


GES-C01 Reliable Exam Vce & GES-C01 Pass Guide



P.S. Free 2025 Snowflake GES-C01 dumps are available on Google Drive shared by PracticeMaterial:
<https://drive.google.com/open?id=1JIO2TmB1Nob-XktbFbNP4QcXyJpHUyHR>

If you are the first time to buy the GES-C01 learning material online, or you have bought them for many times, there may be some problem that puzzle you, if you have any questions about the GES-C01 exam dumps, you can ask our service stuff for help. They have the professional knowledge of GES-C01 Training Materials, and they will be very helpful for solving your problem. In addition, we have free demo for you to try before buying the product, and you can have a try before purchasing.

Only by practising our GES-C01 exam braindumps on a regular base, you will see clear progress happened on you. Besides, rather than waiting for the gain of our GES-C01 practice guide, you can download them immediately after paying for it, so just begin your journey toward success now. With our GES-C01 learning questions, you will find that passing the exam is as easy as pie for our GES-C01 study materials own 100% pass guarantee.

>> GES-C01 Reliable Exam Vce <<

100% Pass Quiz Snowflake - GES-C01 –High Hit-Rate Reliable Exam Vce

About the dynamic change of our GES-C01 guide quiz, they will send the updates to your mailbox according to the trend of the exam. Besides, we understand you may encounter many problems such as payment or downloading GES-C01 practice materials and so on, contact with us, we will be there. Our employees are diligent to deal with your need and willing to do their part 24/7. They always treat customers with courtesy and respect to satisfy your need on our GES-C01 Exam Dumps.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q100-Q105):

NEW QUESTION # 100

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

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 and is available through cross-region inference to accounts in all regions. 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 # 101

A data science team operating in the AWS Europe Central 1 (Frankfurt) region needs to leverage the `snowflake-llama-3.1-405b` model for complex generative AI tasks using the `AI_COMPLETE` function. They've noted that `snowflake-llama-3.1-405b` is not natively available in Frankfurt but is supported via cross-region inference from AWS US regions. The `ACCOUNTADMIN` has already configured the account to include `snowflake-llama-3.1-405b` in the `CORTEX_MODELS_ALLOWLIST`. Despite this, users are consistently encountering errors indicating 'model not found' or 'region not supported'. Which specific Snowflake account parameter, and what value, must be additionally configured by the `ACCOUNTADMIN` to enable the use of `snowflake-llama-3.1-405b` in this scenario?

- A.
The `CORTEX_ENABLED_CROSS_REGION` account parameter must be set to `ANY_REGION` or specifically include `AWS_US_EAST_1` to enable inference requests to be processed in a supported region where the model is available.
- B.
The `ENABLE_CORTEX_ANALYST_MODEL_AZURE_OPENAI` parameter needs to be set to `TRUE` to allow cross-cloud access, even for Snowflake-hosted models.
- C.
A `GRANT USAGE ON LLM snowflake-llama-3.1-405b` statement must be executed for each user role, explicitly enabling cross-region functionality.
- D.
A `GRANT USAGE ON LLM snowflake-llama-3.1-405b` statement must be executed for each user role, explicitly enabling cross-region functionality.
- E. Cross-region inference is automatically managed by Snowflake for allowed models, implying that a new, larger virtual warehouse is required to handle the cross-region data transfer overhead.

Answer: A

Explanation:

Option A is correct. The `CORTEX_ENABLED_CROSS_REGION` account parameter controls whether inference requests can be processed in a different region from the default. To use models available via cross-region inference, this parameter must be set to `ANY_REGION` or a list explicitly including the supported cross-region (e.g., an AWS US region for `snowflake-llama-3.1-405b`). Without this setting, even if the model is in the allowlist, the system won't look for it in other regions. Option B is incorrect as `ENABLE_CORTEX_ANALYST_MODEL_AZURE_OPENAI` specifically pertains to legacy Azure OpenAI models for Cortex Analyst and is not relevant for Snowflake-hosted models or general cross-region inference for LLM functions. Option C is incorrect as access to LLMs is controlled by the `CORTEX_MODELS_ALLOWLIST` parameter and the `SNOWFLAKE.CORTEX_USER` role, not individual `GRANT USAGE ON LLM` statements. Option D is incorrect; the `CORTEX_MODELS_ALLOWLIST` specifies model names, not regions or combinations of regions. Option E is incorrect as cross-region inference does not automatically enable larger warehouses; while latency might be a consideration, it's explicitly controlled by the `CORTEX_ENABLED_CROSS_REGION` parameter, and it does not bypass the need for proper configuration.

NEW QUESTION # 102

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. The retrieved context from Cortex Search should be directly concatenated with the user's prompt as input to the
- **B. Using the**
- **C. To maintain conversational context in a multi-turn chat, the developer should pass all previous user prompts and model responses in the**
- D. For performance and cost optimization, it is always recommended to query Cortex Search and the LLM function within a single
- E. The

`SNOWFLAKE.CORTEX.EMBED_TEXT_768`

Answer: B,C

Explanation:

Option A is incorrect. The user's query is typically embedded (e.g., using `EMBED_TEXT_768`) 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 `prompt_or_history` array to the `COMPLETE` or `AI_COMPLETE` 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 # 103

A data science team is developing an internal LLM to classify legal documents. They previously used a general-purpose LLM, but found its performance for their specific legal domain to be inconsistent, leading to high error rates and increased manual review. They decide to fine-tune a model using Snowflake Cortex Fine-tuning to improve accuracy and reduce latency for real-time document classification. Which base model, among those available for fine-tuning via `SNOWFLAKE.CORTEX.FINETUNE`, is explicitly noted for its low latency and high throughput processing, making it a strong candidate for this use case, especially for multi-page text classification?

- ☐ llama3-70b
known for its general reasoning and large context window, is ideal for comprehensive understanding of legal documents.
- ☐ mixtral-8x7b
optimized for text generation and classification with low memory requirements, translating to higher throughput for enterprise use cases.
- ☐ llama3.1-8b
a cost-effective model, suitable for its large context window, making it adaptable to varied legal document lengths.
- ☐ mistral-7b
specifically highlighted for low latency and high throughput processing for multiple pages of text, with a 32K context window, ideal for summarization, structuration, and question answering tasks that need to be done quickly.
- ☐ snowflake-arctic
Snowflake's top-tier enterprise LLM, excelling at SQL generation and instruction following benchmarks.

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

Answer: B

Explanation:

Option D is correct. The sources explicitly state that mistral-7b

is 'ideal for your simplest summarization, structuration, and question answering tasks that need to be done quickly. It offers low latency and high throughput processing for multiple pages of text with its 32K context window'. This description directly aligns with the scenario's requirement for improved accuracy and reduced latency for real-time document classification, particularly for multi-page legal documents. Option A is incorrect because while

llama3-70b

is available for fine-tuning and suitable for content creation and chat applications, it is not specifically noted for low latency and high throughput processing for multi-page text classification in the same way as

Option B is incorrect because, while

mistral-8x7b

is indeed optimized for low latency with low memory requirements and suitable for classification, the description for

mistral-7b

more directly addresses the 'multiple pages of text' aspect of the classification task. Option C is incorrect;

llama3.1-8b

is a cost-effective model with a large context window, but its performance for low latency and high throughput for multi-page text classification is not highlighted to the same extent as

mistral-7b

Option E is incorrect because

is Snowflake's top-tier enterprise LLM excelling at SQL generation, coding, and instruction following, but it is not listed as a base model available for fine-tuning with SNOWFLAKE.CORTEX.FINETUNE.

NEW QUESTION # 104

A data science team is fine-tuning a mistral-7b model within Snowflake Cortex using proprietary customer interaction logs. Which of the following principles and practices apply to this fine-tuning process concerning data privacy, model ownership, and subsequent inference?

- A. The fine-tuning process occurs entirely within Snowflake's security and governance boundaries, ensuring the data never leaves the Snowflake environment.
- B. The fine-tuned model is fully managed by the Snowflake Model Registry API, allowing programmatic management of its lifecycle.
- C. The fine-tuned model, which is a CORTEX_FINETUNED type, can be shared with other Snowflake accounts using secure data sharing.
- D. The proprietary customer interaction logs used for fine-tuning are leveraged by Snowflake to improve the base mistral-7b model for all customers.
- E. The resulting fine-tuned model is exclusively available to the data science team and cannot be accessed by other Snowflake customers.

Answer: A,C,E

Explanation:

Cortex Fine-tuning is a fully managed service that lets you fine-tune popular LLMs using your data, all within Snowflake. Your Usage and Customer Data (including inputs and outputs) are NOT used to train, re-train, or fine-tune Models made available to others. Fine-tuned Models built using your data can only be used by you. Therefore, Option A is incorrect. Options B and C are correct, as the fine-tuned model is exclusive to the customer, and the process is managed within Snowflake's boundaries. Option D is also correct because models generated with Cortex Fine-tuning (CORTEX_FINETUNED type) can be shared using Data Sharing. Option E is incorrect because Cortex Fine-Tuned LLMs appear in the model registry's Snowsight UI, but are not managed by the model registry API.

NEW QUESTION # 105

.....

Our product's passing rate is 99% which means that you almost can pass the test with no doubts. The reasons why our GES-C01

GES-C01 Pass Guide: <https://www.practicematerial.com/GES-C01-exam-materials.html>

Now, you may be preparing for the GES-C01 Exam Test, If you choose our GES-C01 study questions as your study tool, we can promise that we will try our best to enhance the safety guarantees Exam GES-C01 Overviews and keep your information from revealing, and your privacy will be protected well.

We get information from special channel.

- P.S. Free 2025 Snowflake GES-C01 dumps are available on Google Drive shared by PracticeMaterial: <https://drive.google.com/open?id=1JIO2TmB1Nob-XktbFbNP4QcXyJpHUyHR>

