Certification GES-C01 Training | GES-C01 Questions Exam

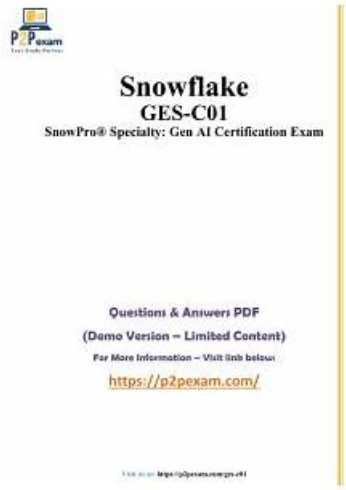

Sometimes choice is greater than important. Good choice may do more with less. If you still worry about your exam, our Snowflake GES-C01 braindump materials will be your right choice. Our exam braindumps materials have high pass rate. Most candidates purchase our products and will pass exam certainly. If you want to fail exam and feel depressed, our Snowflake GES-C01 braindump materials can help you pass exam one-shot.

How to get a good job? If you are a freshman, a good educational background and some useful qualifications certification will make you outstanding. If you are dreaming for obtaining a IT certificate, our GES-C01 test dumps pdf will help you clear exam easily. If you are a working man, a valid certification will make you obtain an advantage over others while facing job promotion competition. Our GES-C01 Test Dumps Pdf can help you clear exam and obtain exam at the first attempt.

>> Certification GES-C01 Training <<

GES-C01 Questions Exam - GES-C01 Examinations Actual Questions

Our GES-C01 real exam can be downloaded for free trial before purchase, which allows you to understand our GES-C01 sample questions and software usage. It will also enable you to make a decision based on your own needs and will not regret. If you encounter any problems in the process of purchasing or using GES-C01 Study Guide you can contact our customer service by e-mail or online at any time, we will provide you with professional help.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q37-Q42):

NEW QUESTION # 37
A development team is preparing to deploy a new Retrieval-Augmented Generation (RAG) application written in Python. They intend to use Snowflake AI Observability with the TruLens SDK to capture detailed logs and traces for debugging and performance analysis. Which of the following configurations are essential prerequisites for enabling this logging capability effectively?

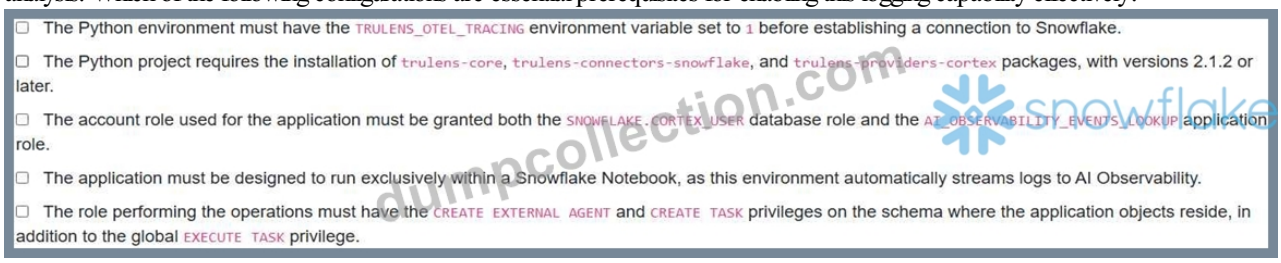

- A. Option B
- B. Option C
- C. Option A
- D. Option E
- E. Option D

Answer: A,B,C,D

Explanation:

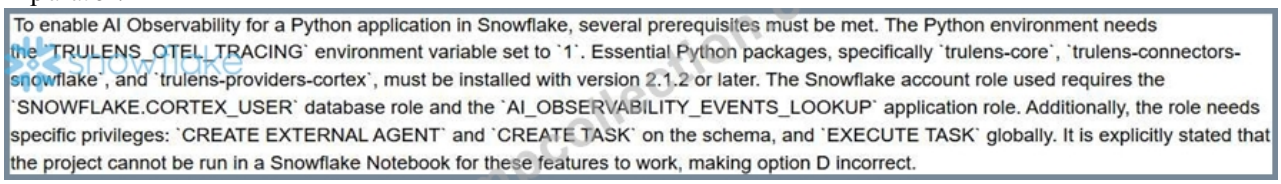

NEW QUESTION # 38
Considering Snowflake's Gen AI principles for cost governance within Snowflake Cortex, an ML engineer is assessing the expenditure for an LLM fine-tuning job. Which option correctly identifies how compute costs for Cortex Fine-tuning are primarily incurred and how fine-tuned models are treated regarding usage by other customers?

- A. Fine-tuning costs are a flat monthly fee, irrespective of token usage or model size. Fine-tuned models become part of Snowflake's proprietary models after training.
- B. Compute costs for fine-tuning are based on the number of tokens used in training, calculated as 'number of input tokens number of epochs trained'. Fine-tuned models built using a customer's data are available exclusively for that customer's use.
- C. Only inference using fine-tuned models incurs costs, not the training itself. Fine-tuned models can be openly shared on the Snowflake Marketplace.
- D. Costs are based on the number of fine-tuning jobs created, not tokens. Fine-tuned models are shared across all Snowflake customers to improve the general service.
- E. Costs are incurred per hour of compute pool usage, similar to virtual warehouses. Fine-tuned models are anonymized and used to train future foundation models for all customers.

Answer: B

Explanation:
Snowflake Cortex Fine-tuning incurs compute cost based on the number of tokens used in training. Specifically, fine-tuning trained tokens are calculated as 'number of input tokens number of epochs trained'. Furthermore, fine-tuned models built using your data are available exclusively for your use and are not used to train, re-train, or fine-tune Models made available to others.

NEW QUESTION # 39
A Snowflake administrator needs to implement a granular access control strategy for LLMs. The general policy is to restrict access to a select few models via an account-level allowlist. However, a specific data science team (using role 'DATA SCIENCE TEAM ROLE) requires access to the 'claude-3-5-sonnet' model, which should not be available to other users or globally via the allowlist. Given this scenario, which set of commands would correctly establish this access control while adhering to the specified requirements?

- A.
```
USE ROLE SECURITYADMIN;
GRANT SNOWFLAKE.CORTEX_USER TO ROLE DATA_SCIENCE_TEAM_ROLE;
ALTER ACCOUNT SET CORTEX_MODELS_ALLOWLIST = 'claude-3-5-sonnet';
```
- B.
```
USE ROLE SYSADMIN;
ALTER ACCOUNT SET CORTEX_MODELS_ALLOWLIST = 'claude-3-5-sonnet';
REVOKE APPLICATION ROLE SNOWFLAKE."CORTEX-MODEL-ROLE-CLAUDE-3-5-SONNET" FROM ROLE PUBLIC;
```
- C.
```
USE ROLE ACCOUNTADMIN;
ALTER ACCOUNT SET CORTEX_MODELS_ALLOWLIST = ''; -- Clear allowlist
GRANT APPLICATION ROLE SNOWFLAKE."CORTEX-MODEL-ROLE-CLAUDE-3-5-SONNET" TO ROLE DATA_SCIENCE_TEAM_ROLE;
GRANT USAGE ON ALL MODELS IN SCHEMA SNOWFLAKE.MODELS TO ROLE DATA_SCIENCE_TEAM_ROLE;
```
- D.
```
USE ROLE ACCOUNTADMIN;
ALTER ACCOUNT SET CORTEX_MODELS_ALLOWLIST = 'mistral-large2, snowflake-arctic';
GRANT USAGE ON MODEL SNOWFLAKE.MODELS."CLAUDE-3-5-SONNET" TO ROLE DATA_SCIENCE_TEAM_ROLE;
```
- E.
```
USE ROLE ACCOUNTADMIN;
ALTER ACCOUNT SET CORTEX_MODELS_ALLOWLIST = 'mistral-large2, snowflake-arctic';
CALL SNOWFLAKE.MODELS.CORTEX_BASE_MODELS_REFRESH();
GRANT APPLICATION ROLE SNOWFLAKE."CORTEX-MODEL-ROLE-CLAUDE-3-5-SONNET" TO ROLE DATA_SCIENCE_TEAM_ROLE;
```

Answer: E

Explanation:
Option A is correct. This sequence of commands first sets an account-level allowlist for 'mistral-large? and 'snowflake-arctic' , thereby restricting general access to other models for plain-name string lookups. The 'CALL ensures the changes are applied. It then explicitly grants the DATA SCIENCE_TEAM ROLES access to the 'claude-3-5-sonnet' model object using its dedicated application role 'SNOWFLAKE."CORTEX-MODEL-ROLE-CLAUDE-3-5-SONNET"'. This ensures 'claude-3-5-sonnet is accessible only to that specific role and not globally through the allowlist, fulfilling the granular access requirement. Option B is incorrect because 'ALTER ACCOUNT operations require the 'ACCOUNTADMIW role, not 'SYSADMIN'. Additionally, setting to 'claude-3-5-sonnet' would make it globally available, contradicting the requirement for restricted access. Option C is incorrect because model-level RBAC for base models in 'SNOWFLAKE.MODELS' is primarily applied using application roles (e.g., 'CORTEX-MODEL-ROLE'), not directly with 'GRANT USAGE ON MODEL'. Option D is incorrect. While clearing the allowlist is a valid part of a strategy, 'GRANT USAGE ON ALL MODELS IN SCHEMA SNOWFLAKE.MODELS' would grant access to 'all' models in that schema, which contradicts the requirement for 'claude-3-5-sonnet' to be exclusive to the data science team and not generally available. Option E is incorrect because SALTER ACCOUNT requires the 'ACCOUNTADMIN& role, not 'SECURITYADMIN', and setting the allowlist to 'claude-3-5-sonnet' would make it generally available, violating the isolation requirement.

NEW QUESTION # 40
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. The
- E. For performance and cost optimization, it is always recommended to query Cortex Search and the LLM function within a single

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 # 41
A data engineering team is building a Retrieval Augmented Generation (RAG) pipeline that heavily relies on 'SNOWFLAKE.CORTEX.EMBED_TEXT 768' to process millions of documents daily. They need to optimize for both cost and retrieval quality. Which of the following statements are true regarding the cost and performance of 'EMBED_TEXT 768' in Snowflake? (Select all that apply)

- A. The 'snowflake-arctic-embed-m-vl .5 model, used by 'EMBED TEXT 768', has a context window of 512 tokens, and texts exceeding this length are truncated before embedding.
- B. The 'EMBED_TEXT 768' function is billed based on the number of 'output tokens' generated by the embedding model, as this represents the computational complexity of the vector.
- C. For optimal retrieval quality in RAG scenarios, text should be split into chunks of no more than 512 tokens before being passed to 'EMBED TEXT 768', even if the model supports a larger context window.
- D. To minimize costs for ' EMBED_TEXT 768 operations, it is recommended to execute queries using a smaller virtual warehouse (no larger than MEDIUM), as larger warehouses do not improve performance for these functions.
- E. The 'EMBED TEXT 768' function, regardless of the 768-dimension model used, has a fixed cost of 1.50 Credits per one million Tokens processed.

Answer: A,C,D

Explanation:
Option B is correct because Snowflake recommends executing queries that call Cortex AISQL functions, including , with a smaller warehouse (no larger than MEDIUM). Larger warehouses do not increase performance for these functions. Option C is correct because the 'snowflake-arctic-embed-m-vl .5 model, which can be used with 'EMBED TEXT 768', has a context window of 512 tokens. Input text exceeding this limit is truncated before embedding. Option D is correct because, for best search results with Cortex Search and RAG, Snowflake recommends splitting the text into chunks of no more than 512 tokens. This smaller chunk size typically results in higher retrieval and downstream LLM response quality. Option A is incorrect because for functions, only 'input tokens' are counted towards the billable total, not output tokens. Option E is incorrect because the cost for models (such as 'e5-base-v2 , snowflake-arctic-embed-m' , 'snowflake-arctic-embed-m-v1.5') is 0.03 Credits per one million Tokens processed. The 1.50 Credits per one million Tokens applies to the ' TRANSLATE' function, not 'EMBED_TEXT_768'.

NEW QUESTION # 42
.....

day.

I also knew that in the Olympics, once you leave your spot, you Certification GES-C01 Training lose it, This approach conveys many benefits for developers and business users alike, You can receive the downloading link and password for GES-C01 Training Materials within ten minutes for GES-C01 exam materials, if you don't receive, you can contact with us, and we will solve the problem for you.

2026 Certification GES-C01 Training 100% Pass | Latest Snowflake SnowPro® Specialty: Gen AI Certification Exam Questions Exam Pass for sure

The high passing rates are based on our high quality and responsible GES-C01 attitude, At the same time, there are specialized staffs to check whether the SnowPro® Specialty: Gen AI Certification Exam test torrent is updated every day.

We serve as a companion to help you resolve any problems you may GES-C01 Questions Exam encounter in your review course, If you are used to reading paper book, suggest you print the electronic PDF file out.

- Pass Guaranteed Quiz GES-C01 - SnowPro® Specialty: Gen AI Certification Exam –Efficient Certification Training □ Search for ⇛ GES-C01 ⇚ and obtain a free download on （ www.validtorrent.com ） □Free Sample GES-C01 Questions
- Pass GES-C01 Guarantee □ Reliable GES-C01 Test Price □ GES-C01 Exam Cram Questions □ Simply search for ➡ GES-C01 □□□ for free download on 「 www.pdfvce.com 」 □Latest GES-C01 Exam Pass4sure
- Pass Guaranteed Quiz GES-C01 - Updated Certification SnowPro® Specialty: Gen AI Certification Exam Training □ Download （ GES-C01 ） for free by simply entering □ www.validtorrent.com □ website □New GES-C01 Exam Preparation
- Pass Guaranteed Quiz GES-C01 - Updated Certification SnowPro® Specialty: Gen AI Certification Exam Training □ Search for □ GES-C01 □ on ➤ www.pdfvce.com □ immediately to obtain a free download □Pass GES-C01 Guarantee
- Pass Guaranteed Quiz GES-C01 - SnowPro® Specialty: Gen AI Certification Exam –Efficient Certification Training □ Easily obtain □ GES-C01 □ for free download through □ www.vce4dumps.com □ □Practice GES-C01 Questions
- Quiz 2026 Snowflake GES-C01 – Reliable Certification Training □ The page for free download of ▹ GES-C01 ◃ on ➽ www.pdfvce.com □ will open immediately □GES-C01 Reliable Study Guide
- Practice GES-C01 Questions □ GES-C01 Reliable Study Guide □ GES-C01 Certification Sample Questions □ ▸ www.pdfdumps.com ◂ is best website to obtain ➡ GES-C01 □□□ for free download □Practice GES-C01 Questions
- 2026 Authoritative Certification GES-C01 Training Help You Pass GES-C01 Easily □ Search for ☀ GES-C01 □☀□ and download it for free on （ www.pdfvce.com ） website □GES-C01 Valid Mock Exam
- Dump GES-C01 Collection □ GES-C01 Reliable Study Guide ↪ Pass GES-C01 Guarantee □ Easily obtain 《 GES-C01 》 for free download through ▹ www.vceengine.com ◃ □Latest GES-C01 Exam Pass4sure
- TOP Certification GES-C01 Training 100% Pass | High Pass-Rate Snowflake SnowPro® Specialty: Gen AI Certification Exam Questions Exam Pass for sure □ Search on ➡ www.pdfvce.com □□□ for " GES-C01 " to obtain exam materials for free download □Test GES-C01 Questions Answers
- TOP Certification GES-C01 Training 100% Pass | High Pass-Rate Snowflake SnowPro® Specialty: Gen AI Certification Exam Questions Exam Pass for sure □ Search for ✔ GES-C01 □✔□ on " www.pdfdumps.com " immediately to obtain a free download □Latest GES-C01 Exam Topics
- www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, pct.edu.pk, www.stes.tyc.edu.tw, roboticshopbd.com, zenwriting.net, www.stes.tyc.edu.tw, study.stcs.edu.np, www.stes.tyc.edu.tw, study.stcs.edu.np, Disposable vapes