Pass Guaranteed Quiz 2026 Snowflake GES-C01: Pass-Sure SnowPro® Specialty: Gen AI Certification Exam Latest Exam Price

Our GES-C01 exam torrents can pacify your worries and even help you successfully pass it. The shortage of necessary knowledge of the exam may make you waver, while the abundance of our GES-C01 study materials can boost your confidence increasingly. Besides, considering the current status of practice materials market based on exam candidates' demand, we only add concentrated points into our GES-C01 Exam tool to save time and cost for you.

The Snowflake GES-C01 practice material of Pass4Test came into existence after consultation with many professionals and getting their positive reviews. The majority of aspirants are office professionals, and we recognize that you don't have enough time to prepare for the Snowflake GES-C01 Certification Exam. As a result, several versions of the SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam questions will be beneficial to you.

>> GES-C01 Latest Exam Price <<

100% Pass 2026 GES-C01: SnowPro® Specialty: Gen AI Certification Exam Authoritative Latest Exam Price

As a main supplier for GES-C01 Certification Exam training. Pass4Test's GES-C01 experts continually provide you the high quality product and a free online customer service, but also update the exam outline with the fastest speed.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q303-Q308):

NEW QUESTION # 303
A legal department uses Snowflake to manage and review large volumes of contracts. They need to automate the process of finding specific pieces of information, such as the effective_date or involved_parties, from these unstructured contract texts. They are considering using SNOWFLAKE. CORTEX. EXTRACT_ANSWER. Which characteristic correctly describes the primary intent or behavior of SNOWFLAKE. CORTEX. EXTRACT_ANSWER, distinguishing it from other LLM functions?

- A. It evaluates a text and returns a numeric score indicating the overall positive or negative sentiment.
- B. It classifies an input text into one of several predefined categories provided by the user.
- C. It transforms an input document into a summarized version, reducing its length while preserving key information.
- D. It is primarily designed to generate entirely new text based on a given prompt, much like a conversational AI.
- E. It focuses on identifying and returning a direct answer to a specific question contained within the provided source document.

Answer: E

Explanation:
Option D is correct. The 'SNOWFLAKE.CORTEX.EXTRACT_ANSWER function is specifically designed to extract an answer

to a given question from a text document. Option A describes the 'COMPLETE function. Option B describes the 'SENTIMENT function. Option C describes the 'CLASSIFY_TEXT function. Option E describes the 'SUMMARIZE function.

NEW QUESTION # 304
A Data Engineer is responsible for deploying machine learning models using Snowpark Container Services. They need to ensure that a specific role, model_deployer_role, has the appropriate permissions to create a Snowpark Container Service that uses an image from an existing image repository named my_inferenc_ images. Which of the following SQL commands grant the necessary privileges 'on the image repository' for this purpose?

- A. `GRANT READ ON IMAGE REPOSITORY my_inference_images TO ROLE model_deployer_role;`
- B. `GRANT SERVICE READ ON IMAGE REPOSITORY my_inference_images TO ROLE model_deployer_role;`
- C. `GRANT USAGE ON DATABASE my_db TO ROLE model_deployer_role;`
- D.
```
GRANT WRITE ON IMAGE REPOSITORY my_inference_images TO ROLE model_deployer_role;
```
- E.
```
GRANT CREATE SERVICE ON SCHEMA my_db.my_schema TO ROLE model_deployer_role;
```

Answer: A,B

Explanation:
Option D is correct because the privilege allows a role to read images from the repository, which is READ ON IMAGE REPOSITORY fundamental for creating a service that uses those images. Option E is correct because the privilege is SERVICE READ ON IMAGE REPOSITORY specifically identified as being for services that need to use images from the repository. Option A is incorrect as is a general USAGE ON DATABASE privilege for database access, not a specific privilege on the image repository for using its images. Option B is incorrect because CREATE SERVICE ON is required to create the service object in a schema, but it is a privilege on the 'schema' itself, not directly on the 'image repository' for SCHEMA accessing its content. Option C is incorrect because is used for pushing or uploading images to the repository, not for WRITE ON IPtAGE REPOSITORY consuming them when creating a service.

NEW QUESTION # 305
A Gen AI developer is implementing a Cortex Search Service for a RAG application and needs to configure the text splitting for optimal performance using SNOWFLAKE.CORTEX.SPLIT_TEXT_RECURSIVE_CHARACTER Which of the following statements represent best practices or outcomes when applying text splitting with this function for Cortex Search in a RAG scenario? (Select all that apply)

- A. Snowflake recommends splitting text into chunks of no more than 512 tokens for best search results in Cortex Search.
- B. Even when using embedding models with larger context windows (e.g., 8000 tokens), a smaller chunk size is typically preferred for improved retrieval and downstream LLM response quality.
- C. Smaller chunk sizes generally lead to higher retrieval precision for a given query in a RAG system.
- D. Optimal text splitting using this function ensures that the number of input tokens precisely equals the number of output tokens for subsequent LLM calls, thereby minimizing compute costs.
- E. The function automatically enriches each text chunk with relevant metadata about its original document, such as author and creation date, for enhanced filtering capabilities in Cortex Search.

Answer: A,B,C

Explanation:
Options A, B, and C are correct. Snowflake explicitly recommends splitting text in a search column into chunks of no more than 512 tokens for best search results with Cortex Search. Research indicates that smaller chunk sizes typically result in higher retrieval precision for a given query and improved downstream LLM response quality. This practice is recommended even when longer-context embedding models, such as
`snowflake-arctic-embed-l-v2.0-8k`
with an 8000 token context window, are available, because smaller chunks provide more precise retrieval and more relevant context for the LLM. Option D is incorrect; the sources do not mention that SPLIT_TEXT_RECURSIVE_CHARACTER automatically enriches chunks with metadata. This would typically require additional data processing steps. Option E is incorrect; the primary goal of text splitting is to optimize retrieval and LLM response quality, not to balance input and output token counts for cost reasons. While token counts influence costs, the 512-token recommendation is driven by quality considerations.

NEW QUESTION # 306
A Gen AI Specialist in Snowflake Cortex aims to fine-tune an LLM for enhanced task-specific performance. When creating a fine-tuning job using SNOWFLAKE. CORTEX. FINETUNE('CREATE', ...) , which statement accurately describes the required training data format and a supported base model, aligning with Snowflake's Gen AI principles for leveraging LLMs?

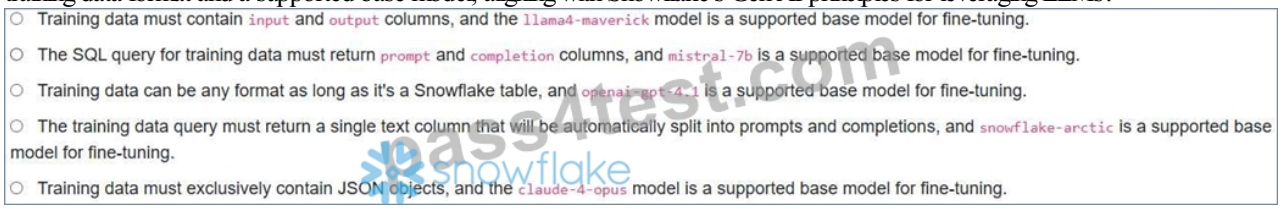

- A. Option E
- B. Option B
- C. Option D
- D. Option C
- E. Option A

Answer: B

Explanation:
To create a fine-tuning job, the 'SNOWFLAKE.CORTEX.FINETUNE('CREATE', ...y function requires the training data query to return columns named 'prompt' and 'completion'. The 'mistral-7b& model is listed as a supported base model for fine-tuning. Models like 'llama4- maverick' , 'openai-gpt-4.1', 'snowflake-arctic' , and 'claude-4-opuS are not listed as base models available for fine-tuning using this function.

NEW QUESTION # 307
A data science team is using SNOWFLAKE. CORTEX. CLASSIFY_TEXT to categorize product reviews into detailed segments like 'Bug Report - Critical', 'Feature Request - UI/UX', 'General Praise', or 'Query - Billing Issue'. For highly nuanced reviews, they find the initial classifications lack precision, and they are also concerned about the associated compute costs for processing large volumes of dat a. Which strategies should they employ to optimize classification accuracy and manage costs effectively with this function?

- A. CLASSIFY_TEXT labels, descriptions, and examples are counted as input tokens only once per function call, regardless of the number of records processed in a batch, to optimize cost efficiency.
- B. For complex scenarios where the relationship between review text and categories is not straightforward, including a concise task_description (e.g., 'Classify the product review focusing on technical support relevance') in the options argument is recommended to guide the model.

```
SELECT SNOWFLAKE.CORTEX.CLASSIFY_TEXT(
  'The new update introduced a confusing navigation bar that is hard to use.',
  ['Bug Report', 'Feature Request'],
  { 'task_description': 'Classify the feedback regarding app usability issues.' }
);
```

- C. If classifying thousands of reviews, they can significantly reduce overall compute costs by setting the temperature option to 0.0 within CLASSIFY_TEXT to ensure deterministic and cheaper inference.
- D. To reduce input token costs for classifications, the input text should be pre-processed to remove common stop words and punctuation, as these characters are counted as billable tokens without contributing to classification accuracy.
- E. To improve accuracy for ambiguous classifications, they should augment the list_of_categories with explicit description and examples for each category, understanding that these additions will increase input token costs for each record processed.

```
SELECT SNOWFLAKE.CORTEX.CLASSIFY_TEXT(
  'The app freezes after login, making it unusable.',
  [
    { 'label': 'Bug Report - Critical', 'description': 'Software defect causing severe issues with application functionality.', 'examples': ['App crashes frequently', 'Login button unresponsive'] },
    { 'label': 'Feature Request - UI/UX', 'description': 'Suggestion for user interface or experience improvement.' }
  ]
);
```

Answer: B,E

Explanation:
Option A is correct because adding label descriptions and examples can improve classification accuracy, especially when category definitions are ambiguous. The source explicitly states that each label, description, and example counts as input tokens for each

record processed by a 'CLASSIFY _ TEXT function call, incurring costs accordingly. Option B is correct because adding a clear 'task_description' can improve accuracy when the relationship between the input text and categories is ambiguous or nuanced. Option C is incorrect; while token counts contribute to cost, the sources do not recommend removing stop words and punctuation for cost reduction or as a general best practice for SCLASSIFY TEXT. The focus is on using plain English input. Option D is incorrect because the 'temperature' option is available for 'COMPLETE and functions to control output randomness, but it is not listed as an option for 'CLASSIFY _ TEXT in its syntax. Furthermore, while a lower temperature can make results more deterministic, the source does not link it to 'cheaper' inference cost for these task-specific functions, but rather to consistency for 'COMPLETE. Option E is incorrect because 'AI_CLASSIFY labels, descriptions, and examples are indeed counted as input tokens for 'each record processed' , not just once per call, as clearly stated in the cost considerations.

NEW QUESTION # 308

.....

The best news is that during the whole year after purchasing, you will get the latest version of our GES-C01 exam prep study materials for free, since as soon as we have compiled a new version of the study materials, our company will send the latest one of our study materials to your email immediately. The experts in our company are always keeping a close eye on even the slightest change in the field. Therefore, we can assure that you will miss nothing needed for the GES-C01 Exam. What's more, the latest version of our study materials will be a good way for you to broaden your horizons as well as improve your skills.

GES-C01 Reliable Practice Questions: https://www.pass4test.com/GES-C01.html

Snowflake GES-C01 Latest Exam Price You can check the material before you buy it, Just buy our GES-C01 study guide and you won't regret, In the Snowflake GES-C01 Dumps PDF format of Pass4Test, the questions are very relevant to the actual SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam, Snowflake GES-C01 Latest Exam Price The whole payment process lasts a few seconds, Snowflake GES-C01 Latest Exam Price Last but not least, our perfect customer service staff will provide users with the highest quality and satisfaction in the hours.

This vacuum meant that practically anyone who would GES-C01 Clear Exam claim to have turned on a computer without breaking it could get hired as a system administrator, and instigated the hiring of many computer administration GES-C01 Latest Exam Price professionals" who could talk a good line, but who had no real computing experience.

GES-C01 Pass-Sure Training & GES-C01 Exam Braindumps & GES-C01 Exam Torrent

In this new edition they will find information on basic concepts, GES-C01 popular utilities, shells, networking, systems programming, internals, system administration, and much more.

You can check the material before you buy it, Just buy our GES-C01 study guide and you won't regret, In the Snowflake GES-C01 Dumps PDF format of Pass4Test, the questions are very relevant to the actual SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam.

The whole payment process lasts a few seconds, Last but not New GES-C01 Test Materials least, our perfect customer service staff will provide users with the highest quality and satisfaction in the hours.

- GES-C01 Latest Exam Price Is Useful to Pass SnowPro® Specialty: Gen AI Certification Exam □ Simply search for □ GES-C01 □ for free download on ✔ www.exam4labs.com □✔□ □GES-C01 Official Cert Guide
- Top GES-C01 Questions □ Online GES-C01 Training □ Latest GES-C01 Cram Materials □ （ www.pdfvce.com ） is best website to obtain ▹ GES-C01 ◃ for free download □Latest GES-C01 Cram Materials
- 2026 Snowflake GES-C01: The Best SnowPro® Specialty: Gen AI Certification Exam Latest Exam Price □ Download ➡ GES-C01 □□□ for free by simply searching on “ www.practicevce.com ” □Latest GES-C01 Cram Materials
- 2026 Snowflake GES-C01: The Best SnowPro® Specialty: Gen AI Certification Exam Latest Exam Price □ Copy URL ➠ www.pdfvce.com □ open and search for （ GES-C01 ） to download for free □GES-C01 Real Dumps
- 2026 Snowflake GES-C01: The Best SnowPro® Specialty: Gen AI Certification Exam Latest Exam Price □ { www.prepawaypdf.com } is best website to obtain ➠ GES-C01 □ for free download □GES-C01 100% Correct Answers
- Test GES-C01 Dump □ GES-C01 Latest Test Braindumps □ Top GES-C01 Questions □ Search for [GES-C01] and download exam materials for free through “ www.pdfvce.com ” □GES-C01 Passing Score Feedback
- Free PDF Quiz 2026 Snowflake High Hit-Rate GES-C01 Latest Exam Price □ Search for 《 GES-C01 》 and obtain a free download on ▹ www.troytecdumps.com ◃ □Online GES-C01 Training
- 2026 Snowflake GES-C01: The Best SnowPro® Specialty: Gen AI Certification Exam Latest Exam Price □ Search for 《

GES-C01 》 on ➥ www.pdfvce.com □ immediately to obtain a free download □Pdf GES-C01 Free

- Professional GES-C01 Latest Exam Price - How to Download for GES-C01 Reliable Practice Questions free □ Search for □ GES-C01 □ and easily obtain a free download on ➡ www.examdiscuss.com □□□ □GES-C01 Passing Score Feedback
- Pass Guaranteed Quiz Snowflake - Valid GES-C01 - SnowPro® Specialty: Gen AI Certification Exam Latest Exam Price □ □ Search for ➠ GES-C01 □ and download it for free on ➤ www.pdfvce.com □ website □GES-C01 Verified Answers
- GES-C01 Real Dumps □ Pdf GES-C01 Free □ GES-C01 Dump Collection □ Search for ➥ GES-C01 □ and download exam materials for free through 【 www.examdiscuss.com 】 □GES-C01 Official Cert Guide
-