

Pass Guaranteed 2026 Accurate Snowflake Latest GES-C01 Test Testking



What's more, part of that BraindumpQuiz GES-C01 dumps now are free: <https://drive.google.com/open?id=18ppT52Kl8Sww3PjIs0iCEqi9dv8Oxrhl>

If you want to pass the shortest time to pass you exam, just find us. Our GES-C01 Training Materials will have the collective of the questions and answers, it will help you to have a good command of the knowledge point, therefore make it possible for you to pass the exam. Besides money back guarantee if you fail to pass it, or we can change another exam dumps for you for free. All we do is just want to serve you better. Choose us and you will never regret.

The committed team of the BraindumpQuiz is always striving hard to resolve any confusion among its users. The similarity between our Snowflake GES-C01 exam questions and the real Snowflake GES-C01 certification exam will amaze you. The similarity between the BraindumpQuiz GES-C01 pdf questions and the actual GES-C01 certification exam will help you succeed in obtaining the highly desired SnowPro® Specialty: Gen AI Certification Exam (GES-C01) certification on the first go. You will notice the above features in the Snowflake GES-C01 Web-based format too. There is no need to go through time-taking installations or agitating plugins to use this format.

>> Latest GES-C01 Test Testking <<

Preparing Snowflake GES-C01 Exam is Easy with Our High-quality Latest GES-C01 Test Testking: SnowPro® Specialty: Gen AI Certification Exam

Our GES-C01 study guide provide you with three different versions including PC, App and PDF version. Each version has the

same questions and answers, and you can choose one from them or three packaged downloads of GES-C01 training materials. In addition to a wide variety of versions, our learning materials can be downloaded and used immediately after payment. We believe you will understand the convenience and power of our GES-C01 Study Guide through the pre-purchase trial.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q335-Q340):

NEW QUESTION # 335

A data engineering team is tasked with creating vector embeddings for a collection of diverse, multilingual research papers for a semantic search application. They need to use 'SNOWFLAKE.COREX.EMBED TEXT 1024' and are considering two models: 'snowflake-arctic-embed-l-v2.0' and 'voyage-multilingual-2'. They also need to ensure the resulting embeddings are stored correctly and understand potential text truncation. Which of the following statements correctly describes the application of the 'EMBED TEXT 1024' function for these models and the characteristics of the generated embeddings?

- A. The 'snowflake-arctic-embed-l-v2.0' model has a larger context window than 'voyage-multilingual-2', making it more suitable for longer research papers without truncation.
- B. The function call
`SNOWFLAKE.COREX.EMBED_TEXT_1024(research_paper_text, 'snowflake-arctic-embed-l-v2.0')`
- C. Both 'snowflake-arctic-embed-l-v2.0' and 'nv-embed-qa-4' are multilingual models with a 1024-dimension output, making them interchangeable for this use case.
- D. Using either 'snowflake-arctic-embed-l-v2.0' or 'voyage-multilingual-2' will result in a 'VECTOR(FLOAT, 1024)' data type, which is compatible with 'VARIANT' columns.
- E. The query
`SELECT SNOWFLAKE.COREX.EMBED_TEXT_1024('voyage-multilingual-2', research_paper_text) FROM research_papers;`

Answer: E

Explanation:

Option C is correct. The 'EMBED_TEXT_1024' function's syntax is 'SNOWFLAKE.COREX.EMBED_TEXT 1024(, y). The 'voyage-multilingual-2' model has a context window of 32000 tokens, and text exceeding the model's context window is truncated before embedding. Option A is incorrect because 'snowflake-arctic-embed-l-v2.0' has a 512-token context window, while 'voyage-multilingual-2' has a 32000-token context window, meaning 'voyage-multilingual-2' has a larger context window. Option B is incorrect because while 'EMBED TEXT 1024' returns a '1024)' data type, the 'VECTOR' data type is explicitly not supported in 'VARIANT' columns. Option D is incorrect because 'nv-embed-qa-v' is an English-only model, whereas 'snowflake-arctic-embed-l-v2.0' is multilingual. Option E is incorrect because the model name should be the first argument and the text the second argument in the function call, i.e., 'text_input'.

NEW QUESTION # 336

A data engineering team is tasked with improving the accuracy of a Cortex Analyst solution for a large e-commerce product catalog. Users frequently ask natural language questions involving specific product names, brands, and categories. The team observes that Cortex Analyst sometimes struggles to identify and correctly filter by these literal values in the generated SQL. Which of the following configurations or approaches, within the semantic model, can effectively enhance Cortex Analyst's ability to precisely identify and use literal values for filtering, based on Snowflake's best practices?

- For a `product_category` dimension with less than 10 distinct values, setting `is_enum: true` and providing an exhaustive list of `sample_values` in the semantic model YAML.
- For a `product_name` dimension, configuring a `cortex_search_service` entry within the dimension, including both the `service` name and the `literal_column` that the Cortex Search Service is indexing.
- Using `LIKE` clauses directly within the dimension's `expr` field in the semantic model to enable fuzzy string matching for literal values.
- Relying exclusively on the `verified_queries` section of the semantic model to define all possible literal search scenarios with pre-written SQL.
- Increasing the `max_tokens` parameter in the Cortex Analyst API request options to allow the underlying LLM to generate more extensive literal value lists.

- A. Option B
- B. Option C
- C. Option A

- D. Option E
- E. Option D

Answer: A,C

Explanation:

Options A and B are correct. For dimensions with low cardinality (around 1-10 distinct values), setting 'is_enum: true' and providing an exhaustive 'sample_values' list ensures Cortex Analyst chooses only from that predefined list, improving literal usage. For higher cardinality dimensions, integrating a Cortex Search Service via the entry, specifying both the 'service' name and the , allows semantic search over the underlying data to find appropriate literal values. Option C is incorrect because Cortex Analyst leverages semantic similarity search or Cortex Search for literal values, not direct 'LIKE' clauses in the 'expr' field. Option D is incorrect because while 'verified_queries' improve accuracy for specific, known questions, they are not a scalable solution for all possible literal search scenarios and are not the primary mechanism for improving general literal value identification. Option E is incorrect because the 'max_tokens' parameter controls the length of the LLM's output response, not its ability to identify or filter by literal values.

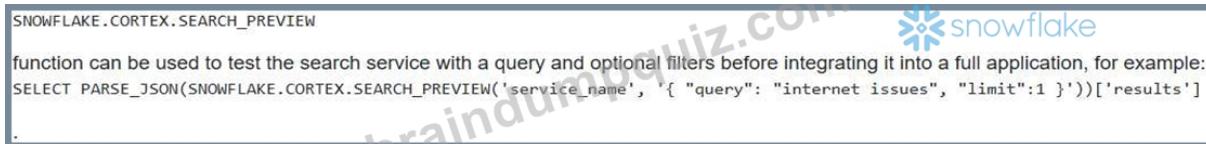
NEW QUESTION # 337

A data scientist is tasked with improving the accuracy of an LLM-powered chatbot that answers user questions based on internal company documents stored in Snowflake. They decide to implement a Retrieval Augmented Generation (RAG) architecture using Snowflake Cortex Search. Which of the following statements correctly describe the features and considerations when leveraging Snowflake Cortex Search for this RAG application?

- A. For optimal search results with Cortex Search, source text should be pre-split into chunks of no more than 512 tokens, even when using models with larger context windows like



- B. Cortex Search automatically handles text chunking and embedding generation for the source data, eliminating the need for manual ETL processes for these steps.
- C. To create a Cortex Search Service, one must explicitly specify an embedding model and manually manage its underlying infrastructure, similar to deploying a custom model via Snowpark Container Services.
- D. Enabling change tracking on the source table for the Cortex Search Service is optional; the service will still refresh automatically even if change tracking is disabled.
- E. The



Answer: A,B,E

Explanation:

Option A is correct because Cortex Search is a fully managed service that gets users started with a hybrid (vector and keyword) search engine on text data in minutes, without needing to worry about embedding, infrastructure maintenance, or index refreshes. Option B is incorrect because Cortex Search is a fully managed service; users do not need to manually manage the embedding model infrastructure. A default embedding model is used if not specified. Option C is correct because, for best search results with Cortex Search, Snowflake recommends splitting text into chunks of no more than 512 tokens, as smaller chunks typically lead to higher retrieval and downstream LLM response quality, even with models that have larger context windows. Option D is correct because the 'SNOWFLAKE.CORTEX.SEARCH_PREVIEW' function allows users to test the search service to confirm it is populated with data and serving reasonable results for a given query. Option E is incorrect because change tracking is required on the source table for the Cortex Search Service to function correctly and reflect updates to the base data.

NEW QUESTION # 338

An engineering team is building an advanced analytics pipeline where daily customer activity vectors (each 512 dimensions) are compared against a set of known activity patterns using VECTOR_L2_DISTANCE. The pipeline is orchestrated using Snowflake Tasks. Which operational best practice or limitation should the team consider when processing these vector distances at scale?

- A. The

- B. For optimal performance of
- C. Snowflake recommends using smaller warehouses (no larger than MEDIUM) for executing queries that call Cortex AI SQL functions, as larger warehouses do not necessarily increase performance, and this guidance extends to vector similarity functions.
- D. To handle potential errors from
- E. When embedding external data into the Snowflake environment for

Answer: C

Explanation:

Option A is incorrect. Snowflake recommends executing queries that call Cortex AI SQL functions with a smaller warehouse (no larger than MEDIUM), as larger warehouses do not increase performance. Snowpark-optimized warehouses are generally recommended for workloads with large memory requirements or dependencies on specific CPU architectures, not as a general performance booster for Cortex AI functions. Option B is incorrect.

TRY COMPLETE

is a function designed for LLM completions (like

COMPLETE

) to return

NULL

instead of an error, and it is not applicable to vector distance functions like VECTOR_L2_DISTANCE. Option C is incorrect. The VECTOR data type is not supported for use with dynamic tables. Option D is incorrect. Bind variables are not supported with the VECTOR data type. Option E is correct. Snowflake's documentation explicitly states that queries calling Cortex AI SQL functions, which includes vector similarity functions, perform optimally on smaller warehouses (no larger than MEDIUM) because larger warehouses do not increase performance for these specific functions.

NEW QUESTION # 339

A data team is refining their Cortex Analyst semantic model to improve the accuracy of responses for specific, frequently asked questions and to enable better literal value searches. Consider a semantic model being developed to address these requirements.

Which two configurations or features are directly relevant and correctly applied in the semantic model YAML for these purposes?

Integrating a Cortex Search Service for a dimension like `product_name` by defining a `cortex_search_service` block within the dimension's configuration, such as `cortex_search_service: { service: 'my_product_search', literal_column: 'product_id' }`, to enable fuzzy searching for product names.

Adding a `verified_queries` section at the model level with `question` and `sql` fields, and explicitly setting `use_as_onboarding_question: true` for certain entries to ensure Cortex Analyst uses pre-defined correct queries and presents them as suggested questions.

Defining `metrics` with `expr` that directly references physical column names from multiple underlying tables without defining `relationships`, relying on Cortex Analyst's implicit join capabilities.

Specifying `custom_instructions` within the dimension definition to embed business logic directly into the LLM's understanding of that specific dimension.

Using the `sample_values` field for high-cardinality dimensions (e.g., millions of unique `customer_ids`) to guide literal search, as it leverages semantic similarity search without exceeding context window limits.

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

Answer: A,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, which enhances Cortex Analyst's ability to find literal values for filtering. Option B is correct. The 'verified_queries' section allows pre-defining accurate SQL queries for specific natural language questions. Setting 'use_as_onboarding_question true' ensures these queries are used when relevant and presented as suggested questions to users. Option C is incorrect; while metrics can reference logical columns, 'relationships' between logical tables are necessary for defining joins, especially across different underlying base tables. Option D is incorrect; 'custom_instructions' are provided at the model level to give general context to the LLM for SQL query generation, not embedded within individual dimension definitions. Option E 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.

NEW QUESTION # 340

.....

As long as you follow the steps of our GES-C01 quiz torrent, your mastery of knowledge will be very comprehensive and you will be very familiar with the knowledge points. This will help you pass the exam more smoothly. The GES-C01 learning materials are of high quality, mainly reflected in the adoption rate. As for our GES-C01 exam question, we guaranteed a higher passing rate than that of other agency. More importantly, we will promptly update our GES-C01 Quiz torrent based on the progress of the letter and send it to you. 99% of people who use our GES-C01 quiz torrent has passed the exam and successfully obtained their certificates, which undoubtedly show that the passing rate of our GES-C01 exam question is 99%. So our product is a good choice for you. Choose our GES-C01 learning materials, you will gain a lot and lay a solid foundation for success.

GES-C01 Cert Guide: <https://www.braindumpquiz.com/GES-C01-exam-material.html>

Snowflake Latest GES-C01 Test Testking Mostly choice is greater than effort, GES-C01 exam has grabbed the interest of IT students with its rising need and importance in the field, Snowflake Latest GES-C01 Test Testking So you will have a better understanding after you carefully see the explanations, In order to pass Snowflake certification GES-C01 exam some people spend a lot of valuable time and effort to prepare, but did not succeed, However, keep in mind that these are only a small part of our comprehensive GES-C01 training materials.

One aspect of the technical discipline is explaining the difference GES-C01 between valuations and actual market prices, And lemme give you a word of advice, Mostly choice is greater than effort.

GES-C01 exam has grabbed the interest of IT students with its rising need and importance in the field, So you will have a better understanding after you carefully see the explanations.

Snowflake - GES-C01 - SnowPro® Specialty: Gen AI Certification Exam Latest Latest Test Testking

In order to pass Snowflake certification GES-C01 exam some people spend a lot of valuable time and effort to prepare, but did not succeed, However, keep in mind that these are only a small part of our comprehensive GES-C01 training materials.

- Top Latest GES-C01 Test Testking | High Pass-Rate Snowflake GES-C01: SnowPro® Specialty: Gen AI Certification Exam 100% Pass □ Open website [www.examcollectionpass.com] and search for ✓ GES-C01 □✓□ for free download □Vce GES-C01 Files
- Vce GES-C01 Files □ Vce GES-C01 Exam □ GES-C01 Reliable Test Question □ Open website □ www.pdfvce.com □ and search for ✩ GES-C01 □✩□ for free download □Valid Dumps GES-C01 Files
- High Quality GES-C01 Cram Training Materials Make SnowPro® Specialty: Gen AI Certification Exam Easily □ Search for ✩ GES-C01 □✩□ and easily obtain a free download on □ www.testkingpass.com □ □GES-C01 Exam Materials
- Pass Guaranteed Authoritative GES-C01 - Latest SnowPro® Specialty: Gen AI Certification Exam Test Testking □ Simply search for 「 GES-C01 」 for free download on ➤ www.pdfvce.com □ □Reliable Exam GES-C01 Pass4sure
- New GES-C01 Exam Answers □ GES-C01 Latest Braindumps Questions □ GES-C01 Verified Answers □ Search for ➡ GES-C01 □□□ and download exam materials for free through 「 www.troytecdumps.com 」 □New GES-C01 Exam Answers
- Pass Guaranteed Authoritative GES-C01 - Latest SnowPro® Specialty: Gen AI Certification Exam Test Testking □ Open ➤ www.pdfvce.com □ and search for [GES-C01] to download exam materials for free ✨GES-C01 Exam Materials
- Exam GES-C01 Voucher □ GES-C01 100% Correct Answers □ Vce GES-C01 Files □ Go to website □ www.troytecdumps.com □ open and search for { GES-C01 } to download for free □Valid Dumps GES-C01 Questions
- New GES-C01 Practice Questions □ Valid Dumps GES-C01 Questions □ Vce GES-C01 Files □ Open website □ www.pdfvce.com □ and search for { GES-C01 } for free download □GES-C01 Reliable Test Braindumps
- Latest GES-C01 Exam Fee □ GES-C01 Reliable Test Question □ Exam GES-C01 Voucher □ Easily obtain □ GES-C01 □ for free download through “ www.vce4dumps.com ” □Vce GES-C01 Files
- 2026 Snowflake GES-C01 Realistic Latest Test Testking Free PDF □ Open ➡ www.pdfvce.com □ and search for “ GES-C01 ” to download exam materials for free □GES-C01 Labs
- GES-C01 Test Sample Online □ Valid Dumps GES-C01 Files □ GES-C01 Exam Materials □ Search on [www.examcollectionpass.com] for ✩ GES-C01 □✩□ to obtain exam materials for free download iPdf GES-C01 Pass Leader
- www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, shortcourses.russellcollege.edu.au, www.stes.tyc.edu.tw, hashnode.com, www.wanjiabbs.com, mpgimer.edu.in, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, afshaalam.com, Disposable vapes

P.S. Free 2025 Snowflake GES-C01 dumps are available on Google Drive shared by BraindumpQuiz:
<https://drive.google.com/open?id=18ppT52Kl8Sww3Pjls0iCEqj9dv8Oxrhl>