# Pass Guaranteed 2026 Snowflake GES-C01: Pass-Sure SnowPro® Specialty: Gen AI Certification Exam Actual Test Pdf


The SnowPro® Specialty: Gen AI Certification Exam (GES-C01) actual questions we sell also come with a free demo. Spend no time, otherwise, you will pass on these fantastic opportunities. Start preparing for the SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam by purchasing the most recent Snowflake GES-C01 exam dumps. You must improve your skills and knowledge to stay current and competitive. You merely need to obtain the GES-C01 Certification Exam badge in order to achieve this. You must pass the SnowPro® Specialty: Gen AI Certification Exam exam to accomplish this, which can only be done with thorough exam preparation. Download the SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam questions right away for immediate and thorough exam preparation.

Keep making progress is a very good thing for all people. If you try your best to improve yourself continuously, you will that you will harvest a lot, including money, happiness and a good job and so on. The GES-C01 preparation exam from our company will help you keep making progress. Choosing our GES-C01 Study Material, you will find that it will be very easy for you to overcome your shortcomings and become a persistent person. Just come and buy our GES-C01 learning guide!

**>> GES-C01 Actual Test Pdf <<**

## Snowflake certification GES-C01 best exam questions and answers

If you want to get satisfaction with the preparation and get desire result in the GES-C01 real exam then you must need to practice our Snowflake braindumps and latest questions because it is very useful for preparation. You will feel the atmosphere of GES-C01 Actual Test with our online test engine and test your ability in any time without any limitation. There are also GES-C01 free demo in our website for you download.

## Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q173-Q178):

**NEW QUESTION # 173**
A data engineer is tasked with defining a semantic model for Cortex Analyst to enable natural language queries over sales dat a. They are creating a YAML file to describe the logical structure. Which of the following statements correctly describe the configuration of this semantic model? (Select all that apply)

- A. The 'tables' section must define logical tables that directly map to Snowflake physical tables or views, with 'base_table' specifying the fully qualified name (database, schema, table).
- B. To ensure high accuracy for specific common questions, 'verified_querieS can be included, where the 'sqr field must reference the underlying physical table and column names (ee.g., not the logical names defined in the semantic model.
- C. Dimensions such as 'product_category' can include a configuration, which can specify an optional 'literal_column' that, if omitted, defaults to the search index column.
- D. A 'metric' like 'total_revenue' can be defined using an 'expr' that references other logical columns, such as 'sum(profity from a logical 'sales_data' table, within the semantic model.

- E. The 'VARIANT data type is a supported data type for columns defined as 'dimensions' or 'facts' within the semantic model, allowing flexible storage of complex attributes.

**Answer: A,C,D**

Explanation:
Option A is correct. A logical table, a foundational concept of Snowflake's semantic model, represents either a physical database table or a view and its 'base_table' field specifies the fully qualified name of the underlying physical table. Option B is correct. Dimensions can specify a block to integrate with Cortex Search, and the 'literal_column' field within this block is optional and defaults to the search index. Option C is incorrect. The 'VARIANT, 'OBJECT, GEOGRAPHY' , and 'ARRAY' data types are explicitly not supported for dimensions, time dimensions, or facts within a semantic model. Option D is incorrect. Verified queries must use the names of the logical tables and columns as defined in the semantic model, not those in the underlying physical dataset. For example, sales_data' for a logical table named 'sales_data'. Option E is correct. Metrics can be defined using an SQL expression ('expr') that can reference logical columns (facts, dimensions, or time dimensions) within the same logical table or from another logical table in the semantic model.

**NEW QUESTION # 174**
A Snowflake team observes consistently high token costs from 'SNOWFLAKE.ACCOUNT USAGE.CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY' for a summarization task using the 'mistral- large? model. The task involves summarizing legal documents, which often exceed the context window of common LLMs. To optimize these token-based costs, which strategy should the team prioritize?

- A. Implement a text splitting mechanism, potentially using SPLIT_TEXT_RECURSIVE_CHARACTER, to break down lengthy documents into smaller chunks before passing them to the summarization function, then aggregate the summaries.
- B. Increase the virtual warehouse size (e.g., from X-SMALL to MEDIUM) used for running the summarization queries to boost performance and reduce overall cost per query.
- C. Set the temperature parameter to 0 in the COMPLETE function options to ensure more deterministic and thus more cost-efficient summarization outputs.
- D. Enable Cortex Guard for the COMPLETE function calls, as its filtering capabilities automatically reduce the number of tokens processed for unsafe content.
- E. Switch from using the COMPLETE function to TRY_COMPLETE to automatically avoid billing for queries that fail due to context window limits, thereby reducing costs.

**Answer: A**

Explanation:
Option C is correct. For summarization of lengthy documents, exceeding the context window or using large inputs significantly increases token consumption. Text splitting, for example using can break documents into smaller, more manageable chunks. This reduces the number of input tokens per LLM call, directly leading to cost optimization, and is recommended for best search results and LLM response quality. Option A is incorrect because for Cortex AISQL functions, Snowflake recommends using a smaller warehouse (no larger than MEDIUM) as larger warehouses do not increase performance but can result in unnecessary costs associated with keeping the warehouse active. The compute cost for Cortex LLM functions is based on tokens processed, not warehouse size performance. Option B is incorrect because only prevents costs for 'failed' operations by returning NULL instead of an error. It does not optimize the token consumption of 'successful' summarization tasks. Option D is incorrect; Cortex Guard processes additional tokens for its filtering, thus 'increasing' token consumption, not reducing it. Option E is incorrect because setting 'temperature' to 0 makes the output more deterministic, which might improve consistency but does not directly reduce the number of input or output tokens processed for a summarization task.

**NEW QUESTION # 175**
A data platform architect is integrating 'SNOWFLAKE.CORTEX.EMBED TEXT 768' into a complex data pipeline for a new search application. The pipeline involves extracting text from various sources, generating embeddings, storing them in Snowflake, and performing semantic searches. Which of the following statements accurately describes a compatibility aspect or limitation when working with 'EMBED TEXT 768' and the resulting 'VECTOR' data type within Snowflake?

- A. The function can be directly integrated into a dynamic table's 'SELECT statement to provide continuous, automated embedding updates for new data.
- B. When is invoked within a Snowpark Python User-Defined Function (UDF) on Snowflake data, the data remains within Snowflake's network boundary during the embedding generation process.
- C. To support diverse embedding dimensions from different models, the 'VECTOR data type can be stored efficiently within

a 'VARIANT column, which automatically handles schema variations.
- D. If the function is not natively available in the account's primary Snowflake region, cross-region inference cannot be enabled, thus preventing its use.
- E. The 'VECTOR' data type, which stores the output of is fully compatible with all Snowflake features, including being used as a primary key in hybrid tables for fast lookups.

**Answer: B**

Explanation:
Option D is correct. When Snowflake Cortex LLM functions, such as 'EMBED_TEXT_768', are called on Snowflake data (e.g., within a Snowpark Python UDF), the data never actually leaves Snowflake's network boundary. This ensures that data governance and security are maintained. Option A is incorrect because Snowflake Cortex functions, including 'EMBED_TEXT_768' , do not support dynamic tables. Option B is incorrect; cross-region inference can be enabled if ' is not natively available in a region, using the 'CORTEX_ENABLED_CROSS_REGION' parameter. Option C is incorrect because the 'VECTOR data type is not supported as primary or secondary index keys in hybrid tables. Option E is incorrect because 'VECTOR data types are explicitly not supported in 'VARIANT' columns.

**NEW QUESTION # 176**
A data engineer is setting up a Document AI pipeline to extract information from scanned invoices stored in an internal stage named 'invoice_stage'. They have created the stage using 'CREATE STAGE and uploaded several PDF documents. However, when attempting to run the extraction query, they encounter an error message: 'File extension does not match actual mime type. Mime-Type: application/octet-stream'. Additionally, they anticipate a privilege issue might arise for pipeline automation. Which of the following conditions must be met to resolve the current error and ensure proper setup for Document AI extraction and subsequent pipeline creation?

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

**Answer: A,D**

Explanation:

**NEW QUESTION # 177**
A data analyst is setting up a new Cortex Analyst-powered conversational app for business users. They want to understand how the "Suggested Questions" feature behaves under different semantic model configurations to ensure an optimal user experience. Which of the following statements accurately describe the behavior of the "Suggested Questions" feature in Cortex Analyst based on the semantic model configuration?

- A. If multiple verified queries in a VQR are marked with
- B. If the semantic model does not include a Verified Query Repository (VQR), Cortex Analyst will always return a blank list of suggested questions.
- C. When a semantic model contains a VQR, Cortex Analyst prioritizes returning up to five suggested questions from the VQR that are semantically similar to the user's input.
- D. In the absence of a VQR, Cortex Analyst uses underlying Large Language Models (LLMs) to generate up to three suggested questions, which are guaranteed to be answerable.
- E. Setting the

**Answer: A,C**

Explanation:
Cortex Analyst's 'Suggested Questions' feature operates in different modes. If a semantic model does not include a Verified Query Repository (VQR), Cortex Analyst uses Large Language Models (LLMs) to generate up to three suggested questions, which may not always be answerable. Therefore, option A is incorrect because it returns a blank list, and option D is incorrect because the LLM-generated questions are not guaranteed to be answerable. When a semantic model has a VQR defined, Cortex Analyst

suggests up to five questions from the VQR based on their similarity to the user's input, making option B correct. Furthermore, in Customizable Mode, if verified queries in the VQR are marked with
, Cortex Analyst will return all of these flagged questions, regardless of their quantity or similarity to user input, making option E correct and option C incorrect.

**NEW QUESTION # 178**
.....

Because there are free trial services provided by our GES-C01 preparation materials, by the free trial services you can get close contact with our products, learn about our GES-C01 real test, and know how to choice the different versions before you buy our products. On the other hand, using free trial downloading before purchasing, I can promise that you will have a good command of the function of our GES-C01 Test Prep. According to free trial downloading, you will know which version is more suitable for you.

**GES-C01 New APP Simulations**: https://www.lead2passexam.com/Snowflake/valid-GES-C01-exam-dumps.html

Snowflake GES-C01 Actual Test Pdf If we accidentally miss your question, please contact us again and we will keep in touch with you, Snowflake GES-C01 Actual Test Pdf Our staff will be online for 24 hours, It is a good tool for the candidates to learn more knowledge and to practice and improve their capability of dealing with all kinds of questions in real Snowflake GES-C01 exam, Our GES-C01 learning quiz can relieve you of the issue within limited time.

Ordering Graphic Elements, It will contain a detailed scope of work GES-C01 and related exceptions, If we accidentally miss your question, please contact us again and we will keep in touch with you.

## Newest GES-C01 Actual Test Pdf & Complete GES-C01 New APP Simulations & Free Download GES-C01 New Braindumps Pdf

Our staff will be online for 24 hours, It is a good tool for the candidates to learn more knowledge and to practice and improve their capability of dealing with all kinds of questions in real Snowflake GES-C01 Exam.

Our GES-C01 learning quiz can relieve you of the issue within limited time, It will let you close to your success, and into your dream paradise step by step.

- Top GES-C01 Questions □ Top GES-C01 Questions □ Detailed GES-C01 Answers □ Search for " GES-C01 " and obtain a free download on □ www.prep4sures.top □ □GES-C01 Real Dump
- Detailed GES-C01 Answers □ GES-C01 Latest Test Pdf □ GES-C01 Dump □ Open website ➥ www.pdfvce.com □ and search for ➤ GES-C01 □ for free download □Valid GES-C01 Test Duration
- Snowflake - GES-C01 - Marvelous SnowPro® Specialty: Gen AI Certification Exam Actual Test Pdf □ Search for ☀ GES-C01 □☀□ and download exam materials for free through ➡ www.prep4sures.top □□□ □GES-C01 Dumps Questions
- Top GES-C01 Questions □ GES-C01 Pdf Torrent □ GES-C01 Pdf Pass Leader □ Open { www.pdfvce.com } enter [ GES-C01 ] and obtain a free download □GES-C01 Real Torrent
- New GES-C01 Practice Materials □ GES-C01 Pdf Torrent □ GES-C01 Pdf Torrent □ Search for ▸ GES-C01 ◂ and download it for free on { www.torrentvce.com } website □Valid GES-C01 Test Sample
- GES-C01 Latest Test Pdf □ GES-C01 Standard Answers □ Official GES-C01 Practice Test □ The page for free download of 【 GES-C01 】 on ➤ www.pdfvce.com □ will open immediately □GES-C01 Real Torrent
- GES-C01 Dump □ New GES-C01 Practice Materials □ GES-C01 Dump □ Enter □ www.prep4away.com □ and search for 《 GES-C01 》 to download for free □Online GES-C01 Lab Simulation
- GES-C01 Study Guide - GES-C01 Free Download pdf - GES-C01 Latest Pdf Vce □ Easily obtain [ GES-C01 ] for free download through □ www.pdfvce.com □ ❤□GES-C01 Best Study Material
- GES-C01 Pdf Pass Leader □ GES-C01 Dump □ GES-C01 Latest Test Pdf □ Search for " GES-C01 " on ➤ www.practicevce.com □ immediately to obtain a free download □New GES-C01 Practice Materials
- GES-C01 Study Guide - GES-C01 Free Download pdf - GES-C01 Latest Pdf Vce □ Immediately open □ www.pdfvce.com □ and search for ➥ GES-C01 □ to obtain a free download □Reliable GES-C01 Exam Test
- The Best Accurate Trustable GES-C01 Actual Test Pdf Covers the Entire Syllabus of GES-C01 □ Easily obtain free download of □ GES-C01 □ by searching on ➽ www.troytecdumps.com □ □Official GES-C01 Practice Test
- www.stes.tyc.edu.tw, 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, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, best100courses.com, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, writeablog.net, 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, myportal.utt.edu.tt, Disposable vapes