warehouse charges for managing evaluation runs and for queries that compute evaluation metrics. Option E is correct because the view, with a filter for 'SERVICE _ TYPE ILIKE , provides a comprehensive daily credit usage report for all AI services, which would include AI Observability's components. Option B is incorrect; the view is specifically for Document AI processing functions like '!PREDICT and 'AI_EXTRACT , not for general LLM judge usage in AI Observability. The view is more appropriate for tracking individual Cortex function calls. Option C is incorrect because while prompt sizes might be similar, the pricing for different LLMs (e.g., 'mistral-7b' at 0.12 credits per million tokens vs. 'claude-3-5-sonnet' at 2.55 credits per million tokens for AI Complete) will still result in different billed amounts due to varying per-token costs, even if the number of tokens is the same.

NEW QUESTION # 345
A company is developing a Streamlit application leveraging Snowflake Cortex Analyst for natural language querying over sales data. They want to implement a robust multi-turn conversational experience where users can ask follow-up questions. Which of the following statements accurately describe the design and cost implications of supporting multi-turn conversations in Cortex Analyst? (Select all that apply)

- A. When an LLM judge is used to evaluate the summarization quality for multi-turn conversations, a smaller model like Llama 3.1 8B is generally preferred over Llama 3.1 70B to minimise latency, even if it leads to a slightly higher error rate in rewritten questions.
- B. The cost for Cortex Analyst's multi-turn conversational support is primarily based on the number of messages processed, and the number of tokens within each message does not directly affect the per-message cost.
- C. An internal LLM summarization agent is automatically employed by Cortex Analyst before its original workflow to reframe follow-up questions based on conversation history, optimising LLM processing for each agent.
- D. Developers can manually implement multi-turn conversations in their applications by using the
- E. Cortex Analyst supports multi-turn conversations by simply passing the entire conversation history directly to every LLM call within its agentic workflow, which is the most efficient method for maintaining context.

Answer: B,C,D

Explanation:
Option A is incorrect. Cortex Analyst does not simply pass the entire conversation history to every LLM call. This primitive approach would lead to longer inference times, more non-determinism, and degraded performance. Instead, it employs an LLM summarization agent. Option B is correct. To support multi-turn conversations, Cortex Analyst adds an additional LLM summarization agent before its original workflow to reframe questions based on the conversation history. Option C is incorrect. While there's a latency-performance tradeoff, Llama 3.1 70B was found to be sufficient for the summarization task in Cortex Analyst, achieving 96.5% good ratings from an LLM judge, significantly outperforming Llama 3.1 8B which had an approximate 5% error rate in rewritten questions. Option D is correct. The credit rate usage for Cortex Analyst is based on the number of messages processed (67 Credits per 1 ,000 messages or 0.067 Credits per message), and the number of tokens in each message does not affect this per-message cost. Option E is correct. The 'SNOWFLAKE.CORTEX.COMPLETE function, and its 'TRY COMPLETE' variant, can be used to provide a stateful conversational experience by passing an array of objects for the argument, where each object contains a 'role' and 'content' for previous user prompts and model responses.

NEW QUESTION # 346
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' C!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. 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.
- C. 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.
- D. Scrape PDF data from an external website, load unstructured PDF files to an internal stage, then use a 'UDE 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.
- E. 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.

Answer: A,D

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 # 347
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?

☐ The REST API call with `"model": "claude-3-5-sonnet"` will successfully execute due to the `DEV_ROLE` having `SNOWFLAKE."CORTEX-MODEL-ROLE-CLAUDE-3-5-SONNET"`.

☐ The `CORTEX_MODELS_ALLOWLIST` parameter restricts access to models for `Cortex LLM REST API` calls.

☐ To successfully use `claude-3-5-sonnet` via the REST API, the `CORTEX_MODELS_ALLOWLIST` must be updated to include it, or the request should specify `"model": "SNOWFLAKE.MODELS.\"CLAUDE-3-5-SONNET\""`.

☐ An attempt to use an unallowed model via the `Cortex LLM REST API` would result in an HTTP `403 Not Authorized` error, even if the user's role has `SNOWFLAKE.CORTEX_USER`.

☐ The `ENABLE_CORTEX_ANALYST_MODEL_AZURE_OPENAI` parameter could be enabled to allow access to `claude-3-5-sonnet` for this `COMPLETE` REST API call.

- A. Option E
- B. Option C

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

Answer: B,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_OPENAF parameter is specifically for Cortex Analyst and determines if it can use Azure OpenAl 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 # 348
A developer is building a client application that interacts with a Snowflake Cortex Agent using its REST API. They are implementing multi- turn conversation support. Which of the following is the most critical aspect for maintaining conversational context over multiple API calls?

- A. A specific session_id parameter, generated at the start of the conversation, must be sent with each request to link turns.
- B. The agent internally summarizes long conversations, and the client receives only a condensed summary_token to pass for subsequent turns.
- C. The REST API automatically identifies and reuses context from previous requests based on the Authorization token, eliminating the need to pass full history.
- D. Only the last user prompt and the immediately preceding agent response should be sent to conserve token usage.
- E. The client application must include all prior user prompts and assistant/analyst responses in a messages array in each new API request.

Answer: E

Explanation:
Option A is the most critical and correct method for maintaining conversational state in multi-turn interactions with Cortex LLM functions, including those used by Cortex Analyst and by extension Cortex Agents. The underlying COMPLETE function (and its REST API equivalent) does not retain state from one call to the next; therefore, the client application must explicitly pass all previous user prompts and model (assistant/analyst) responses in chronological order within the 'messages' (or array for each new request to provide a stateful experience. Option B is incorrect as Cortex LLM functions do not automatically retain state across calls. Option C is incorrect; there is no documented 'session_id' parameter for implicitly managing conversation history for these APIs. Option D is incorrect; passing only partial history would lead to a loss of full conversational context and degrade the quality of follow-up responses. Option E is incorrect; while Cortex Analyst internally uses a summarization agent for long conversations to reframe questions, the client application is still responsible for managing and sending the full conversation history in the 'messages' array via the API.

NEW QUESTION # 349
.....

Simulations GES-C01 Pdf: https://www.itexamreview.com/GES-C01-exam-dumps.html

ITexamReview provides free trial for Snowflake Specialist GES-C01 exam (Snowflake Exam), Practicing for the GES-C01 practice test, again and again, can be nerve-wracking, so in this situation Exams, Snowflake GES-C01 Dump Collection As for a qualified worker and graduate, you need to learn many useful skills to meet the demands of the modern world, Snowflake GES-C01 Dump Collection It is essential to equip yourself with international admitted certifications.

In this process I also removed the brick wall that was on the right GES-C01 Exam Training side of the image, We had recently completed the Design Patterns book, but we all knew it was far from the final word on the topic.

Why do you need to Trust ITexamReview Snowflake GES-C01 Exam Questions?

ITexamReview provides free trial for Snowflake Specialist GES-C01 Exam (Snowflake Exam), Practicing for the GES-C01 practice test, again and again, can be nerve-wracking, so in this situation Exams.

As for a qualified worker and graduate, you need to learn many useful GES-C01 skills to meet the demands of the modern world, It is essential to equip yourself with international admitted certifications.

You have a wide choice without worrying about the Snowflake exam.

- Latest GES-C01 Mock Test □ GES-C01 Valid Exam Dumps □ Valid GES-C01 Test Review □ The page for free download of ⇛ GES-C01 ⇚ on 【 www.dumps4pdf.com 】 will open immediately □Reliable GES-C01 Test Cost
- 100% Pass Quiz 2025 GES-C01: Pass-Sure SnowPro® Specialty: Gen AI Certification Exam Dump Collection □ Immediately open " www.pdfvce.com " and search for 《 GES-C01 》 to obtain a free download □New GES-C01 Exam Experience
- Pass Guaranteed GES-C01 - SnowPro® Specialty: Gen AI Certification Exam –High-quality Dump Collection □ Copy URL ▹ www.lead1pass.com ◃ open and search for 《 GES-C01 》 to download for free □Reliable GES-C01 Test Cost
- GES-C01 Test Simulator Online □ GES-C01 Test Sample Online □ Exam GES-C01 Quick Prep □ Search for □ GES-C01 □ and download exam materials for free through 《 www.pdfvce.com 》 □Reliable GES-C01 Test Cost
- Board Your Capacities By Updated Snowflake GES-C01 Exam Dumps !! Open □ www.examcollectionpass.com □ and search for " GES-C01 " to download exam materials for free □GES-C01 Valid Exam Dumps
- Latest GES-C01 Test Preparation □ Test GES-C01 Tutorials □ Latest GES-C01 Test Preparation □ Open ➠ www.pdfvce.com □ enter { GES-C01 } and obtain a free download □GES-C01 New Braindumps
- Board Your Capacities By Updated Snowflake GES-C01 Exam Dumps □ Easily obtain free download of 《 GES-C01 》 by searching on 【 www.real4dumps.com 】 □GES-C01 New Braindumps
- Latest GES-C01 Study Plan □ Reliable GES-C01 Test Cost □ Reliable GES-C01 Test Cost □ Download ⇛ GES-C01 ⇚ for free by simply searching on { www.pdfvce.com } □GES-C01 PDF
- Free Download GES-C01 Dump Collection - Trustable GES-C01 Exam Tool Guarantee Purchasing Safety □ Simply search for ➥ GES-C01 □ for free download on " www.testsdumps.com " □GES-C01 Reliable Learning Materials
- Snowflake GES-C01 PDF Questions-Turn Your Exam Fear Into Confidence ✶ The page for free download of ➡ GES-C01 □□□ on 「 www.pdfvce.com 」 will open immediately ↙ Cheap GES-C01 Dumps
- Test GES-C01 Tutorials □ GES-C01 Reliable Learning Materials □ GES-C01 Online Tests □ Search for （ GES-C01 ） on ▹ www.torrentvce.com ◃ immediately to obtain a free download □GES-C01 Test Sample Online
- www.wcs.edu.eu, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, lms.ait.edu.za, www.myvrgame.cn, libict.org, samston182.bloginwi.com, motionentrance.edu.np, 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, lekoltoupatou.com, myskilluniversity.com, Disposable vapes