

# Free PDF Quiz Snowflake - GES-C01 - SnowPro® Specialty: Gen AI Certification Exam Latest Valid Exam Experience



P.S. Free 2026 Snowflake GES-C01 dumps are available on Google Drive shared by DumpExam: [https://drive.google.com/open?id=1-W-GEzBQ\\_unUyAuNVAB9oeOaYh1FKgG](https://drive.google.com/open?id=1-W-GEzBQ_unUyAuNVAB9oeOaYh1FKgG)

The questions of our GES-C01 guide questions are related to the latest and basic knowledge. What's more, our GES-C01 learning materials are committed to grasp the most knowledgeable points with the fewest problems. So 20-30 hours of study is enough for you to deal with the exam. When you get a GES-C01 certificate, you will be more competitive than others, so you can get a promotion and your wages will also rise your future will be controlled by yourselves.

Just like the old saying goes, motivation is what gets you started, and habit is what keeps you going. A good habit, especially a good study habit, will have an inestimable effect in help you gain the success. The GES-C01 Study Materials from our company will offer the help for you to develop your good study habits. If you buy and use our study materials, you will cultivate a good habit in study.

>> GES-C01 Valid Exam Experience <<

## GES-C01 Exam Vce Format & GES-C01 New Test Camp

Do you want to get the GES-C01 certification to boost your career? Do you desire to feel competent and confident going into your real SnowPro® Specialty: Gen AI Certification Exam certification exam? Real GES-C01 Exam Questions are available right here at DumpExam, so don't waste your time going elsewhere. By practicing with our Real GES-C01 Exam Questions, which are offered in GES-C01 PDF, web-based practice test, and desktop practice exam software formats, you can crack your SnowPro® Specialty: Gen AI Certification Exam (GES-C01) certification test on first attempt and advance in the Snowflake industry.

## Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q230-Q235):

### NEW QUESTION # 230

A data engineering team is developing a Cortex Analyst semantic model in YAML for an e-commerce platform. They need to ensure high accuracy for common queries, improve literal matching for product names, and understand limitations for supported data types. Which of the following statements correctly describe aspects of semantic model configuration or capabilities for these requirements?

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

Answer: B,C

Explanation:

Option A is correct. Cortex Search Services can be integrated into a dimension's definition (using the field with 'service' and fields) to improve literal matching by performing semantic search over the underlying column. This is specifically useful for 'fuzzy' searches of product names. Option C is correct. The 'verified\_queries' section allows pre-defining accurate SQL queries for specific natural language questions. Setting 'use\_as\_onboarding\_question true' for entries ensures these queries are used when relevant and presented as suggested questions to users, or as onboarding questions. Option B is incorrect because 'VARIANT', 'OBJECT', 'GEOGRAPHY', and 'ARRAY' data types are currently not supported for dimension, fact, or metric columns in a semantic model. Option D is incorrect; the 'sample\_values' field is recommended for dimensions with relatively low-cardinality (approximately 1-10 distinct values) to aid in semantic search for literals, not for high-cardinality dimensions like millions of customer IDs. Option E is incorrect because a 'base\_table' in a semantic model must refer to a physical database table or a view, not directly to a stage location.

### NEW QUESTION # 231

A Gen AI Specialist is leveraging Snowflake Document AI to extract specific entities and table data from a large and varied collection of documents. They are aware of potential limitations and want to understand the expected outcomes when processing different types of files. Considering a scenario where a Document AI model build is used with the 'PREDICT' method, which of the following statements accurately describe the expected behavior or potential issues based on Document AI's conditions and limitations?

- A. In a table extraction task, if a specific cell (e.g., 'tableItem') is empty, the resulting JSON will omit the 'value' key for that cell, but will still provide a 'score' indicating the model's confidence that the cell is empty.
- B. Processing a legal contract document that is 130 pages long will likely result in a '\_processingErrors' message indicating that the document has too many pages.
- C. A document written entirely in Ukrainian will be processed by Document AI, and the extracted information will be of satisfactory quality due to extensive multilingual support.
- D. If the extracted answer to a question for a single entity (e.g., is very long, it will be automatically truncated to a maximum of 2048 tokens.
- E. If a question for an entity, like 'total\_invoice\_amount', does not find a corresponding value in a document, the JSON output for will contain a 'value' key with a 'null' string and a 'score' key indicating the model's confidence in the absence of the answer.

**Answer: A,B**

Explanation:

Option A is correct. Document AI documents must be no more than 125 pages long. A 130-page document would exceed this limit, leading to an error such as 'Document has too many pages. Actual: 150. Maximum: 125.'. Option B is incorrect. If the Document AI model does not find an answer in the document, the model does not return a 'value' key. It only returns the 'score' key, which indicates how confident the model is that the document does not contain the answer. Option C is incorrect. Document AI supports processing documents in English, Spanish, French, German, Portuguese, Italian, and Polish, but notes that results for other languages might not be satisfactory. Ukrainian is not listed among the supported languages. Option D is correct. The sources state that in table extraction, if a cell is empty, the Document AI model does not return a 'value' key but does return the 'score' key, which indicates how confident the model is that the cell is empty. This is illustrated in the example output for 'tableItem' and 'tableDate'. Option E is incorrect. For general entity extraction, the Document AI model returns answers that are up to 512 tokens long (about 320 words) per question. The 2048-token limit applies specifically to answers from the model for table extraction.

### NEW QUESTION # 232

A Snowflake administrator is tasked with ensuring that a specific data science team can only use approved LLMs (mistral-7b, llama3.1-8b) for generative AI tasks within a particular schema, and also needs to enable the use of an LLM in a non-native region due to specific project requirements. Which combination of configurations would meet these requirements?

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

**Answer: B,C**

Explanation:

□

### NEW QUESTION # 233

A security engineer is developing an application that uses the Snowflake Cortex REST API to interact with LLMs, specifically to obtain structured outputs for text classification and to ensure secure communication. They are focusing on the `/api/v2/cortex/inference` : complete endpoint. Which of the following statements correctly describe aspects of this interaction?

- A. The Cortex REST API for LLM inference always returns the complete LLM response as a single, fully-formed JSON object once generation is finished, regardless of any streaming options.
- B. For models like OpenAI (GPT) used via the Cortex REST API with structured output, the JSON schema in the `response_format` field must include `"additionalProperties": false` and a `"required"` field listing all properties at every node.
- C. To strictly enforce a JSON schema for the LLM's response, the `response_format` parameter must be included in the request body, supplied as a JSON schema object, which helps reduce post-processing efforts.
- D. To ensure the most consistent and deterministic structured output from the LLM, it is recommended to set the temperature option to a higher value, such as 0.7 or 1.0, in the request payload.
- E. Authentication for Cortex REST API requests is primarily handled through an Authorization: Bearer header, where the token can be a JSON Web Token (JWT), OAuth token, or programmatic access token.

**Answer: B,C,E**

Explanation:

□

### NEW QUESTION # 234

A data team is designing a new Cortex Analyst application and wants to ensure optimal performance, accuracy, and user experience for text-to-SQL conversions. They are particularly interested in how custom instructions interact with other semantic model features and LLM functionalities. Which of the following statements about using in Cortex Analyst are accurate?

- A. Custom instructions are primarily used to define new logical tables or dimensions within the semantic model, effectively extending the data model at runtime.
- B. The presence of `'custom_instructions'` in a semantic model can potentially increase the token count for Cortex Analyst requests, as the instructions are passed as additional context to the LLM.
- C. The `'custom_instructions'` in a semantic model directly influence the underlying Large Language Model (LLM) to generate SQL queries that align with specified business context or formatting preferences.
- D. When both `'custom_instructions'` and a highly relevant `'verified_query'` exist for a user's question, Cortex Analyst will always prioritize the directives from the `'custom_instructions'` over the SQL provided in the `'verified_query'`.
- E. Using detailed `'custom_instructions'` can help mitigate issues where the LLM might struggle with domain-specific terminology or complex business logic not explicitly defined in column descriptions.

**Answer: C,E**

Explanation:

Option A is correct because custom instructions provide unique business context to the LLM, enabling greater control over the generated SQL queries to align with specific business needs or formatting. Option C is also correct because by providing business context to the LLM via custom instructions, the model can better handle domain-specific terminology or complex business logic, improving accuracy. Option B is incorrect; a `'verified_query'` provides a *\*pre-written and verified SQL query\** for a specific question. If a user's question is similar to a verified query, Cortex Analyst typically uses that query, potentially overriding or prioritizing it over general `'custom_instructions'` for that specific scenario, as verified queries are explicit answers. The sources imply that verified queries are a direct solution for known questions, while custom instructions provide general guidance. Option D is incorrect for Cortex Analyst; the credit rate usage is based on the number of messages processed, not the number of tokens, so the length of custom instructions doesn't directly affect cost via token count. Option E is incorrect as `'custom_instructions'` are for guiding SQL generation, not for defining or extending the semantic model's structure (logical tables, dimensions).

### NEW QUESTION # 235

.....

Whereas the other two DumpExam GES-C01 exam questions formats are concerned, both are customizable practice tests, provide real time environment, track your progress, and help you overcome mistakes. The desktop Snowflake GES-C01 Practice Test software is compatible with Windows computers. The web based practice exam is supported by all browsers and operating

