

Realistic Top GES-C01 Questions - Win Your Snowflake Certificate with Top Score



Snowflake GES-C01

SnowPro® Specialty: Gen AI Certification Exam

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/ges-c01>

BTW, DOWNLOAD part of Rea4Prep GES-C01 dumps from Cloud Storage: <https://drive.google.com/open?id=1UTs-622hJo0tC9g8AumwOkKz1LaknacG>

The Snowflake GES-C01 pdf questions learning material provided to the customers from Rea4Prep is in three different formats. The first format is PDF format which is printable and portable. It means it can be accessed from tablets, laptops, and smartphones to prepare for the SnowPro® Specialty: Gen AI Certification Exam exam. The Snowflake GES-C01 Pdf Format can be used offline, and candidates can even prepare for it in the classroom or library by printing questions or on their smart devices.

On Rea4Prep website you can free download part of the exam questions and answers about Snowflake Certification GES-C01 Exam to quiz our reliability. Rea4Prep's products can 100% put you onto a success away, then the pinnacle of IT is a step closer to you.

>> Top GES-C01 Questions <<

Latest GES-C01 Test Questions & Free GES-C01 Practice

Providing our customers with up to 1 year of free Snowflake GES-C01 questions updates is also our offer. These Snowflake GES-C01 free dumps updates will help you prepare according to the latest GES-C01 test syllabus in case of changes. 24/7 customer support is available at Rea4Prep to assist users of the GES-C01 Exam Questions through the journey. Above all, Rea4Prep also offers a full refund guarantee (terms and conditions apply) to our customers. Don't miss these amazing offers. Download SnowPro® Specialty: Gen AI Certification Exam (GES-C01) actual exam Dumps today!

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q111-Q116):

NEW QUESTION # 111

A development team is constructing a Gen AI application using Snowflake Cortex LLM functions, particularly for conversational and text generation tasks. They are concerned about potential high costs due to token consumption. Which of the following strategies would most effectively help minimize token usage and optimize costs when working with these Cortex LLM functions?

- For multi-turn conversational experiences using `SNOWFLAKE.CORTEX.COMPLETE`, only send the most recent user prompt in each API call, as the model automatically retains previous context.
- When employing `AI.COMPLETE` for structured output tasks, providing concise and highly descriptive explanations for each field within the JSON schema will reduce the input tokens required for the LLM to understand and adhere to the schema accurately.
- Utilize the `COUNT_TOKENS (SNOWFLAKE.CORTEX)` helper function to pre-validate the prompt length against the model's context window, thereby preventing truncation errors and subsequent re-runs.
- To encourage more succinct LLM responses and reduce `completion_tokens`, configure the `temperature` option to a higher value (e.g., 0.7) in `COMPLETE` function calls.
- In multi-turn conversations within Cortex Analyst, integrate a dedicated LLM summarization agent to rephrase follow-up questions, which reduces the total conversational history passed as context to the main LLM.

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

Answer: A,B,D

Explanation:

Option B is correct because while schema validation itself doesn't incur extra cost, a large or complex schema can increase token consumption. Providing precise and concise descriptions for schema fields helps the LLM understand and adhere to the desired format more efficiently, potentially reducing the overall tokens consumed for accurate responses. Option C is correct as the 'COUNT_TOKENS' function allows developers to determine the token count of an input prompt for a specific model, enabling them to pre-emptively avoid exceeding the model's context window, thus preventing errors and wasted compute from re-runs. Option E is correct because for multi-turn conversations in Cortex Analyst, a summarization agent is specifically used to rephrase follow-up questions by incorporating previous context, without passing the entire, potentially long, conversation history. This significantly reduces the 'prompt_tokens' sent to the main LLM for each turn and optimizes inference times. Option A is incorrect because 'COMPLETE' (and 'TRY_COMPLETE') functions are stateless; to maintain conversational context, all previous user prompts and model responses must be included in the array, which increases token count proportionally. Simply sending the latest prompt would lose context. Option D is incorrect as setting a higher 'temperature' value (e.g., 0.7) increases the 'randomness and diversity' of the LLM's output, not necessarily its conciseness for cost optimization. For the most consistent (and often direct) results, a 'temperature' of 0 is recommended.

NEW QUESTION # 112

A data pipeline processes thousands of scanned legal contracts daily using a Document AI model via the 'PREDICT' function. The process involves fetching presigned URLs for documents from an internal stage using 'GET PRESIGNED URL'. Recently, the pipeline has started failing intermittently, returning the error:

The data engineering team confirms network connectivity and privileges are stable. Which of the following is the most likely cause of this issue?

- A. The documents contain handwritten annotations, which Document AI struggles to process effectively.
- B. The virtual warehouse assigned to the Document AI pipeline is undersized (e.g., X-Small), leading to slow processing that indirectly causes URL expiration.
- C. The internal stage is not configured with
- D. The Document AI model build has not been published to the production environment.
- E. The

Answer: E

Explanation:

The error message 'Received HTTP 403 response for presigned URL. URL may be expired.' is a direct indication that the presigned URL, generated by the function, has exceeded its default expiration time before the Document AI model could complete processing the document. The GET PRESIGNED URL' function has a default expiration time of 60 minutes. If the query processing time for a batch of documents, or any delay in the pipeline, exceeds this window, the URLs become invalid. Option A would typically result in errors related to file format or mime type if not using 'SNOWFLAKE SSE' for an internal stage. Option B would prevent the 'PREDICT function from executing successfully, likely with a different error message indicating the model build is not ready or accessible. Option D describes a potential performance bottleneck, but the immediate cause of a '403 URL expired' error is the time limit, regardless of underlying warehouse efficiency. Option E would affect the accuracy of extraction, not cause a presigned URL access error.

NEW QUESTION # 113

An enterprise is deploying a new RAG application using Snowflake Cortex Search on a large dataset of customer support tickets. The operations team is concerned about managing compute costs and ensuring efficient index refreshes for the Cortex Search Service, which needs to be updated hourly. Which of the following considerations and configurations are relevant for optimizing cost and performance of the Cortex Search Service in this scenario?

- **A. For embedding text, selecting a model like snowflake-arctic-embed-m-v1.5 (0.03 credits/million tokens) over voyage-multilingual-2 (0.07 credits/million tokens) could significantly reduce EMBED_TEXT_TOKENS**
- B. The primary cost driver for Cortex Search is the number of search queries executed against the service, with the volume of indexed data (GB/month) having a minimal impact on overall billing.
- **C. For optimal performance and cost efficiency, Snowflake recommends using a dedicated warehouse of size no larger than MEDIUM for each Cortex Search Service.**
- D. The **CREATE CORTX SEARCH SERVICE**
- **E. CHANGE_TRACKING**

Answer: A,C,D,E

Explanation:


Option A is correct because a Cortex Search Service requires a virtual warehouse to refresh the service, which runs queries against base objects when they are initialized and refreshed, incurring compute costs. Option B is correct because the cost of embedding models varies. For example, 'snowflake-arctic-embed-m-v1.5 costs 0.03 credits per million tokens, while 'voyage-multilingual-2 costs 0.07 credits per million tokens. Choosing a more cost-effective model like 'snowflake-arctic-embed-m-v1 for English-only data can reduce token costs. Option C is correct because Snowflake recommends using a dedicated warehouse of size no larger than MEDIUM for each Cortex Search Service to achieve optimal performance. Option D is correct because change tracking is required for the Cortex Search Service to be able to detect and process updates to the base table, enabling incremental refreshes that are more efficient than full re-indexing. Option E is incorrect because Cortex Search Services incur costs based on virtual warehouse compute for refreshes, 'EMBED TEXT TOKENS' cost per input token, and a charge of 6.3 Credits per GB/mo of indexed data. The volume of indexed data has a significant impact, not minimal.

NEW QUESTION # 114

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 # 115

A data analyst needs to use SNOWFLAKE.CORTEX.EXTRACT_ANSWER to streamline information retrieval from various contract documents. They are new to Cortex functions and want to understand access requirements and optimal usage. Which of the following statements about using SNOWFLAKE.CORTEX.EXTRACT_ANSWER are correct?

- A. EXTRACT_ANSWER is the most current and recommended function for all text extraction tasks, offering multi-label and image extraction capabilities.
- B. It is generally recommended to process multiple documents by applying EXTRACT_ANSWER to a table column containing the document texts, allowing for efficient batch processing.
- C. If EXTRACT_ANSWER encounters an unresolvable issue during processing, it returns NULL instead of an error, similar to TRY_COMPLETE.
- D. For optimal accuracy, the source_document input should be in plain English, and the question should be specific, asking for a single value.
- E. The analyst's role must be granted the SNOWFLAKE.CORTEX_USER database role to execute EXTRACT_ANSWER functions.

Answer: B,D,E

Explanation:

Option A is correct because users must use a role that has been granted the 'SNOWFLAKE.CORTEX_USER' database role to access EXTRACT_ANSWER and other Cortex AI functions. Option B is correct because for optimal performance and accurate responses, it is recommended to use plain English text for input and categories, and questions should be specific and ask for a single value. Option C is incorrect; EXTRACT_ANSWER would typically raise an error if an operation cannot be performed. The 'TRY_COMPLETE' function is specifically designed to return 'NULL' instead of an error in such cases. Option D is incorrect; 'EXTRACT_ANSWER' is an older version of this function, and 'AI_EXTRACT' is the latest version, which supports additional capabilities like image and multi-language extraction. Option E is correct; EXTRACT_ANSWER can be called on a table column, enabling efficient batch processing of multiple documents or text entries. This is a common pattern for integrating Cortex functions into data pipelines.

NEW QUESTION # 116

.....

It's not easy for most people to get the GES-C01 guide torrent, but I believe that you can easily and efficiently obtain qualification GES-C01 certificates as long as you choose our products. After you choose our study materials, you can master the examination point from the GES-C01 Guide question. Then, you will have enough confidence to pass your exam. As for the safe environment and effective product, why don't you have a try for our GES-C01 question torrent, never let you down!

Latest GES-C01 Test Questions: <https://www.real4prep.com/GES-C01-exam.html>

As long as you have questions, you can send e-mail to us, we have online staff responsible for ensuring 24-hour service to help you solve all the problems about our GES-C01 test prep, Fortunately, our website can offer you the most comprehensive GES-C01 dumps pdf to help you pass the GES-C01 valid test quickly, The training material provided by the Real4Prep is trusted by more 90 DAYS FREE UPDATES FOR GES-C01 EXAM.

That is, even though a process may be unpredictable on its own, it can still be controlled with feedback, GES-C01 pdf training contains the comprehensive knowledge of the actual test.

As long as you have questions, you can send e-mail to us, we have online staff responsible for ensuring 24-hour service to help you solve all the problems about our GES-C01 Test Prep.

Prominent Features of Snowflake GES-C01 Practice Test Questions

Fortunately, our website can offer you the most comprehensive GES-C01 dumps pdf to help you pass the GES-C01 valid test quickly, The training material provided by the Real4Prep is trusted by more 90 DAYS FREE UPDATES FOR GES-C01 EXAM.

Our free demo will show you the actual GES-C01 certification exam, If for any reason you do not pass your exam, PassLeader.com will provide you with a full refund or another GES-C01 exam of your choice absolutely free within 30 days from the date of purchase.

- www.dumpsmaterials.com Snowflake GES-C01 Exam Questions Formats www.dumpsmaterials.com is best website to obtain ➔ GES-C01 for free download GES-C01 Exam Labs
- GES-C01 Test Prep GES-C01 Exams Collection Latest GES-C01 Version Download GES-C01 for free by simply searching on ▶ www.pdfvce.com ◀ GES-C01 Exam Labs
- Hot Top GES-C01 Questions | Efficient Latest GES-C01 Test Questions: SnowPro® Specialty: Gen AI Certification Exam 100% Pass Open ➔ www.dumpsmaterials.com and search for GES-C01 to download exam materials for free Reliable GES-C01 Exam Materials
- Reliable GES-C01 Test Price GES-C01 Actual Questions Examcollection GES-C01 Vce Immediately open ▷ www.pdfvce.com ◁ and search for ➔ GES-C01 to obtain a free download GES-C01 Exams Collection
- Snowflake GES-C01 Desktop Practice Exam Questions Software Simply search for ⇒ GES-C01 ⇐ for free download on « www.pdfdumps.com » GES-C01 Customizable Exam Mode
- Latest GES-C01 Version Examcollection GES-C01 Dumps Torrent Examcollection GES-C01 Vce Download GES-C01 for free by simply entering { www.pdfvce.com } website GES-C01 Test Online
- Pass Guaranteed Snowflake - GES-C01 - SnowPro® Specialty: Gen AI Certification Exam –The Best Top Questions Search for ⇒ GES-C01 ⇐ and easily obtain a free download on « www.vceengine.com » GES-C01 Test Prep
- Free PDF 2026 GES-C01: High Hit-Rate Top SnowPro® Specialty: Gen AI Certification Exam Questions ➔ www.pdfvce.com is best website to obtain ▷ GES-C01 ◁ for free download GES-C01 Customizable Exam Mode
- Free PDF 2026 GES-C01: High Hit-Rate Top SnowPro® Specialty: Gen AI Certification Exam Questions Open website (www.prepawayexam.com) and search for « GES-C01 » for free download ↔GES-C01 Actual Questions
- GES-C01 Exam Labs GES-C01 Test Preparation GES-C01 Reliable Test Sims Search on (www.pdfvce.com) for ▷ GES-C01 ◁ to obtain exam materials for free download GES-C01 Actualtest
- GES-C01 Actual Questions Examcollection GES-C01 Vce GES-C01 Exam Labs Enter www.practicevce.com and search for ▷ GES-C01 ◁ to download for free Latest GES-C01 Test Testking
- course.hkmhf.org, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, writeablog.net, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, Disposable vapes

BTW, DOWNLOAD part of Rea4Prep GES-C01 dumps from Cloud Storage: <https://drive.google.com/open?id=1UTs-622hJo0tC9g8AumwOkKz1LaknacG>