

Snowflake GES-C01資格模擬 & GES-C01日本語受験教科書

Snowflake GES-C01 Exam

SnowPro® Specialty: Gen AI Certification Exam

<https://www.passquestion.com/ges-c01.html>



Pass GES-C01 Exam with PassQuestion GES-C01 questions and answers in the first attempt.

<https://www.passquestion.com/>

1 / 23

2026年JPNTTestの最新GES-C01 PDFダンプおよびGES-C01試験エンジンの無料共有：<https://drive.google.com/open?id=1xntpDWWY3mDAovIYS90djY8oa9khQsK9>

お客様に最も信頼性の高いバックアップを提供するという信念から当社のGES-C01試験問題を作成し、優れた結果により、試験受験者の機能に対する心を捉えました。GES-C01練習資料は、3つのバージョンに分類できます。これらのバージョンの使用はすべて、彼らに受け入れられています。これらのバージョンのGES-C01模擬練習には大きな格差はありませんが、能力を強化し、レビュープロセスをスピードアップして試験に関する知識を習得するのに役立ちます。そのため、レビュープロセスは妨げられません。

GES-C01試験ダンプは効率的かつ先鋭的に設計されているため、ユーザーはセクションを完了した後、タイムリーに学習効果を確認できます。GES-C01クイズガイドの成功率に関する優れた実践は、知識を習得したことを完全に示しているわけではありません。したがって、GES-C01テスト資料は、ユーザーが学習コンテンツをできるだけ多く統合することを可能にします、しかしそれは良い統合知識の結果を達成することができます。さらに重要なことは、GES-C01試験に合格し、夢のGES-C01認定を取得できることです。

>> Snowflake GES-C01資格模擬 <<

GES-C01試験の準備方法 | 最新のGES-C01資格模擬試験 | 実際的な SnowPro® Specialty: Gen AI Certification Exam日本語受験教科書

JPNTestは、このような効率的な学習計画を設計して、今後の開発のために効率の高い学習態度を構築できるようにすることを期待しています。私たちのGES-C01研究急流は、あなたが学生や事務員、緑の手、または長年の経験のあるスタッフであっても、すべての候補者に対応します。したがって、GES-C01試験に合格できるかどうかを心配する必要はありません。当社の技術力で成功することが保証されているからです。GES-C01試験問題の言語はわかりやすく、GES-C01学習ガイドの合格率は99%~100%です。

Snowflake SnowPro® Specialty: Gen AI Certification Exam 認定 GES-C01 試験問題 (Q89-Q94):

質問 # 89

A development team is preparing to deploy a new Retrieval-Augmented Generation (RAG) application written in Python. They intend to use Snowflake AI Observability 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 A
- B. Option E
- C. Option C
- D. Option D
- E. Option B

正解: A、B、C、E

解説:

質問 # 90

A data processing team is using Snowflake Document AI to extract data from incoming supplier invoices. They observe that many documents are failing to process, and successful extractions are taking longer than expected, leading to increased costs. Upon investigation, they find error messages such as

. Additionally, their 'X-LARGE' virtual warehouse is constantly active, contributing to higher-than-anticipated bills. Which two of the following actions are essential steps to troubleshoot and address the root causes of these processing errors and optimize their Document AI pipeline?

- A. Increase the 'max_tokenS' parameter within the '!PREDICT' function options to accommodate longer document responses from the model.
- B. Scale down the virtual warehouse to 'X-SMALL' or 'SMALL' size, as larger warehouses do not increase Document AI query processing speed and incur unnecessary costs.
- C. Configure the internal stage used for storing invoices with 'ENCRYPTION = (TYPE = 'SNOWFLAKE_SSE')'.
- D. Redefine extraction questions to be more generic and encompassing, reducing the number of distinct questions needed per document.
- E. Implement a pre-processing step to split documents exceeding 125 pages or 50 MB into smaller, compliant files before loading to the stage.

正解: C、E

解説:

The error messages 'Document has too many pages. Actual: 130. Maximum: 125.' and 'File exceeds maximum size. Actual: 54096026 bytes. Maximum: 50000000 bytes.' directly indicate that the documents do not meet Document AI's input requirements, which specify a maximum of 125 pages and 50 MB file size. Therefore, implementing a pre-processing step to split or resize these documents is an essential solution (Option B). The error 'cannot identify image file <_io.BytesIO object at 0x...>' is a known issue that occurs when an internal stage used for Document AI is not configured with 'SNOWFLAKE_SSE' encryption. Correctly configuring the stage with this encryption type is crucial for resolving this processing error (Option D). Option A, while addressing cost optimization, is not a root cause of the 'processing errors' themselves, although it is a best practice for cost governance as larger warehouses do not increase Document AI query processing speed. Option C is incorrect; best practices for question optimization suggest being specific, not generic. Option E is incorrect as 'max_tokenS' relates to the length of the model's output, not the input document's size or page limits.

質問 # 91

A security auditor needs to access and analyze logs generated by Snowflake AI Observability for compliance auditing and to track

the activity of generative AI applications. They need to understand how to reliably query this data and its temporal characteristics within Snowflake. Which of the following statements accurately describes the access and characteristics of this logged data?

- A. Access to these detailed event tables is implicitly granted to roles holding the SNOWFLAKE.CORTEX_USER database role and the AI_OBSERVABILITY_EVENTS_LOOKUP application role.
- B. The logs are automatically purged after 7 days of being recorded, requiring a separate process for long-term data retention.
- C. Logs are exclusively available for analysis through pre-built dashboards in Snowsight and cannot be accessed via direct SQL queries.
- D. Logged data from AI Observability's event tables becomes visible within a small latency, typically 1-2 minutes, after a request is made.
- E. Detailed request and response bodies, along with the generated SQL, are stored and can be directly queried using standard SQL.

正解: A、D、E

解説:

Snowflake AI Observability features logging of application traces and Cortex Analyst logs requests to an event table in the Snowflake database. There is a small latency of '1-2 minutes' before these logged requests are visible, making option A correct. The logs include detailed information such as 'Generated SQL' and 'Request and response bodies', which are stored and can be queried, validating option C. The necessary roles for AI Observability, including 'SNOWFLAKE.CORTEX_USER and EVENTS_LOOKUP', are required for creating and executing runs, which implies they grant access to the generated logs for monitoring, making option D correct. Option B is incorrect as the sources do not mention an automatic 7-day purge for these logs. Option E is incorrect because the documentation includes a subheading 'Querying logs with SQL' for Cortex Analyst administrator monitoring, indicating that direct SQL access is supported.

質問 # 92

A data engineering team is setting up an automated pipeline in Snowflake to process call center transcripts. These transcripts, once loaded into a raw table, need to be enriched by extracting specific entities like the customer's name, the primary issue reported, and the proposed resolution. The extracted data must be stored in a structured JSON format in a processed table. The pipeline leverages a SQL task that processes new records from a stream. Which of the following SQL snippets and approaches, utilizing Snowflake Cortex LLM functions, would most effectively extract this information and guarantee a structured JSON output for each transcript?

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

正解: B

解説:

To guarantee a structured JSON output for entity extraction, (the updated version of 'COMPLETE()') with the response_format argument and a specified JSON schema is the most effective approach. This mechanism enforces that the LLM's output strictly conforms to the predefined structure, including data types and required fields, significantly reducing the need for post-processing and improving data quality within the pipeline. Option A requires multiple calls and manual JSON assembly, which is less efficient. Option B relies on the LLM's 'natural ability' to generate JSON, which might not be consistently structured without explicit 'response_format'. Option D uses , which is for generating summaries, not structured entity extraction. Option E involves external LLM API calls and Python UDFs, which, while possible, is less direct than using native 'AI_COMPLETE structured outputs within a SQL pipeline in Snowflake Cortex for this specific goal.

質問 # 93

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 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.
- B. 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.
- C. The Cortex Search Service can effectively be used as a RAG engine for LLM chatbots by leveraging semantic search capabilities to provide customized and contextualized responses from the text data.
- D. 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.
- 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.

正解: A、B、C

解説:

Option A is correct because change tracking is required for the Cortex Search Service to monitor updates to the base table. 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. Options D and E are correct as Cortex Search provides low-latency, high-quality hybrid (vector and keyword) search, handling underlying complexities, and is primarily used as a RAG engine for LLM chatbots leveraging semantic search.

質問 #94

.....

最も専門的な専門家によって編集された当社のSnowflake練習資料は、成功のために高品質で正確なGES-C01練習資料を提供します。これまで、Snowflake試験トレントをサポートする世界中の何万人ものお客様がいます。GES-C01学習教材に不慣れな場合は、参考のために無料のデモをダウンロードしてください。また、一部の未学習の試験受験者には、Snowflake実践教材で必要事項をすぐにマスターできます。

GES-C01日本語受験教科書: <https://www.jpntest.com/shiken/GES-C01-mondaishu>

テストGES-C01認定に関する最も重要な情報を収集し、業界の上級専門家および認定講師および著者によって作成およびコンパイルされた新しい知識ポイントを補足します、ワークロードに圧倒され、息を吸うことができない場合、GES-C01準備トレントを選択してみませんか、業界での先駆者として、我々の目的はSnowflake GES-C01試験に参加する皆様に助けを提供することです、Snowflake GES-C01資格模擬 また、不明な点がある場合は、他の人に簡単に質問したり話したりできます、Snowflake GES-C01資格模擬 もちろんそんなことはないです、JPNTTestのSnowflakeのGES-C01試験トレーニング資料はIT認証試験を受ける全ての受験生が試験に合格することを助けるもので、受験生からの良い評価をたくさんもらいました。

博士はそこにひとりで住んでいる、甲斐のファスナーが下がったと思ったら、あっけないくらい素早く下着ごとズボンを下げられている、テストGES-C01認定に関する最も重要な情報を収集し、業界の上級専門家および認定講師および著者によって作成およびコンパイルされた新しい知識ポイントを補足します。

有効的な GES-C01資格模擬 & 保証する Snowflake GES-C01 公認された試験の成功GES-C01日本語受験教科書

ワークロードに圧倒され、息を吸うことができない場合、GES-C01準備トレントを選択してみませんか、業界での先駆者として、我々の目的はSnowflake GES-C01試験に参加する皆様に助けを提供することです、また、不明な点がある場合は、他の人に簡単に質問したり話したりできます。

もちろんそんなことはないです。

- GES-C01認定デベロッパー GES-C01試験解説 GES-C01認定資格 検索するだけで www.xhs1991.com から {GES-C01} を無料でダウンロードGES-C01認定資格
- 最新の高効率のSnowflake GES-C01試験問題集、100%合格できます! 《www.goshiken.com》で GES-C01 を検索して、無料で簡単にダウンロードできますGES-C01赤本勉強
- 最新の高効率のSnowflake GES-C01試験問題集、100%合格できます! jp.fast2test.com には無料の GES-C01 問題集がありますGES-C01テキスト
- 最新の高効率のSnowflake GES-C01試験問題集、100%合格できます! GES-C01 の試験問題は (www.goshiken.com) で無料配信中GES-C01入門知識

