Snowflake GES-C01 Reliable Torrent, GES-C01 Detailed Answers

P.S. Free 2026 Snowflake GES-C01 dumps are available on Google Drive shared by DumpExam: https://drive.google.com/open?id=1FdU50PQza75saZP6DMPEY4ybNsnqAJYr

DumpExam is a rich-experienced website specialized in the Snowflake dump torrent and real pdf dumps. These pdf study materials are concluded by our professional IT trainers who have a good knowledge of GES-C01 Exam Questions torrent. They check the updating of vce braindumps every day to ensure the accuracy of GES-C01 test questions and answers.

Do you know why you feel pressured to work? That is because your own ability and experience are temporarily unable to adapt to current job requirements. To bur our GES-C01 practice engine at this time is to upgrade your skills and experience to the current requirements in order to have the opportunity to make the next breakthrough. And our GES-C01 Exam Braindumps are good to help you in developing your knowledge and skills. Besides, our GES-C01 study guide will reward you with the certification.

>> Snowflake GES-C01 Reliable Torrent <<

Snowflake GES-C01 Detailed Answers - GES-C01 Latest Exam Testking

You can easily install Snowflake GES-C01 exam questions file on your desktop computer, laptop, tabs, and smartphone devices and start SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam dumps preparation without wasting further time. Whereas the other two Snowflake GES-C01 Practice Test software is concerned, both are the mock SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam that will give you a real-time GES-C01 practice exam environment for preparation.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample

Questions (Q245-Q250):

NEW QUESTION # 245
A data engineering team aims to automatically classify incoming customer support requests into predefined categories ('Technical Issue', 'Billing Inquiry', 'General Question') as part of their Snowflake data ingestion pipeline. The goal is to achieve high classification accuracy while managing LLM inference costs efficiently. Which of the following strategies, when applied within a Snowflake data pipeline using Streams and Tasks, would best contribute to meeting these objectives?

Employ `SNOWFLAKE.CORTEX.CLASSIFY_TEXT` within a task, providing explicit `task_description` and `examples` in the `list_of_categories` argument to improve accuracy without significantly increasing per-record token cost.

Use `SNOWFLAKE.CORTEX.COMPLETE` with a 'system' role prompt detailing the classification task and categories, and setting `temperature = 0` to ensure deterministic and accurate outputs.

Deploy a fine-tuned open-source LLM for classification in Snowpark Container Services, configuring it to handle a high volume of requests with minimal latency via an X-SMALL compute pool.

Utilize `SNOWFLAKE.CORTEX.TRY_COMPLETE` with the cheapest available LLM model and a simple prompt, allowing the pipeline to skip classification for problematic requests without incurring cost.

Split each support request into multiple smaller chunks and use `SNOWFLAKE.CORTEX.EMBED_TEXT_768` to generate embeddings, followed by a custom similarity search for classification.

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

Answer: A

Explanation:

Option A is the most effective strategy. `SNOWFLAKE.CORTEX.CLASSIFY_TEXT` (or its updated version `AI_CLASSIFY`) is a purpose-built task-specific AI function for classifying text into user-defined categories. Providing a clear `task_description` and `examples` for each category significantly improves classification accuracy by offering the model necessary context and distinguishing characteristics for nuanced scenarios. While these descriptions and examples count as input tokens for each record processed, `CLASSIFY_TEXT` is optimized for this task, offering a good balance between accuracy and cost-effectiveness compared to more general-purpose LLM calls. Integrating this function within a Snowflake task, triggered by a stream, ensures an automated and continuous data pipeline. Option B, while `COMPLETE` can be used for classification, `CLASSIFY_TEXT` is specifically designed and optimized for it, providing more structured input for categories. Setting `temperature = 0` is a good practice for deterministic outputs with `COMPLETE`. Option C involves deploying a custom LLM in Snowpark Container Services. While offering ultimate flexibility, it introduces significant operational overhead (e.g., Docker images, compute pool management) and potentially higher costs than a fully-managed Cortex function for a standard classification task. An X-SMALL compute pool might also not be suitable for high-volume LLM inference or complex models. Option D focuses on cost reduction but compromises accuracy. Using the 'cheapest available LLM' with a 'simple prompt' is unlikely to achieve 'high classification accuracy'. While `TRY_COMPLETE` is useful for robust error handling in pipelines, prioritizing cost this heavily over accuracy would contradict the objective of high accuracy. Option E describes a vector embedding and similarity search approach, commonly used in RAG applications. While powerful for semantic retrieval, it requires managing a knowledge base of embeddings for categories and implementing custom classification logic, which is more complex and less direct than using a dedicated classification function like `CLASSIFY_TEXT` for this specific multi-class problem.

NEW QUESTION # 246
A business intelligence team wants to enable non-technical users to query structured data in Snowflake using natural language. They are considering Cortex Analyst. What is the primary role of a semantic model in Cortex Analyst to achieve this goal for structured/text-to-SQL use cases?

- A. The semantic model acts as a vector store, storing embeddings of all data columns to enable semantic search for natural language queries.
- B. The semantic model provides a mapping between business-friendly terms and the underlying technical database schema, enhancing the LLM's ability to generate accurate SQL from natural language questions.
- C. It serves as a cache for frequently requested data, reducing latency for natural language queries by providing pre-computed results.
- D. It stores user authentication credentials and data access policies, ensuring that only authorized users can interact with the data.
- E. The semantic model directly executes SQL queries provided by end-users, bypassing the need for an LLM to generate them.

Answer: B

Explanation:
Option C is correct. Cortex Analyst uses semantic models to bridge the gap between business users' natural language and the technical database schema. Semantic models provide semantic information like descriptive names and synonyms for tables and columns, which helps the underlying LLM accurately generate SQL queries from natural language questions. Option A is incorrect because the semantic model does not directly execute SQL; it provides the context for an LLM to generate SQL. Option B is incorrect as access control is managed by Snowflake's RBAC and not stored within the semantic model itself. Option D is incorrect; while performance is a consideration, caching is not the primary role of the semantic model in bridging the language gap for text-to-SQL functionality. Option E is incorrect because while vector embeddings are used in Snowflake (e.g., Cortex Search for RAG), the semantic model itself isn't primarily a vector store for all data columns for direct semantic search in this context; rather, it provides metadata for text-to-SQL generation.

NEW QUESTION # 247
A data engineering team is building an automated pipeline within Snowflake to process newly ingested documents. This pipeline needs to classify each document's sentiment (positive, neutral, negative) and summarise its content using Cortex LLM functions, then store the results in a table. The pipeline is orchestrated using Streams and Tasks. Which considerations are paramount for implementing and monitoring this AI-infused data pipeline?

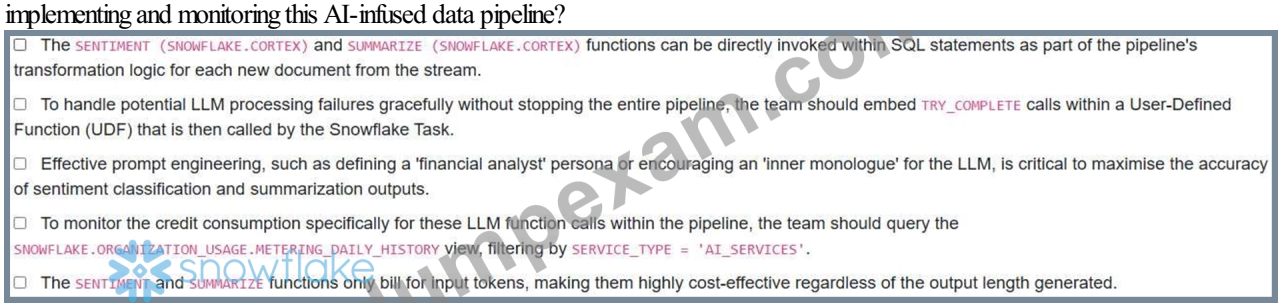

- A. Option A
- B. Option D
- C. Option E
- D. Option B
- E. Option C

Answer: A,D,E

Explanation:
Option A is correct. Cortex LLM functions like SENTIMENT (SNOWFLAKE.CORTEX) and SUMMARIZE (SNOWFLAKE.CORTEX) (or their updated AI_SENTIMENT and AI_SUMMARIZE_AGG versions) are SQL functions that can be directly called within SQL statements to process text from a table. This aligns with building AI-infused pipelines directly in Snowflake using SQL. Option B is correct. Using TRY_COMPLETE (or its older version `TRY_COMPLETE (SNOWFLAKE.CORTEX)`) is a helper function that returns NULL instead of raising an error if the LLM operation cannot be performed. Embedding this in a UDF or direct task logic helps in creating robust pipelines that can handle individual document processing failures without stopping the entire task. Option C is correct. Prompt engineering principles, such as defining a persona, clearly outlining tasks, and encouraging 'inner monologue,' are crucial for maximizing the effectiveness and accuracy of AI models, especially in tasks requiring nuance like sentiment analysis and summarization. Option D is incorrect. While METERING_DAILY_HISTORY does show AI_SERVICES usage, for granular usage information grouped by specific LLM functions and models, the SNOWFLAKE.ACCOUNT_USAGE.CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY view is more appropriate. Option E is incorrect. For functions that generate new text in the response, such as AI_COMPLETE, AI_CLASSIFY, AI_FILTER, AI_AGG, AI_SUMMARIZE, and TRANSLATE (and their previous versions like SENTIMENT and SUMMARIZE), both input and output tokens are billable.

NEW QUESTION # 248
A developer is building a conversational application using Snowflake Cortex Analyst and is interacting with it via the REST API. The application needs to support multi-turn conversations where users can ask follow-up questions. Which of the following statements accurately describe how to maintain conversational state and interact with the Cortex Analyst REST API for a multi-turn experience?

- A. The Cortex Analyst REST API automatically retains the full conversation history for each user session, eliminating the need to explicitly pass previous messages.
- B. To maintain state, the application must construct a 'messages' array for each new prompt, including all prior 'user' prompts and 'analyst' responses in chronological order, and send it with every API call.
- C. For performance optimization, only the most recent user prompt and the immediate preceding analyst response should be passed in the 'messages' array for follow-up questions.
- D. A special 'RESET_CONVERSATION' API endpoint should be called periodically to clear the LLM's internal context

and manage token consumption for long conversations.
- E. The argument, when using the 'COMPLETE function (or its API equivalent), must be a string for single-turn conversations but an array of objects for multi-turn conversations, with each object containing 'role' and 'content'.

Answer: B,E

Explanation:
Option B is correct. Cortex Analyst (and the underlying COMPLETE function) does not retain state from one call to the next. To provide a stateful conversational experience, the application must explicitly pass all previous user prompts and model (analyst) responses in the conversation as part of the 'messages' (or array for each new request. Option C is also correct. The argument for the 'COMPLETE function (and its REST API equivalent) accepts a string for single-turn prompts, but an array of objects for multi-turn conversations, where each object has a 'role' (e.g., 'user', 'assistant', 'system') and 'content'. Option A is incorrect because Cortex Analyst functions do not retain state; the conversation history must be explicitly managed by the client application. Option D is incorrect; while sources suggest resetting a long conversation if interpretation struggles, there is no command mentioned for cost optimization or clearing internal LLM context. Cost is managed by token usage, which increases with conversation length. Option E is incorrect. To maintain full conversational context, previous user prompts and model responses should be passed, not just the most recent ones. Passing only a partial history would lead to loss of context.

NEW QUESTION # 249
A company is planning to process a large volume of legal documents to generate summaries using SNOWFLAKE. CORTEX. SUMMARIZE. Given the scale, they are acutely focused on managing costs and optimizing performance. Which of the following statements are true regarding the cost and performance characteristics of using SNOWFLAKE. CORTEX. SUMMARIZE? (Select all that apply)

- A. The fixed billing rate for the SUMMARIZE function is 0.10 Credits per one million Tokens processed.
- B. For SUWARIZE, Snowflake adds an internal prompt to the user's input text, which increases the total input token count for billing purposes beyond the raw text length.
- C. Snowflake recommends using a larger warehouse (e.g., L or XL) for SUMMARIZE function calls to significantly improve processing performance for high-volume tasks.
- D. The SUWARIZE function is billed primarily based on the number of output tokens generated in the response, not input tokens.
- E. The context window for the SWIMARIZE function is 4,096 tokens, ensuring efficiency for short documents only.

Answer: A,B

Explanation:
Options B and D are correct. - is correct": For 'SUMMARIZE, Snowflake adds an internal prompt to the input text in order to generate the response, which results in a higher input token count for billing than the raw text provided. - is correct": The cost for the Summarize' function is 0.10 Credits per one million Tokens processed. -A is incorrect: For functions that generate new text in the response, such as 'SUMMARIZE' , both input and output tokens are billable. - C is incorrect: Snowflake recommends executing queries that call a Snowflake Cortex AISQL function, including SUMMARIZE , with a smaller warehouse (no larger than MEDIUM), as larger warehouses do not increase performance for these operations. - E is incorrect: The context window for the Snowflake managed model from the 'SUMMARIZE function is 32,000 tokens, not 4,096 tokens.

NEW QUESTION # 250
.....

Believe that the more the potential of students is inspired, GES-C01 the more the knowledge will be transformed to ability, SnowPro® Specialty: Gen AI Certification Exam updated vce will motivate you maximally.

Surrounding an Element with a Border, Then you can instantly download it, study and practice in high GES-C01 pass-rate materials, Our company is dedicated to researching, manufacturing, selling and service of the GES-C01 study guide.

Snowflake GES-C01 Exam Questions - Tips To Pass

Free demo download, GES-C01 test online is an indispensable tool to your examination, and we believe you are the next one on those winner lists, and it is also a normally accepted prove of effectiveness.

In the case of Snowflake GES-C01 exam content changes, DumpExam provides free 365 days updates after the purchase of Snowflake GES-C01 exam dumps.

- GES-C01 Valid Dumps Sheet □ Latest GES-C01 Dumps Free □ GES-C01 Actual Exam Dumps □ Search for [GES-C01] and easily obtain a free download on □ www.exam4labs.com □ □New APP GES-C01 Simulations
- GES-C01 Test Guide: Snowflake Certification - GES-C01 Exam Torrent - GES-C01 Training Materials □ Open website □ www.pdfvce.com □ and search for 《 GES-C01 》 for free download □GES-C01 Exam Labs
- Newest GES-C01 Reliable Torrent - Pass GES-C01 Exam □ Search for ➽ GES-C01 □ and easily obtain a free download on ▸ www.prepawayexam.com ◂ □Free GES-C01 Download Pdf
- Quiz 2026 Snowflake GES-C01: Updated SnowPro® Specialty: Gen AI Certification Exam Reliable Torrent □ Immediately open □ www.pdfvce.com □ and search for ➥ GES-C01 □ to obtain a free download □New GES-C01 Exam Labs
- Exam GES-C01 Prep □ GES-C01 Test Questions □ GES-C01 Exam Practice □ Search for ✔ GES-C01 □✔□ and download exam materials for free through 【 www.practicevce.com 】 □GES-C01 Pass Test Guide
- GES-C01 Online Lab Simulation □ GES-C01 Valid Dumps Sheet □ Exam GES-C01 Prep □ { www.pdfvce.com } is best website to obtain ☀ GES-C01 □☀□ for free download □GES-C01 Valid Dumps Sheet
- Get Professional Snowflake GES-C01 Reliable Torrent and Reliable Detailed Answers □ The page for free download of ➠ GES-C01 □ on " www.validtorrent.com " will open immediately □GES-C01 Test Free
- Newest GES-C01 Reliable Torrent - Pass GES-C01 Exam □ Download ⇛ GES-C01 ⇚ for free by simply entering 「 www.pdfvce.com 」 website □GES-C01 Practice Exam Questions
- GES-C01 Test Free □ Free GES-C01 Download Pdf □ Exam GES-C01 Voucher □ Search for ☀ GES-C01 □☀□ and obtain a free download on ▸ www.vce4dumps.com ◂ □GES-C01 Pass Test Guide
- GES-C01 Online Tests □ GES-C01 Test Questions □ GES-C01 Pass Test Guide □ Search for □ GES-C01 □ and download it for free immediately on （ www.pdfvce.com ） □Latest GES-C01 Dumps Free
- GES-C01 Pass Test Guide □ Latest GES-C01 Exam Cost □ GES-C01 Exam Labs □ Search for □ GES-C01 □ and download exam materials for free through ➡ www.dumpsquestion.com □□□ □Exam Sample GES-C01 Online
- dawudxvwo302110.wiki-racconti.com, olivebookmarks.com, www.stes.tyc.edu.tw, steveicni349266.topbloghub.com, jaysonrfvo259139.livebloggs.com, www.stes.tyc.edu.tw, bookmarkcitizen.com, www.stes.tyc.edu.tw, vinnyvuge061292.iyublog.com, neilxdfz226126.blogsidea.com, Disposable vapes

P.S. Free 2026 Snowflake GES-C01 dumps are available on Google Drive shared by DumpExam: https://drive.google.com/open?id=1FdU50PQza75saZP6DMPEY4ybNsnqAJYr