

New Release GES-C01 PDF Questions [2026] - Snowflake GES-C01 Exam Dumps



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>

BONUS!!! Download part of Lead2Passed GES-C01 dumps for free: <https://drive.google.com/open?id=179AmbQFyfOk2IMHbrUBG1MvixFZT6xxP>

We attract customers by our fabulous GES-C01 certification material and high pass rate, which are the most powerful evidence to show our strength. We are so proud to tell you that according to the statistics from our customers' feedback, the pass rate of our GES-C01 exam questions among our customers who prepared for the exam with our GES-C01 Test Guide have reached as high as 99%, which definitely ranks the top among our peers. Hence one can see that the GES-C01 learn tool compiled by our company are definitely the best choice for you.

What does it mean to win a competition? Users of our GES-C01 actual exam can give you good answers. They have improved their strength and proved their strength. Now they have more opportunities and they have the right to choose. Of course, the effective learning methods they learned during the use of our GES-C01 Preparation materials also greatly enhanced their work. All of them had praised that our GES-C01 exam questions are the best choice they had made to buy. So what are you waiting for? Just rush to buy our GES-C01 practice guide!

>> GES-C01 Certification Questions <<

Snowflake GES-C01 Practice Test | GES-C01 Latest Version

If you choose to sign up to participate in Snowflake certification GES-C01 exams, you should choose a good learning material or training course to prepare for the examination right now. Because Snowflake Certification GES-C01 Exam is difficult to pass. If you want to pass the exam, you must have a good preparation for the exam.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q92-Q97):

NEW QUESTION # 92

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

Answer: B,C,D,E

Explanation:

NEW QUESTION # 93

A Snowflake Gen AI Specialist is defining a semantic model for Cortex Analyst to improve text-to-SQL accuracy. They are adding entries to the `verified_queries` section of their YAML file. Consider the following semantic model snippet and a proposed `verified_query` entry. Which of the following statements correctly identifies an issue or a best practice not followed in the `sql` field of the proposed `verified_query` entry, based on Cortex Analyst VQR guidelines? Semantic Model Snippet:

Proposed `verified_query` entry:

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

Answer: B,E

Explanation:

NEW QUESTION # 94

A development team is evaluating Snowpark Container Services (SPCS) for deploying various AI/ML workloads, including custom LLMs and GPU-accelerated model training. They need to understand its core benefits and operational characteristics compared to traditional container orchestration platforms. Which of the following statements accurately describe the benefits and/or operational characteristics of Snowpark Container Services for deploying third-party models and AI applications?

- A. SPCS primarily supports applications written in Python and Java, with limited experimental support for other programming languages.
- B. Snowpark Container Services ensures that data remains within Snowflake's security and governance boundaries, eliminating the need to move data out of the environment for processing.
- C. Compute clusters within SPCS are designed to auto-scale dynamically based on workload demand, automatically adjusting the number of instances.
- D. SPCS provides a fully managed OCI runtime execution environment, allowing users to run containerized workloads directly within Snowflake without managing underlying Docker or Kubernetes infrastructure.
- E. SPCS supports both long-running services (e.g., web applications) and finite-lifespan job services (e.g., GPU-accelerated machine learning model training).

Answer: B,D,E

Explanation:

Option A is correct. Snowpark Container Services is designed to run containerized workloads directly within Snowflake, ensuring that data does not need to be moved out of the Snowflake environment for processing, thereby maintaining security and governance.

Option B is correct. SPCS is a fully managed container offering that provides an OCI runtime execution environment optimized for Snowflake, abstracting away the complexities of traditional container orchestration platforms like Docker or Kubernetes. Option C is incorrect. Snowpark Container Services allows application code to be written in 'any language' (e.g., C++) and to use 'any libraries', not just Python and Java. Option D is correct. SPCS supports long-running services, which explicitly stop when no longer needed (like web services), and job services, which have a finite lifespan and run to completion (like GPU-accelerated ML model training). Option E is incorrect. The size of compute clusters in Snowpark Container Services does not auto-scale; users must manually alter the number of instances at runtime using commands like 'ALTER SERVICE MIN INSTANCES =!'.

NEW QUESTION # 95

A data engineering manager needs to audit Cortex LLM function costs to identify specific SQL queries that are unexpectedly high in token consumption for the 'lama3.1-8b' model. They require granular analysis of prompt, completion, and guardrail token usage for these queries. Which of the following Snowflake methods or views would provide the necessary insights?

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

Answer: A,B

Explanation:

NEW QUESTION # 96

A data engineering team is setting up a Retrieval Augmented Generation (RAG) application using Snowflake Cortex Search to provide contextual answers from customer support transcripts. The transcripts are stored in a Snowflake table named SUPPORT_TRANSCRIPTS. Which of the following statements are crucial considerations or accurate facts regarding the initial setup and configuration of the Cortex Search Service for this use case?

- A. Cortex Search Services currently support replication and cloning, allowing for easy disaster recovery and geographical distribution of the search index.
- B. Columns specified in the ATTRIBUTES field during service creation are only used for filtering search results and do not need to be present in the source query.
- C. The CREATE CORTEX SEARCH SERVICE command requires that CHANGE_TRACKING = TRUE be enabled on the source table, especially if the role creating the service is not the table owner. This ensures that the service can track updates to the base data.
- D. Cortex Search is designed to get users up and running quickly with a hybrid (vector and keyword) search engine on text data, handling embedding, infrastructure maintenance, and search quality parameter tuning automatically.
- E. Snowflake recommends using a dedicated virtual warehouse of any size, including X-Large or 2X-Large, for each Cortex Search Service to ensure the fastest possible materialization of search indexes during creation and refresh.

Answer: C,D

Explanation:

Option A is correct because change tracking is required for the Cortex Search Service to monitor updates to the base table, particularly if the service creator is not the table owner. Option B is incorrect; Snowflake recommends using a dedicated warehouse no larger than MEDIUM for each service, as larger warehouses do not necessarily increase performance for index materialization. Option C is incorrect because columns in the ATTRIBUTES field must be included in the source query. Option D is correct as Cortex Search provides a low-latency, high-quality hybrid (vector and keyword) search engine that automatically manages embedding, infrastructure, and search quality tuning. Option E is incorrect because Cortex Search Services currently do not support replication or cloning.

NEW QUESTION # 97

.....

GES-C01 test questions have a mock examination system with a timing function, which provides you with the same examination environment as the real exam. Although some of the hard copy materials contain mock examination papers, they do not have the

DOWNLOAD the newest Lead2Passed GES-C01 PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=179AmbQFyfOk2lMHbrUBG1MvixFZT6xxP>