

# 実地的なGES-C01最新関連参考書 &合格スムーズ GES-C01的中合格問題集 |最新のGES-C01模擬体験



無料でクラウドストレージから最新のPass4Test GES-C01 PDFダンプをダウンロードする：<https://drive.google.com/open?id=1TZfLK67iwppFUGANfvLybryGvJc8TxgA>

SnowflakePDFバージョン、PCバージョン、APPオンラインバージョンなど、3つの異なるバージョンのSnowPro® Specialty: Gen AI Certification Exam prepトレントを選択できます。異なるバージョンには独自の利点とユーザー数があります。PDFバージョンの機能をご紹介します。Pass4TestのGES-C01試験トレントのPDFは、主にデモの利便性のために、若者の中で最も一般的なバージョンであることに疑いの余地はありません。SnowPro® Specialty: Gen AI Certification Exam印刷してメモを取ることができます。

SnowflakeのGES-C01試験の合格書は君の仕事の上で更に一步の昇進と生活条件の向上を助けられて、大きな財産に相当します。SnowflakeのGES-C01認定試験はIT専門知識のレベルの考察として、とっても重要な地位になりつつます。Pass4Testは最も正確なSnowflakeのGES-C01試験資料を追求しています。

>> GES-C01最新関連参考書 <<

## 高品質なGES-C01最新関連参考書試験-試験の準備方法-効率的なGES-C01的中合格問題集

GES-C01試験ダンプは効率の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 SnowPro® Specialty: Gen AI Certification Exam 認定 GES-C01 試験問題 (Q247-Q252):

### 質問 # 247

A data science team is deploying a custom real-time inference service for a fine-tuned LLM using Snowpark Container Services (SPCS). They have a Docker image in their Snowflake image repository. They need to define the service using a YAML specification file. Which of the following are "essential" components or configurations that must be included in the 'spec.yaml' file for a long-running service that uses this image, custom environment variables, and requires external access?

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

正解: B、D、E

解説:

Option A is correct because the 'container' block, specifying the 'name' of the container and the Docker 'image' path from the Snowflake registry, is fundamental for defining any containerized service in SPCS. Option B is correct as the 'endpoints' block, explicitly defining a 'name', 'port', and setting 'public: true', is essential for a long-running service that requires external access via a web browser or API. Option C is incorrect because the 'compute\_pool' is specified during the 'CREATE SERVICE SQL' command, not within the 'spec.yaml' file itself. Option D is incorrect as a 'volumes' block, while potentially useful for persistent storage, is not listed as an essential component in the provided spec.yaml examples for basic service deployment. Option E is correct as the 'env' block within a container definition is used to set custom environment variables, which is explicitly shown in the Jupyterlab example for configuring the application within the container.

#### 質問 # 248

A Snowflake administrator is tasked with monitoring the efficiency and cost-effectiveness of their Cortex Analyst deployments. They need to identify if certain semantic models are generating a high volume of failed or expensive queries. Which of the following approaches or statements are crucial for effectively monitoring and identifying issues with Cortex Analyst usage and associated costs?

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

正解: A、C、E

解説:

Option A is correct because Cortex Analyst logs requests to an event table, and the function can be used to query these logs, which include generated SQL and errors/warnings, helping identify issues. Option C is correct as the 'METERIN\_DAILY\_HISTORY' view tracks daily credit usage for services, including 'AI\_SERVICES', which encompasses Cortex Analyst. Option E is correct as there is a reported 1-2 minute lag for Cortex Analyst requests to become visible in the event table logs. Option B is incorrect because 'CORTEX\_DOCUMENT\_PROCESSING\_USAGE\_HISTORY' specifically displays Document AI processing function activity, not all AI Services or Cortex Analyst. Option D is incorrect because Cortex Analyst costs are based on messages, not tokens, so 'prompt\_tokens' and 'completion\_tokens' would not be relevant for direct cost assessment in this context.

#### 質問 # 249

A financial institution wants to automate the extraction of key entities (e.g., invoice number, total amount, list of invoice items) from incoming PDF financial statements into a structured JSON format within their Snowflake data pipeline. The extracted data must conform to a specified JSON schema for seamless downstream integration. Which Snowflake Cortex capabilities, when combined, can best achieve this data augmentation and ensure schema adherence in a continuous processing pipeline?

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

正解: B、D

解説:

#### 質問 # 250

A data engineering team is planning to build a real-time data pipeline using Snowflake's dynamic tables to process incoming log data. They want to use SNOWFLAKE.CORTEX.EXTRACT\_ANSWER to pull out specific error codes and timestamps from log entries. They are also mindful of the operational costs. Which of the following statements accurately describes limitations or cost considerations for using SNOWFLAKE.CORTEX.EXTRACT\_ANSWER in this scenario?

- A. The billing for EXTRACT\_ANSWER is based on the combined token count of both the source\_document (log entry) and the question used for extraction.

- B. The `EXTRACT_ANSWER` function does not support dynamic tables, which will prevent its direct use in this pipeline design.
- C. If a `source_document` input exceeds the model's token limit, `EXTRACT_ANSWER` will automatically truncate the text, potentially impacting results but avoiding an error.
- D. To optimize performance and reduce cost for large-scale log processing, it is recommended to execute `EXTRACT_ANSWER` queries on a Snowpark-optimized warehouse of at least 'LARGE' size.
- E. `EXTRACT_ANSWER` is explicitly designed for multi-language extraction, making it suitable for logs that might contain various languages without affecting cost.

正解: A、B

解説:

Option A is correct because Snowflake Cortex functions, including `EXTRACT_ANSWER`, do not support dynamic tables. Option B is correct because for `EXTRACT_ANSWER`, the number of billable tokens is the sum of the number of tokens in the (source document) and 'question' fields. Option C is incorrect; Snowflake recommends executing queries that call Cortex AISQL functions with a smaller warehouse (no larger than MEDIUM) as larger warehouses do not increase performance. Option D is incorrect; inputs that exceed the model's token limit result in an error, rather than automatic truncation. Option E is incorrect; while the newer 'AI\_EXTRACT' supports multiple languages, the documentation for "EXTRACT\_ANSWER" does not explicitly state multi-language support and generally refers to plain English text. Costs are incurred per token regardless of language.

#### 質問 # 251

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. 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.
- 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. 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.
- 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. 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.

正解: B、D、E

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 # 252

.....

今多くのIT技術会社は職員がSnowflakeのGES-C01資格認定を持つのを要求します。SnowflakeのGES-C01試験に合格するのは必要なこととなります。速く試験に合格して資格認証を取得したいなら、我々Pass4TestのGES-C01問題集を使ってください。弊社はあなたに相応しくて品質高いGES-C01問題集を提供します。また、あなたの持っている問題集は一年間の無料更新を得られています。あなたは十分の時間でGES-C01試験を準備することができます。

GES-C01的中合格問題集: <https://www.pass4test.jp/GES-C01.html>

