


Pass Guaranteed Quiz Snowflake - High Hit-Rate GES-C01 Valid Dumps Book



Are you still feeling uncomfortable about giving up a lot of time to entertain, work or accompany your family and friends in preparation for the exam? Using GES-C01 quiz torrent, you can spend less time and effort reviewing and preparing, which will help you save a lot of time and energy. Whether you are a worker or student, you will save much time to do something whatever you want. It only needs 5-10 minutes after you pay for our GES-C01 learn torrent that you can learn it to prepare for your exam. Actually, if you can guarantee that your effective learning time with GES-C01 test preps are up to 20-30 hours, you can pass the exam.

To attain all these you just need to enroll in the Snowflake GES-C01 certification exam and put in all your efforts and prepare well to crack the Snowflake GES-C01 exam easily. For the perfect and instant Snowflake GES-C01 preparation, you can get help from Snowflake GES-C01 Questions. The Prep4cram GES-C01 exam questions are real and will entirely assist you in GES-C01 exam preparation and you can easily pass the final Snowflake GES-C01 certification exam.

>> GES-C01 Valid Dumps Book <<

New GES-C01 Exam Notes & GES-C01 Passing Score

The Prep4cram is one of the leading platforms that has been offering real and valid SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam practice test questions. These SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam questions are designed and verified by Snowflake GES-C01 subject matter experts. They work closely together and put all their expertise to check the Snowflake GES-C01 exam questions one by one.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q90-Q95):

NEW QUESTION # 90

A financial institution is deploying a sentiment analysis application that uses Snowflake Cortex 'SENTIMENT' and 'COMPLETE' functions, with different LLMs, for processing customer feedback. They are using AI Observability (Public Preview) to compare the cost- efficiency of using 'mistral-7b' versus 'claude-3-5-sonnet' as LLM judges for evaluation metrics, and also tracking the overall cost of their AI Observability usage. Which statements accurately reflect the cost implications and monitoring tools for this scenario?

- ☐ AI Observability incurs charges for LLM judges (e.g., 'mistral-7b', 'claude-3-5-sonnet') invoked via 'COMPLETE (SNOWFLAKE.CORTEX)' calls to compute evaluation metrics, and these charges are based on 'tokens processed'.
- ☐ The 'CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY' view should be used to monitor the credit consumption of the LLM judges specifically, as it tracks all Cortex function calls.
- ☐ Comparing models with vastly different context windows (e.g., 'mistral-7b' at 32k tokens vs. 'claude-3-5-sonnet' at 200k tokens) using AI Observability will not impact the billed 'tokens processed' if the actual prompt sizes are small and similar.
- ☐ In addition to LLM judge costs, warehouse charges are incurred for tasks managing evaluation runs and for queries used to compute evaluation metrics within AI Observability.
- ☐ The 'METERING_DAILY_HISTORY' view, filtering by 'SERVICE_TYPE ILIKE '%ai_services%', can provide an overview of daily credit usage for all AI services, including AI Observability LLM judge activity and associated warehouse costs.

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

Answer: A,B,D

Explanation:

Option A is correct because AI Observability utilizes LLM judges (such as 'mistral-7b' or 'claude-3-5-sonnet') through 'COMPLETE (SNOWFLAKE.CORTEX)' function calls to compute evaluation metrics, and these calls incur charges based on the 'tokens processed'. Option D is correct as, beyond LLM judge costs, AI Observability also incurs warehouse charges for managing evaluation runs and for queries that compute evaluation metrics. Option E is correct because the view, with a filter for 'SERVICE _ TYPE ILIKE ', provides a comprehensive daily credit usage report for all AI services, which would include AI Observability's components. Option B is incorrect; the view is specifically for Document AI processing functions like 'PREDICT and 'AI_EXTRACT', not for general LLM judge usage in AI Observability. The view is more appropriate for tracking individual Cortex function calls. Option C is incorrect because while prompt sizes might be similar, the pricing for different LLMs (e.g., 'mistral-7b' at 0.12 credits per million tokens vs. 'claude-3-5-sonnet' at 2.55 credits per million tokens for AI Complete) will still result in different billed amounts due to varying per-token costs, even if the number of tokens is the same.

NEW QUESTION # 91

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. 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.
- B. 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_querv'.
- 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. 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.
- E. Custom instructions are primarily used to define new logical tables or dimensions within the semantic model, effectively extending the data model at runtime.

Answer: C,D

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_querv' 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 # 92

A global analytics firm is developing a Retrieval Augmented Generation (RAG) system in Snowflake to answer customer queries across a large repository of technical documentation, which includes documents in English, German, and Spanish. They are looking to use a Snowflake Cortex embedding model to convert document chunks into vector embeddings for their Cortex Search Service. Which of the following considerations are critical when selecting an appropriate embedding model to optimize for both query relevance and cost-efficiency for their multilingual RAG application? (Select all that apply)

`snowflake-arctic-embed-m-v1.5`

Model is suitable as it offers the fastest indexing times and is highly cost-effective for all supported languages.

☐ For multilingual support,

`snowflake-arctic-embed-l-v2.0`

or

`voyage-multilingual-2`

are appropriate choices, with

`snowflake-arctic-embed-l-v2.0-8k`

Offering an extended context window (8192 tokens) for longer document chunks without an increased cost per million tokens.

☐ Smaller chunk sizes (e.g., no more than 512 tokens) are generally recommended for higher retrieval precision and improved downstream LLM response quality, even with embedding models that support larger context windows.

☐

`EMBED_TEXT_1024`

Functions are always more cost-efficient than

`EMBED_TEXT_768`

Functions due to their higher output dimensions, making them universally preferred for RAG applications.

☐ The context window of the chosen embedding model directly dictates the maximum number of pages a document can have for processing in Cortex Search

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

Answer: B,D

Explanation:

Option B is correct because both

`snowflake-arctic-embed-l-v2.0`

and

`voyage-multilingual-2`

are explicitly listed as multilingual embedding models. The

`snowflake-arctic-embed-l-v2.0-8k`

model provides an increased context window of 8000 tokens while maintaining the same cost per million tokens (0.05 credits) as the 512-token version of

Option C is correct because Snowflake recommends splitting text into chunks of no more than 512 tokens for best search results with Cortex Search, as research shows this typically leads to higher retrieval precision and improved downstream LLM response quality, even when using longer-context embedding models. Option A is incorrect because

`snowflake-arctic-embed-m-v1.5`

is an English-only embedding model, which does not meet the requirement for multilingual documentation. Option D is incorrect; the cost per million tokens for `EMBED_TEXT_1024` models (e.g., 0.05-0.07 credits) is not inherently more cost-efficient than `EMBED_TEXT_768` models (e.g., 0.03 credits), and cost-efficiency depends on the specific model and use case, not just output

dimensions. Option E is incorrect; the context window of an embedding model refers to the maximum length of a text input (chunk) it can process. The maximum pages a document can have (e.g., 300 pages for Document AI) is a separate document requirement, not directly determined by the embedding model's context window.

NEW QUESTION # 93

A Document AI administrator is investigating why a new 'finance_automation_role' cannot create Document AI model builds in 'finance_analytics.quarterly_reports_schema'. The role has the 'SNOWFLAKE.DOCUMENT_INTELLIGENCE_CREATOR' database role, and 'USAGE' on the database and schema. The warehouse is active and accessible. The error message received is: 'Unable to create a build on the specified database and schema. Please check the documentation to learn more.'

Which 'combination of missing schema-level privileges' is explicitly cited in the documentation as a direct cause for this error, assuming a unique model build name?

- A.

Only 'CREATE STAGE' is missing on 'finance_analytics.quarterly_reports_schema'.

- B.

- C.

Only 'CREATE SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE' is missing on 'finance_analytics.quarterly_reports_schema'.

- D.

Only 'CREATE MODEL' is missing on 'finance_analytics.quarterly_reports_schema'.

- E.

Both 'CREATE FUNCTION' and 'CREATE PROCEDURE' are missing on 'finance_analytics.quarterly_reports_schema'.

Answer: B

Explanation:

The troubleshooting documentation for the error message 'Unable to create a build on the specified database and schema' explicitly lists two primary causes related to missing schema-level privileges: 'The 'CREATE SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE' privilege is not granted to your role' and 'Your role has not been granted the 'CREATE MODEL' privilege on the schema that uses the model'. Both of these privileges are required on the schema to prepare a DocumentAI model build. Therefore, the combination of both missing would directly lead to this specific error. Options C and D are individual components of the correct answer, but the question asks for the 'combination of missing schema-level privileges' as cited in the documentation.


NEW QUESTION # 94

A data application developer is building a Streamlit chat application within Snowflake. This application uses a RAG pattern to answer user questions about a knowledge base, leveraging a Cortex Search Service for retrieval and an LLM for generating responses. The developer wants to ensure responses are relevant, concise, and structured. Which of the following practices are crucial when integrating Cortex Search with Snowflake Cortex LLM functions like AI_COMPLETE for this RAG chatbot?


- A. Using the



- B. For performance and cost optimization, it is always recommended to query Cortex Search and the LLM function within a single
- C. The



- D. To maintain conversational context in a multi-turn chat, the developer should pass all previous user prompts and model responses in the



- E. The retrieved context from Cortex Search should be directly concatenated with the user's prompt as input to the

Answer: A,D

Explanation:

Option A is incorrect. The user's query is typically embedded (e.g., using to perform a similarity search against the Cortex Search Service. The 'retrieved documents' (context) are then passed to the 'AI COMPLETE function, not the embedding function itself. Option B is correct because to provide a stateful, conversational experience, all previous user prompts and model responses should be passed in the array to the 'COMPLETE or function. Option C is incorrect. While concatenation is a method, for better accuracy and control, the retrieved context should be integrated into a well-engineered prompt, often using tags or specific instructions, rather than just raw concatenation, to guide the LLM's response. Option D is correct because 'AI_COMPLETE Structured Outputs' allows you to supply a JSON schema that completion responses must follow, reducing the need for post-processing and enabling seamless integration with systems requiring deterministic responses. Option E is incorrect. While keeping processing within Snowflake is good for data governance, complex RAG pipelines often involve multiple distinct steps (query embedding, search, retrieval, LLM completion) that may benefit from a staged approach rather than a single monolithic SQL statement. The optimal approach depends on the specific complexity and performance requirements, and a single 'SELECT for the 'entire' RAG flow might not always be the most efficient or practical solution.

NEW QUESTION # 95

.....

Maybe you often come up with great new ideas from daydream, but you can not do anything. Do you have some trouble passing Snowflake GES-C01 Exam? Turn on your computer, click Prep4cram. Then, you will find the dumps torrent you need. After you purchase our products, we provide free updates for a year. 100% guarantee to get the certification.

New GES-C01 Exam Notes: https://www.prep4cram.com/GES-C01_exam-questions.html

Snowflake GES-C01 Valid Dumps Book We are truly a dream team, we believe in talent and professionalism, and, what's important, we're always hiring, Snowflake GES-C01 Valid Dumps Book They will offer as the smartest way to succeed in limited time, If you buy the GES-C01 latest questions of our company, you will have the right to enjoy all the GES-C01 certification training materials from our company, Firstly, all types of questions are included in our GES-C01 training material that wide coverage will be helpful for you to pass exam

Select the Bevel and Emboss Layer Effect on the ellipse shape layer, and delete GES-C01 it this way, Working as root, We are truly a dream team, we believe in talent and professionalism, and, what's important, we're always hiring!

Free PDF Quiz 2025 Reliable Snowflake GES-C01 Valid Dumps Book

They will offer as the smartest way to succeed in limited time, If you buy the GES-C01 Latest Questions of our company, you will have the right to enjoy all the GES-C01 certification training materials from our company.

Firstly, all types of questions are included in our GES-C01 training material that wide coverage will be helpful for you to pass exam, You can take multiple GES-C01 SnowPro® Specialty: Gen AI Certification Exam practice exam attempts and identify and overcome your mistakes.

- Authoritative GES-C01 Valid Dumps Book - Find Shortcut to Pass GES-C01 Exam ☐ Open website [www.prep4pass.com] and search for ☐ GES-C01 ☐ for free download ☐ GES-C01 Clearer Explanation
- GES-C01 Valid Exam Vce ➡ Certification GES-C01 Test Questions ☐ GES-C01 Clearer Explanation ☐ Immediately open “ www.pdfvce.com ” and search for ☐ GES-C01 ☐ to obtain a free download ☐ GES-C01 Interactive Practice Exam
- Free GES-C01 Practice ☺ GES-C01 Valid Dumps Ppt ☐ GES-C01 Reliable Test Simulator ☐ Search for ➡ GES-C01

- and download it for free on □ www.exam4pdf.com □ website □ Valid GES-C01 Exam Experience
- Latest SnowPro® Specialty: Gen AI Certification Exam dumps pdf - GES-C01 examsboost review □ Search for [GES-C01] and download it for free immediately on □ www.pdfvce.com □ □GES-C01 Interactive Practice Exam
- The Best Accurate GES-C01 Valid Dumps Book – Find Shortcut to Pass GES-C01 Exam □ The page for free download of ➡ GES-C01 □ on 【 www.torrentvce.com 】 will open immediately □GES-C01 Interactive Practice Exam
- Newly Released Snowflake GES-C01 Dumps in Three Formats [2025] □ Open website ⇒ www.pdfvce.com ⇐ and search for ➡ GES-C01 □ for free download □GES-C01 Reliable Test Topics
- Pass Guaranteed Quiz Reliable Snowflake - GES-C01 - SnowPro® Specialty: Gen AI Certification Exam Valid Dumps Book □ Immediately open □ www.real4dumps.com □ and search for ⇒ GES-C01 ⇐ to obtain a free download □Exam GES-C01 Tests
- Valid GES-C01 Exam Experience □ GES-C01 Interactive Practice Exam □ Best GES-C01 Study Material ☞ Open ► www.pdfvce.com ◀ and search for ➡ GES-C01 □ to download exam materials for free □GES-C01 Interactive Practice Exam
- Certification GES-C01 Test Questions □ GES-C01 Valid Dumps Ppt □ GES-C01 Clearer Explanation □ [www.passcollection.com] is best website to obtain { GES-C01 } for free download ↘ Training GES-C01 Solutions
- Pass Guaranteed Quiz Reliable Snowflake - GES-C01 - SnowPro® Specialty: Gen AI Certification Exam Valid Dumps Book □ Search for “GES-C01 ” and download it for free immediately on 《 www.pdfvce.com 》 □GES-C01 Valid Dumps Ppt
- Express Greetings to a Useful Future by Getting Snowflake GES-C01 Dumps □ Search for □ GES-C01 □ on (www.lead1pass.com) immediately to obtain a free download □GES-C01 Guaranteed Success
- 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, bbs.synwit.cn, study.stcs.edu.np, tutor.foodshops.ng, digitalrepublix.com, www.jobskillstraining.org, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, fluencyfocus.in, Disposable vapes