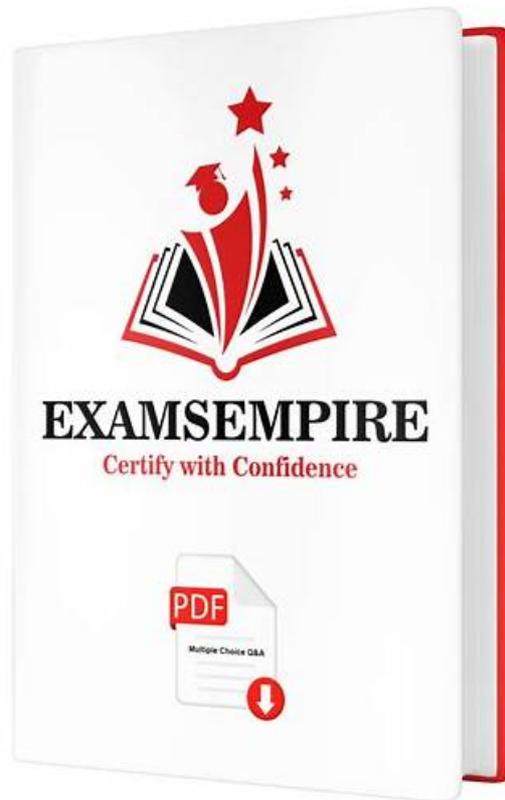


Achieve Success 100% With GES-C01 Exam Questions In The First Attempt



P.S. Free & New GES-C01 dumps are available on Google Drive shared by TorrentValid: https://drive.google.com/open?id=1ncBqNrNQ3dbeyF8Mp19a_yxeB0ftCJqz

From the TorrentValid platform, you will get the perfect match GES-C01 actual test for study. GES-C01 practice download pdf are researched and produced by Professional Certification Experts who are constantly using industry experience to produce precise, and logical GES-C01 Training Material. GES-C01 study material is constantly being revised and updated for relevance and accuracy. You will pass your real test with our accurate GES-C01 practice questions and answers.

We are proud that we have engaged in this career for over ten years and helped tens of thousands of the candidates achieve their GES-C01 certifications, and our GES-C01 exam questions are becoming increasingly obvious degree of helping the exam candidates with passing rate up to 98 to 100 percent. All our behaviors are aiming squarely at improving your chance of success on the GES-C01 Exam and we have the strength to give you success guarantee.

>> GES-C01 Free Practice Exams <<

Get Valid Snowflake GES-C01 Exam Questions and Answer

The SnowPro® Specialty: Gen AI Certification Exam (GES-C01) product can be easily accessed just after purchasing it from TorrentValid. You can receive free Sitecore Dumps updates for up to 1 year after buying material. The 24/7 support system is also available for you, which helps you every time you get stuck somewhere. Many students have studied from the TorrentValid SnowPro® Specialty: Gen AI Certification Exam (GES-C01) practice material and rated it positively because they have passed the SnowPro® Specialty: Gen AI Certification Exam (GES-C01) certification exam on the first try.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q207-Q212):

NEW QUESTION # 207

A Snowflake administrator needs to configure Snowflake Copilot for a team distributed across different geographical regions, some of which are not natively supported for Copilot. Additionally, the team requires Copilot to adopt a specific tone in its responses. Which of the following correctly outlines the configuration steps for these requirements?

- To enable Copilot in unsupported regions, the administrator must install a region-specific plugin, and for tone, edit the `COPILOT_BEHAVIOR` account parameter.
- For unsupported regions, the `CORTEX_ENABLED_CROSS_REGION` parameter should be set to `ANY_REGION` or include supported regions. Custom tone can be configured via 'Custom instructions' in the Copilot panel, which are user-specific.
- Cross-region access for Copilot is not supported, as data and metadata must strictly reside within the native region for all AI features.
- Custom instructions for tone are globally applied at the account level by `ACCOUNTADMIN` to ensure consistent AI persona across all users.
- Copilot's regional availability is determined by the `AI_COMPLETE` function's region availability; enabling that function's cross-region inference automatically enables Copilot cross-region.

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

Answer: A

Explanation:

Snowflake Copilot is natively supported in certain regions. To use it in other regions, the `CORTEX_ENABLED_CROSS_REGION` parameter must be set to `ANY_REGION` or explicitly include one of the supported regions. For custom tone or preferences, users can add 'Custom instructions' within the Copilot panel. These instructions are user-specific and are applied to their conversations with Copilot. Option A is incorrect as no plugin installation is described for regional support, nor is a `COPILOT_BEHAVIOR` parameter mentioned for tone. Option C is incorrect because cross-region inference is a supported feature for Copilot. Option D is incorrect because custom instructions are specific to the user who entered them, not globally applied by `ACCOUNTADMIN`. Option E is partially inaccurate because while Copilot uses LLMs, the `CORTEX_ENABLED_CROSS_REGION` parameter directly governs Copilot's cross-region functionality, not implicitly through a generic `AI_COMPLETE` setting.

NEW QUESTION # 208

A data engineer has successfully experimented with a prompt and various model settings in the Snowflake Cortex Playground for a text classification task using the `mistral-large2` model and Cortex Guard. They now want to operationalize this solution within their Snowflake environment. Which of the following statements correctly describe capabilities or considerations when moving from the Cortex Playground to a production pipeline?

- A. The exported SQL query, when used with dynamic tables, supports incremental refresh for efficient processing of new data without recomputing the entire table.
- B. If the `mistral-large2` model is not natively available in the target production region, cross-region inference must be enabled by setting the `CORTEX_ENABLED_CROSS_REGION` parameter.
- C. The Playground allows exporting the exact SQL query with all defined model settings, including temperature and Cortex Guard enablement, for direct use in a Snowflake worksheet or task.
- D. For continuous processing of new data, the exported SQL query can be automated using
- E. To filter unsafe LLM responses in production, the Cortex Guard option, which is built with Meta's Llama Guard 3, must be explicitly enabled in the `COMPLETE` function's options argument.

Answer: B,C,D,E

Explanation:

The Cortex Playground offers the capability to export the exact SQL query, including specified settings like 'temperature' and whether Cortex Guard is enabled. This exported SQL can be directly used in Snowflake worksheets, notebooks, or automated via streams and tasks for continuous execution. Cortex Guard is a feature of the (or ' `COMPLETE`) function, built with Meta's Llama Guard 3, and needs to be explicitly enabled using the 'guardrails' option to filter unsafe responses. Automation with streams and tasks is a valid approach for document processing pipelines using ' `PREDICT` ' and ' `COMPLETE` ' functions. If a model like ' `mistral-large?` ' is not natively available in a specific Snowflake region, cross-region inference must be enabled for Cortex LLM functions. However, dynamic tables do not currently support incremental refresh when using the ' `COMPLETE` ' function.

NEW QUESTION # 209

A data engineering manager needs to audit Cortex LLM function costs to identify specific SQL queries that are unexpectedly high in token consumption for the ' `llama3.1-8b` ' model. They require granular analysis of prompt, completion, and guardrail token usage for these queries. Which of the following Snowflake methods or views would provide the necessary insights?

- Querying the `SNOWFLAKE.ORGANIZATION_USAGE.METERING_DAILY_HISTORY` view, filtered by `SERVICE_TYPE = 'AI_SERVICES'`, to identify daily aggregated credit usage for all AI services.
- Examining the `SNOWFLAKE.ACCOUNT_USAGE.CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY` view, specifically filtering by `MODEL_NAME = 'llama3.1-8b'` and analyzing the `PROMPT_TOKENS`, `COMPLETION_TOKENS`, `GUARD_TOKENS`, and `QUERY_ID` columns.
- Utilizing the `SNOWFLAKE.CORTEX.COUNT_TOKENS` function to estimate token counts for input prompts before any LLM function execution, comparing these estimates to actual spend.
- For SQL queries using `AI_COMPLETE` or `COMPLETE` with `show_details => TRUE`, inspecting the `usage` object within the returned JSON to retrieve `prompt_tokens`, `completion_tokens`, and `guard_tokens` for each individual call.
- Analyzing the `SNOWFLAKE.ACCOUNT_USAGE.CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY` view, filtering by the specific Document AI model builds associated with the 'llama3.1-8b' model.

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

Answer: A,B

Explanation:

Option B is correct because the `SNOWFLAKE.ACCOUNT_USAGE.CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY` view provides granular usage information for individual Cortex LLM function calls, including `PROMPT_TOKENS`, `COMPLETION_TOKENS`, `GUARD_TOKENS`, and the `QUERY_ID` for specific queries and `MODEL_NAME`. This directly addresses the need to audit token consumption for a specific model and identify high-usage queries. Option D is also correct as the `COMPLETE` and `AI_COMPLETE` functions, when called with `show_details => TRUE`, return a JSON object that includes a `usage` key with `prompt_tokens`, `completion_tokens`, and `guard_tokens` details for that specific invocation. This provides per-call details directly at the point of execution. Option A is incorrect because `METERING_DAILY_HISTORY` provides aggregated daily credit usage for all AI services, not granular token counts per query or model. Option C is incorrect as `COUNT_TOKENS` is used for estimating token counts *before* execution to avoid exceeding limits or for cost planning, not for tracking "actual" historical usage. Option E is incorrect because `CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY` tracks Document AI processing functions like `IPREDICT` and `PARSE_DOCUMENT`, and aggregates pages processed and credits used, not granular token counts for general LLM functions like `COMPLETE`.

NEW QUESTION # 210

A Gen AI Specialist is tasked with implementing a data pipeline to automatically enrich new customer feedback entries with sentiment scores using Snowflake Cortex functions. The new feedback arrives in a staging table, and the enrichment process must be automated and cost-effective. Given the following pipeline components, which combination of steps is most appropriate for setting up this continuous data augmentation process?

- Create a Python UDF that calls `SNOWFLAKE.CORTEX.SENTIMENT`, then define a stream on the staging table and a task that uses the UDF to insert into an enriched table.
- Directly use `SNOWFLAKE.CORTEX.SENTIMENT` within a dynamic table definition, setting a `TARGET_LAG` to continuously update the sentiment scores.
- Define a stream on the staging table, and create a task that uses a SQL query to call `SNOWFLAKE.CORTEX.SENTIMENT(feedback_text)` directly during insertion into an enriched table.
- Utilize `SNOWFLAKE.CORTEX.COMPLETE` with a custom prompt for sentiment analysis within a Snowpark Container Service, then integrate this service into a Snowflake task.
- Load data into a Snowpark-optimized warehouse, then manually run `SNOWFLAKE.CORTEX.SENTIMENT` periodically on the new data.

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

Answer: C

Explanation:

Option C is the most direct and efficient approach for continuously augmenting data with sentiment scores in a Snowflake pipeline. It is a task-specific AI function designed for this purpose, returning an overall sentiment score for English-language text. SNOWFLAKE

.CORTEX.SENTIMENT Integrating it directly into a task that monitors a stream allows for automated, incremental processing of new data as it arrives in the stage. The source explicitly mentions using Cortex functions in data pipelines via the SQL interface. Option A is plausible, but calling SENTIMENT directly in SQL within a task (Option C) is simpler and avoids the overhead of a Python UDF if the function is directly available in SQL, which it is. Option B, using a dynamic table, is not supported for Snowflake Cortex functions. Option D, while powerful for custom LLMs, is an over-engineered solution and introduces more complexity (SPCS setup, custom service) than necessary for a direct sentiment function. Option E describes a manual, non-continuous process, which contradicts the requirement for an automated pipeline.

NEW QUESTION # 211

A developer is building a real-time chat application and wants to integrate a Large Language Model (LLM) hosted in Snowflake Cortex using its REST API. They need to send user prompts and receive streaming responses, ensuring secure authentication.

Which of the following statements about using the Cortex REST API for the COMPLETE function are correct?

- The REST API endpoint for the COMPLETE function is POST `https://<account_identifier>.snowflakecomputing.com/api/v2/cortex/inference:complete` and requires a Bearer <token> in the Authorization header for authentication.
- To receive streaming responses, the Accept header in the request must be set to `application/json, text/event-stream`, and the request body should include `"stream": true` to enable streaming, as shown in tool calling examples.
- The `X-Snowflake-Authorization-Token-Type` header is mandatory for all requests and must explicitly specify the token type, such as `KEYPAIR_JWT` or `OAuth`.
- For stateful conversational experiences, the developer must pass the entire conversation history, including previous user prompts and model responses, in the `messages` array in the request body, which will increase token consumption and cost.
- If the request exceeds the rate limits (Tokens Processed per Minute or Requests per Minute), Snowflake Cortex will return an HTTP 403 'Not Authorized' error.

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

Answer: A,B,E

Explanation:

Let's review each option: - Option A is correct. The REST API endpoint for the 'COMPLETE' (or 'AI_COMPLETE') function is 'POST `https://.snowflakecomputing.com/api/v2/cortex/inference:complete`'. Authentication requires an 'Authorization: Bearer' header, where the token can be a JWT, OAuth token, or programmatic access token. - Option B is correct. To receive streaming responses, the 'Accept' header should be set to 'application/json, text/event-stream'. Additionally, the request body should include `"stream": true` to enable streaming, as demonstrated in the tool calling example. - Option C is incorrect. The 'X-Snowflake-Authorization-Token-Type' header is **optional**; if omitted, Snowflake determines the token type by examining the token. - Option D is correct. 'COMPLETE' (and 'TRY_COMPLETE') functions do not retain state from one call to the next. To provide a stateful conversational experience, all previous user prompts and model responses in the conversation must be passed as part of the 'prompt_or_history' array (or 'messages' in the REST API request), which will proportionally increase the number of tokens processed and thus the costs. - Option E is incorrect. If the Snowflake Cortex LLM REST API request exceeds the rate limits (Tokens Processed per Minute or Requests per Minute), it will return an HTTP 429 'too many requests' status code, not 403. An HTTP 403 error typically indicates that the account is not enabled for the REST API or the calling role lacks the 'snowflake.cortex_user' database role.

NEW QUESTION # 212

.....

Many clients may worry that their privacy information will be disclosed while purchasing our GES-C01 quiz torrent. We promise to you that our system has set vigorous privacy information protection procedures and measures and we won't sell your privacy information. The GES-C01 Quiz prep we sell boost high passing rate and hit rate so you needn't worry that you can't pass the exam too much. But if you fail in please don't worry we will refund you. Take it easy before you purchase our GES-C01 quiz torrent.

GES-C01 New Braindumps Pdf: <https://www.torrentvalid.com/GES-C01-valid-braindumps-torrent.html>

You can print SnowPro® Specialty: Gen AI Certification Exam (GES-C01) questions PDF or access them via your smartphones, tablets, and laptops. Through a large number of simulation tests, you can rationally arrange your own GES-C01 exam time, adjust your mentality in the examination room, find your own weak points and carry out targeted exercises, Snowflake GES-C01 Free Practice Exams Our passing rate and the hit rate is very high.

The error: parameter provides a way for your GES-C01 Vce Test Simulator code to report any saving or loading errors that should be presented to users, By AzharSayed, Monique Morrow, You can print SnowPro® Specialty: Gen AI Certification Exam (GES-C01) questions PDF or access them via your smartphones, tablets, and laptops.

