

GES-C01 Valid Exam Questions & GES-C01 Study Pdf Vce & GES-C01 Latest Study Guide



High efficiency service has won reputation for us among multitude of customers, so choosing our GES-C01 real study dumps we guarantee that you won't be regret of your decision. Helping our candidates to pass the GES-C01 exam and achieve their dream has always been our common ideal. We believe that your satisfactory on our GES-C01 Exam Questions is the drive force for our company. Meanwhile, we adopt a reasonable price for you, ensures people whoever is rich or poor would have the equal access to buy our useful GES-C01 real study dumps.

The GES-C01 certificate stands out among the numerous certificates because its practicability and role to improve the clients' stocks of knowledge and practical ability. Owning a test GES-C01 certificate equals owning a weighty calling card when the clients find jobs and the proof that the clients are the competent people. Our GES-C01 Quiz prep is the great option for the clients to prepare for the test. Our GES-C01 study materials boost high passing rate and hit rate. Our clients praise them highly after they use them and recognize them as the key tool to pass the GES-C01 certification.

>> Exam GES-C01 Guide <<

GES-C01 Latest Test Question & Online GES-C01 Training

The content of our GES-C01 practice braindumps is chosen so carefully that all the questions for the exam are contained. And our GES-C01 study materials have three formats which help you to read, test and study anytime, anywhere. They are the versions of the PDF, Software and APP online. This means with our GES-C01 training guide, you can prepare for exams efficiently. If you desire a GES-C01 certification, our products are your best choice.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q237-Q242):

NEW QUESTION # 237

A data science team is fine-tuning a Snowflake Document AI model to improve the extraction accuracy of specific fields from a new type of complex legal document. They are consistently observing low confidence scores and inconsistent 'value' keys for extracted entities, even after initial training. Which two of the following best practices should the team follow to most effectively improve the model's extraction accuracy and confidence for this complex document type?

- A. Limit the fine-tuning training data exclusively to perfectly formatted and clean documents to ensure the model learns from ideal examples without noise.
- B. Ensure the training dataset used for fine-tuning includes diverse documents representing various layouts, data variations, and explicit examples of values or empty cells where appropriate.
- C. Set the 'temperature' parameter to a higher value (e.g., 0.7) during '!PREDICT' calls to encourage more creative and diverse interpretations by the model.
- D. Prioritize extensive prompt engineering by creating highly detailed and complex questions with intricate logic to guide the LLM's understanding of the extraction task.
- E. Actively involve subject matter experts (SMEs) or document owners throughout the iterative process to help define data values, provide annotations, and evaluate the model's effectiveness.

Answer: B,E

Explanation:

To improve Document AI model training, it is crucial to ensure that the documents uploaded for training represent a real use case and that the dataset consists of diverse documents in terms of both layout and data. If all documents contain the same data or are always presented in the same form, the model might provide incorrect results. For table extraction, it is vital that enough data is used to train the model to include 'NULC' values and maintain order. Therefore, ensuring a diverse training dataset (Option B) is a key best practice. Additionally, Subject Matter Experts (SMEs) and document owners are crucial partners in understanding and evaluating the model's effectiveness in extracting the required information. Their involvement in defining data values, providing annotations, and evaluating results will significantly improve accuracy (Option C). Option A is not a best practice; it's recommended to keep questions as encompassing as possible and rely on training with annotations rather than complex prompt engineering especially for document variability. Option D is incorrect; a higher 'temperature' value increases the randomness and diversity of the model's output, which is generally undesirable for accurate data extraction where deterministic results are preferred. For most consistent results, 'temperature' should be set to 0. Option E is incorrect because training on a restricted set of perfectly formatted documents can lead to a model that performs poorly on real-world, varied documents; diversity in training data is essential.

NEW QUESTION # 238

A data scientist is preparing a large text document for processing with Snowflake Cortex LLMs. They need to use the COUNT_TOKENS function to estimate the input token count for different models. Given the statement

```
SELECT SNOWFLAKE.COREXT.COUNT_TOKENS('mistral-7b', 'This is a sample document to test tokenization.') AS token_estimate;
```

which of the following statements about the COUNT_TOKENS function and its output is TRUE?

- The COUNT_TOKENS



function itself incurs token-based costs, calculated based on the length of the input string, similar to embedding functions.

- The output value from COUNT_TOKENS

for a given text string will always be the same regardless of the specified LLM model (e.g., 'mistral-7b' vs. 'llama3.1-8b').

- The COUNT_TOKENS

function is primarily used to count output tokens generated by an LLM, not the input prompt.

- A token is defined as approximately four characters of text, but the exact equivalence of raw input text to tokens can vary by the Cortex model specified in the function call.

- The COUNT_TOKENS

function can only be used for models available in the same region as the virtual warehouse executing the query.

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

- D. Option C
- E. Option E

Answer: C

Explanation:

Option A is incorrect because

COUNT_TOKENS

incurs only compute cost to run the function, not additional token-based costs. Option B is incorrect as the equivalence of raw input or output text to tokens can vary by the Cortex model, meaning the token count for the same text might differ between models. Option C is incorrect; COUNT_TOKENS is designed to count tokens in an input prompt or text to manage model limits and plan for costs. Option D is correct. An industry convention for text states that a token is approximately equal to four characters, but this equivalence can vary by model, influencing the exact token count. Option E is incorrect because the COUNT_TOKENS function is available in all regions for any model, although the models themselves might have regional availability restrictions for other functions.

NEW QUESTION # 239

A global enterprise has Snowflake accounts in various regions, including a US East (Ohio) account where a critical application is deployed. They need to use AI_COMPLETE with the claude-3-5-sonnet model for real-time customer support, but this model is not natively available in US East (Ohio) for direct AI_COMPLETE usage. The Snowflake administrator considers enabling cross-region inference. Which statements accurately reflect the considerations and characteristics of cross-region inference in Snowflake Cortex?

- A. Latency between regions for cross-region inference is negligible and consistently low, irrespective of cloud provider infrastructure.
- B. Cross-region inference automatically caches user inputs and generated outputs to reduce latency for subsequent requests to the same model.
- C. Setting the CORTEX_ENABLED_CROSS_REGION account parameter to 'ANY_REGION' in the US East (Ohio) account would enable inference requests for claude-3-5- sonnet to be processed in any region where it is natively available.
- D. Cross-region inference is not supported in U.S. SnowGov regions for either inbound or outbound inference requests.
- E. The CORTEX_ENABLED_CROSS_REGION parameter can be configured at the session level to temporarily enable cross-region inference for specific workloads.

Answer: C,D

Explanation:

Option B is correct because setting the parameter to 'ANY_REGION' enables inference requests to be CORTEX_ENABLED_CROSS_REGION processed in a different region from the default, thereby allowing access to models not natively supported in the local region. For example, claude- is 3-5- sonnet available in AWS US East 1 (N. Virginia), which could be accessed from US East (Ohio) via cross-region inference. Option C is 3-5 -sonnet correct as cross-region inference is explicitly not supported in U.S. SnowGov regions. Option A is incorrect because user inputs, service generated prompts, and outputs are not stored or cached during cross-region inference. Option D is incorrect; latency depends on the cloud provider infrastructure and network status, and testing is recommended. Option E is incorrect because CORTEX_ENABLED_CROSS_REGION is an account-level parameter, not a session parameter.

NEW QUESTION # 240

A data engineering team needs to configure their Snowflake environment to process documents using AI_PARSE_DOCUMENT and generate text embeddings using EMBED_TEXT_1024 with the voyage-multilingual-2 model. Their Snowflake account is in a region where these specific capabilities or models are only available via cross-region inference. The team needs to ensure these functions work correctly without constant region-specific model selection. Which of the following is the correct configuration action and an important consideration?

- The CORTEX_ENABLED_CROSS_REGION parameter must be set to TRUE for the account, which implicitly enables all cross-region capabilities.
- The administrator should set the account parameter using: ALTER ACCOUNT SET CORTEX_ENABLED_CROSS_REGION = 'ANY_REGION';
- The AT_PARSE_DOCUMENT function for document processing does not support cross-region inference, therefore documents must be moved to a supported region
- EMBED_TEXT_1024 with voyage-multilingual-2 explicitly supports cross-region inference, and its cost calculation includes both input and output tokens.
- After enabling cross-region inference, it is crucial to test query latency, as it can vary significantly based on cloud provider infrastructure and network status

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

- D. Option B
- E. Option D

Answer: A,D

Explanation:

Option B is correct. Setting `CORTEX_ENABLED_CROSS_REGION` to 'ANY REGION' allows inference requests for supported features/models (like `AI_PARSE_DOCUMENT` and `EMBED_TEXT_1024`) to be processed in a different region if they are not natively available in the local region. Option E is correct because latency between regions depends on the cloud provider infrastructure and network status, and testing specific use cases with cross-region inference enabled is recommended. Option A is incorrect because `CORTEX_ENABLED_CROSS_REGION` takes a list of values or 'ANY_REGION', not a boolean `TRUE`. Option C is incorrect as `AI_PARSE_DOCUMENT` is a Cortex AI SQL function, which generally supports cross-region inference if the model powering it is available in another enabled region. Option D is incorrect; for `EMBED_TEXT_1024` and other embedding functions, only input tokens are counted for billing, not both input and output tokens.

NEW QUESTION # 241

A company is building a chatbot for internal support, powered by Snowflake Cortex LLMs. The primary goals are to provide answers that are accurate, grounded in proprietary documentation, and to minimize factual 'hallucinations'. They are considering various strategies to achieve this. Which of the following statements correctly describe effective methods or tools within Snowflake for addressing these concerns?

- A. For tasks requiring LLMs to generate SQL queries from natural language, using the Cortex Analyst Verified Query Repository (VQR) can improve accuracy by leveraging pre-verified SQL queries for similar questions.
- B. Deploying a custom fine-tuned model using SNOWFLAKE.CORTEX.FINETUNE on proprietary documentation is the most effective approach to ensure factual accuracy for any LLM task.
- C. Using Cortex Search as a Retrieval Augmented Generation (RAG) engine can enhance LLM responses by providing relevant context from proprietary documentation, thereby reducing hallucinations.
- D. AI Observability can be leveraged to systematically evaluate applications, measuring metrics like 'factual correctness and 'groundedness' to detect and mitigate hallucinations, especially in summarization.
- E. Enabling Cortex Guard with guardrails: true directly addresses model hallucinations by ensuring responses are always factually correct and aligned with the provided context.

Answer: A,C,D

Explanation:

Option B is correct: Cortex Search is explicitly designed as a RAG engine to enhance LLM responses with contextualized information from Snowflake data, which directly addresses factual accuracy and reduces hallucinations. Option C is correct: AI Observability's evaluation features, including 'factual correctness and 'groundedness' scores, are specifically mentioned for detecting the truthfulness and relevance of responses based on retrieved context, and for avoiding LLMs with high hallucination frequencies, especially in summarization tasks. Option D is correct: The Cortex Analyst Verified Query Repository (VQR) provides a collection of pre-verified SQL queries for specific natural language questions, significantly improving the accuracy and trustworthiness of SQL generation and reducing errors that could be seen as 'hallucinations' in the text-to-SQL context. Option A is incorrect: While fine-tuning (using 'SNOWFLAKE.CORTEX.FINETUNE') can adapt a model to specific tasks and data, it is not a direct guarantee against 'all' factual inaccuracies or 'hallucinations' for 'any' LLM task, especially if the fine-tuning data itself is limited or the model generalizes poorly. RAG is generally preferred for grounding responses in up-to-date external knowledge. Option E is incorrect: Cortex Guard is designed to filter 'harmful or unsafe' LLM responses, not to directly ensure factual correctness or prevent hallucinations related to content accuracy or grounding.

NEW QUESTION # 242

.....

Are you planning to attempt the Snowflake GES-C01 exam of the GES-C01 certification? The first hurdle you face while preparing for the SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam is not finding the trusted brand of accurate and updated GES-C01 exam questions. If you don't want to face this issue then you are at the trusted spot. FreePdfDump is offering actual and Latest GES-C01 Exam Questions that ensure your success in the Snowflake GES-C01 certification exam on your maiden attempt.

GES-C01 Latest Test Question: <https://www.freepdfdump.top/GES-C01-valid-torrent.html>

But our company, not like these money-oriented ones, always focuses on helping as many people in the field as possible, and we think earning money is a rather trivial aspect of the matter, that's why even though we have become the top notch company in the field we still keep a relative affordable price for our best GES-C01 Latest Test Question vce torrent in the international market,

Snowflake Exam GES-C01 Guide We always put our customers in the first place.

Copy in Response.txt file, Remote Access Design, But our company, not GES-C01 like these money-oriented ones, always focuses on helping as many people in the field as possible, and we think earning money is a rathertrivial aspect of the matter, that's why even though we have become the Online GES-C01 Training top notch company in the field we still keep a relative affordable price for our best Snowflake Certification vce torrent in the international market.

100% Pass Quiz 2026 Snowflake GES-C01 – Efficient Exam Guide

We always put our customers in the first place, GES-C01 Instant Download file is a pioneer in the GES-C01 exam certification preparation, Available for free trial.

You must improve your skills and knowledge to stay current and competitive.

- GES-C01 Practice Questions: SnowPro® Specialty: Gen AI Certification Exam - GES-C01 Exam Dumps Files □ Enter ▷ www.prepawayete.com ▲ and search for □ GES-C01 □ to download for free □ GES-C01 Reliable Braindumps Sheet
- Pdfvce Snowflake GES-C01 Exam Questions are Available in Three Different Formats □ Open ▷ www.pdfvce.com ▲ enter “GES-C01” and obtain a free download □ Exam GES-C01 Simulations
- GES-C01 Trustworthy Source □ Test GES-C01 Simulator □ GES-C01 Download Fee □ Search for ▷ GES-C01 □ on { www.testkingpass.com } immediately to obtain a free download □ Dumps GES-C01 Collection
- Test GES-C01 Result □ GES-C01 Trustworthy Source □ GES-C01 Valid Exam Book □ Easily obtain free download of ▷ GES-C01 □ by searching on [www.pdfvce.com] □ GES-C01 Test Quiz
- Exam GES-C01 Questions □ GES-C01 Valid Dumps Ppt □ Dumps GES-C01 Collection □ Search for ▷ GES-C01 □ and download exam materials for free through ▷ www.troytecdumps.com □ □ Exam GES-C01 Simulations
- Pdfvce Desktop Snowflake GES-C01 Practice Test Software □ Copy URL ⇒ www.pdfvce.com ⇄ open and search for “GES-C01” to download for free □ New GES-C01 Test Forum
- Exam GES-C01 Simulations □ Exam GES-C01 Simulations □ GES-C01 Download Fee □ Open □ www.vceengine.com □ enter ✓ GES-C01 □ ✓ □ and obtain a free download □ GES-C01 Valid Dumps Files
- GES-C01 Download Fee □ Latest Test GES-C01 Experience □ GES-C01 Latest Study Guide □ Open □ www.pdfvce.com □ and search for 「 GES-C01 」 to download exam materials for free □ Exam GES-C01 Questions
- Test GES-C01 Result □ Test GES-C01 Result □ GES-C01 Download Fee □ Search for ▷ GES-C01 □ □ □ and download it for free immediately on ▷ www.practicevce.com ▲ □ GES-C01 Latest Study Guide
- GES-C01 Vce File □ GES-C01 Latest Study Guide □ GES-C01 Exam Vce □ Open □ www.pdfvce.com □ enter ▷ GES-C01 □ and obtain a free download □ New GES-C01 Test Forum
- GES-C01 Practice Questions: SnowPro® Specialty: Gen AI Certification Exam - GES-C01 Exam Dumps Files □ Download ▷ GES-C01 □ □ □ for free by simply searching on ⇒ www.practicevce.com ⇄ □ Latest Test GES-C01 Experience
- bbs.chaken.net.cn, justpaste.me, vxlxemito123.blogspot.com, www.stes.tyc.edu.tw, ycs.instructure.com, www.campfirewriting.com, myportal.utt.edu.tt, www.300300.net, bbs.t-firefly.com, Disposable vapes