

Hot GES-C01 Study Demo | Professional Snowflake GES-C01: SnowPro® Specialty: Gen AI Certification Exam 100% Pass



BONUS!!! Download part of ITExamReview GES-C01 dumps for free: <https://drive.google.com/open?id=1La4A6ZyxQkj8D8F-aa3BvniDqi-Z00hY>

We are well acknowledged for we have a fantastic advantage over other vendors - We offer you the simulation test with the Soft version of our GES-C01 exam engine: in order to let you be familiar with the environment of GES-C01 test as soon as possible. Under the help of the real simulation, you can have a good command of key points which are more likely to be tested in the real GES-C01 test. Therefore that adds more confidence for you to make a full preparation of the upcoming GES-C01 exam.

The ITExamReview wants to win the trust of SnowPro® Specialty: Gen AI Certification Exam (GES-C01) certification exam candidates. To achieve this objective ITExamReview is presenting Valid, Real, and Updated SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam questions in three different formats. These formats have high demand in the market and offer the easiest and quick way for SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam preparation.

>> GES-C01 Study Demo <<

New GES-C01 Exam Labs & GES-C01 Exam Preview

The Snowflake GES-C01 practice exam material is available in three different formats i.e Snowflake GES-C01 dumps PDF format, web-based practice test software, and desktop GES-C01 practice exam software. PDF format is pretty much easy to use for the ones who always have their smart devices and love to prepare for GES-C01 Exam from them. Applicants can also make notes of printed SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam material so they can use it anywhere in order to pass Snowflake GES-C01 Certification with a good score.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q201-Q206):

NEW QUESTION # 201

A machine learning team is leveraging the Snowflake Model Registry to manage diverse models, including a custom Python utility for data preprocessing that they wish to make available as a model method. Which of the following statements accurately describe capabilities or considerations when logging models and their associated artifacts and methods in the Model Registry?

The Snowflake Model Registry supports built-in types such as Scikit-learn, XGBoost, and PyTorch, but does not allow logging custom Python objects or processing code directly as models.

To include additional local files, such as configuration files or custom scripts, with a logged model, the `user_files` argument must be used in `log_model`, mapping stage subdirectories to local file paths.

Once a model version is logged, its methods can be invoked using either `mv.run()` in Python or through service functions named `<service_name>!<method_name>` in SQL, after the model has been deployed to SPCS.

The `function_type` option within `method_options` in `log_model` allows specifying whether a model method should be exposed as a `FUNCTION` or `TABLE_FUNCTION` in SQL, influencing how data is processed.

The maximum total model size for models deployed to a Snowflake warehouse is 5 GB, whereas models deployed to SPCS have no such size limitations.

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

Answer: A,C,D,E

Explanation:

Option B is correct because the `user_files` argument in `log_model` is specifically designed for including additional local files (e.g., configuration, scripts) with the model, by mapping them to stage subdirectories. Option C is correct because after a model is deployed to SPCS, Snowflake Model Serving creates service functions named `<service_name>!<method_name>` that allow invoking the model's methods from SQL. Model methods can also be called via `mv.run()` in Python. Option D is correct because the `function_type` option within `method_options` allows developers to specify if a model method should be registered as a `FUNCTION` or `TABLE_FUNCTION` when exposed in SQL, affecting its input/output processing. Option E is correct. The maximum total model size for warehouse-deployed models is 5 GB. For models running on Snowpark Container Services, these size restrictions are eased or eliminated, allowing for much larger models. Option A is incorrect; the Model Registry is flexible enough to support not only built-in types (like Scikit-learn, XGBoost, PyTorch) but also custom processing code and previously-trained models.

NEW QUESTION # 202

A data engineer is tasked with establishing a robust MLOps pipeline using the Snowflake Model Registry. They have trained a scikit-learn model and need to log it. Which of the following statements correctly describes a 'required' step or privilege for successfully logging a model using the 'Registry.log_model' method?

The `snowflake.ml.registry.Registry` object must always be initialized by explicitly providing both the `database_name` and `schema_name` arguments, even if the session context is already set.

The role used to log the model must be granted the `CREATE MODEL` privilege on the specific schema where the model will be stored.

All model dependencies must be provided exclusively through the `conda_dependencies` argument to ensure compatibility with Snowflake's internal package management.

The `model` argument in `log_model` can accept either a Python model object or a string path to a serialized model file (e.g., `.pkl`).

The `log_model` method requires that the `target_platforms` argument is always set to include `WAREHOUSE`, even if the model is exclusively intended for Snowpark Container Services.

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

Answer: E

Explanation:

Option B is correct because to create a model (which is what logging a model does), the role used must have the `CREATE MODEL` privilege on the schema where the model is being created, or 'OWNERSHIP' on that schema. Option A is incorrect because `database_name` and `schema_name` are optional arguments when initializing the `Registry` object; if not provided, the current database and schema of the session are used. Option C is incorrect as the `log_model` method supports both `conda_dependencies` and `pip_requirements`. Option D is incorrect; the `model` argument expects the Python model object itself, not a path to a serialized file. Serialization is handled automatically for supported types like scikit-learn. Option E is incorrect; the `target_platforms` argument can be set to `['SNOWPARK_CONTAINER_SERVICES']` for models intended solely for SPCS. If 'WAREHOUSE' is specified and the model is ineligible for warehouse deployment, `log_model()` will fail.

NEW QUESTION # 203

An AI developer is building a Snowflake data pipeline to prepare unstructured data for a RAG application. The pipeline involves

extracting text, splitting it into chunks, generating embeddings, and then indexing for Cortex Search. Considering the role of helper functions like SNOWFLAKE.CORTEX.SPLIT_TEXT_RECURSIVE_CHARACTER, which of the following statements accurately describes its typical operational placement and interaction within this Gen AI pipeline?

- A. The function's recursive nature enables it to automatically detect and correct factual inconsistencies or 'hallucinations' present in the original large text documents before they are embedded.
- B. It is typically applied after an embedding function (e.g.,
- C. Its output, consisting of smaller text chunks, serves as the direct input for text embedding functions that then convert these chunks into vector representations for semantic indexing.
- D. It replaces the need for
- E. It is a post-processing step for LLM-generated responses, used to break down long answers into digestible paragraphs for user display in chat interfaces.

Answer: C

Explanation:

Option B is correct.

is a helper function used to divide large text documents into smaller chunks. These smaller text chunks are then processed by embedding functions, such as

`SNOWFLAKE.CORTEX.EMBED_TEXT_768`

or



, to create vector embeddings that are subsequently used for indexing and semantic search in RAG applications. Option A is incorrect because text splitting (chunking) happens *before* embedding generation, not after, as it prepares the raw text for vectorization. Option C is incorrect; COUNT_TOKENS is a separate helper function specifically designed to return the token count of input text. SPLIT_TEXT_RECURSIVE_CHARACTER does not implicitly provide token counts for its output chunks. Option D is incorrect; the function's purpose is text splitting, not hallucination detection or correction, which pertains to LLM output quality. Option E is incorrect; while LLM responses might be formatted, the primary role of SPLIT_TEXT_RECURSIVE_CHARACTER is in preparing input documents for RAG, not post-processing LLM outputs for display.

NEW QUESTION # 204

A security auditor needs to access and analyze logs generated by Snowflake AI Observability for compliance auditing and to track the activity of generative AI applications. They need to understand how to reliably query this data and its temporal characteristics within Snowflake. Which of the following statements accurately describes the access and characteristics of this logged data?

- A. Detailed request and response bodies, along with the generated SQL, are stored and can be directly queried using standard SQL.
- B. Logged data from AI Observability's event tables becomes visible within a small latency, typically 1-2 minutes, after a request is made.
- C. Access to these detailed event tables is implicitly granted to roles holding the SNOWFLAKE.CORTEX_USER database role and the AI_OBSERVABILITY_EVENTS_LOOKUP application role.
- D. The logs are automatically purged after 7 days of being recorded, requiring a separate process for long-term data retention.
- E. Logs are exclusively available for analysis through pre-built dashboards in Snowsight and cannot be accessed via direct SQL queries.

Answer: A,B,C

Explanation:

Snowflake AI Observability features logging of application traces and Cortex Analyst logs requests to an event table in the Snowflake database. There is a small latency of '1-2 minutes' before these logged requests are visible, making option A correct. The logs include detailed information such as 'Generated SQL' and 'Request and response bodies', which are stored and can be queried, validating option C. The necessary roles for AI Observability, including 'SNOWFLAKE.CORTEX_USER' and 'EVENTS_LOOKUP', are required for creating and executing runs, which implies they grant access to the generated logs for monitoring, making option D correct. Option B is incorrect as the sources do not mention an automatic 7-day purge for these logs. Option E is incorrect because the documentation includes a subheading 'Querying logs with SQL' for Cortex Analyst administrator

monitoring, indicating that direct SQL access is supported.

NEW QUESTION # 205

A data engineering team is building an automated pipeline in Snowflake to process customer reviews. They need to use AI_COMPLETE to extract specific details like product, sentiment, and issue type, and store them in a strictly defined JSON format for seamless downstream integration. They aim to maximize the accuracy of the structured output and manage potential model limitations. Which statements accurately reflect the best practices and characteristics when using AI_COMPLETE with structured outputs for this scenario?

- Specifying the `response_format` argument in `AI_COMPLETE` incurs an additional, separate compute cost for the overhead of verifying each token against the provided JSON schema.
- For complex data extraction tasks requiring high accuracy, it is recommended to explicitly include phrases like 'Respond in JSON' within the prompt, in addition to defining the `response_format` schema.
- When using OpenAI (GPT) models with structured outputs, the JSON schema defined in `response_format` must explicitly include `"additionalProperties": false` in every node and ensure the `"required"` field lists all properties in the schema.
- If the model's generated response slightly deviates from the specified JSON schema, `AI_COMPLETE` with structured outputs will automatically reformat the output to conform, preventing validation errors.
- To achieve the most consistent and deterministic results when using `AI_COMPLETE` for structured outputs, the `temperature` option should be set to `0`.

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

Answer: A,C,E

Explanation:

Option A is incorrect because Structured Outputs do not incur additional compute cost for the overhead of verifying each AI_COMPLETE token against the supplied JSON schema, though the number of tokens processed (and thus billed) can increase with schema complexity. Option B is correct because for complex reasoning tasks, it is recommended to use the most powerful models and explicitly add 'Respond in JSON' to the prompt to optimize accuracy. Option C is correct as for OpenAI (GPT) models, the schema has specific requirements: `response_format` must be set to `in` in every node, and the field must include the names of every property in the schema. Option D `additionalProperties false` required is incorrect because verifies each generated token against the JSON schema to ensure conformity, and if the model cannot generate a AI_COMPLETE response that matches the schema, it will result in a validation error. Option E is correct as setting the option to `0` is recommended for temperature the most consistent results, regardless of the task or model, especially for structured outputs.

NEW QUESTION # 206

.....

It is known to us that the error correction is very important for these people who are preparing for the GES-C01 exam in the review stage. It is very useful and helpful for a lot of people to learn from their mistakes, because many people will make mistakes in the same way, and it is very bad for these people to improve their accuracy. If you want to correct your mistakes when you are preparing for the GES-C01 Exam, the study materials from our company will be the best choice for you.

New GES-C01 Exam Labs: <https://www.itexamreview.com/GES-C01-exam-dumps.html>

For candidates who are going to buy GES-C01 learning materials online, they may pay more attention to that money safety, SnowPro® Specialty: Gen AI Certification Exam: Exam Ref GES-C01, Snowflake GES-C01 Study Demo Please keep focus on your email boxes, The exam GES-C01 simulation will mark your mistakes and help you play well in GES-C01 pdf review guide, Our GES-C01 practice questions are the best study materials for the exam in this field, we will spare no effort to help you pass the exam as well as getting the related certification.

For these great merits we can promise to you that if you buy our GES-C01 study materials you will pass the test without difficulties, Our GES-C01 Study Materials could make a difference to your employment prospects.

Complete coverage GES-C01 Online Learning Environment

For candidates who are going to buy GES-C01 learning materials online, they may pay more attention to that money safety,

SnowPro® Specialty: Gen AI Certification Exam: Exam Ref GES-C01, Please keep focus on your email boxes.

The exam GES-C01 simulation will mark your mistakes and help you play well in GES-C01 pdf review guide, Our GES-C01 practice questions are the best study materials for the exam in this field, GES-C01 we will spare no effort to help you pass the exam as well as getting the related certification.

BONUS!!! Download part of ITExamReview GES-C01 dumps for free: <https://drive.google.com/open?id=1La4A6ZyxQkj8D8F-aa3BvniDqi-Z00hY>