

GES-C01 Probesfragen & GES-C01 Fragen Und Antworten

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

Übrigens, Sie können die vollständige Version der ExamFragen GES-C01 Prüfungsfragen aus dem Cloud-Speicher herunterladen: <https://drive.google.com/open?id=1bu1xjjLeqbvMXPdA132eeqxZCbSAhnA9>

Die Snowflake GES-C01 Zertifizierungsprüfung ist eigentlich eine Prüfung für die Technik-Experten. Die Snowflake GES-C01 Zertifizierungsprüfung kann den IT-Fachleuten helfen, eine bessere Berufskarriere zu haben. So können Sie dem Staat und Unternehmen große Gewinne bringen und die wirtschaftliche Entwicklung unseres Landes fördern. Wenn alle Fachleute das machen, ist unser Staat sicher reicher geworden. Unsere Schulungsunterlagen zur Snowflake GES-C01 Zertifizierungsprüfung können dieses Ziel der IT-Fachleute erreichen. Wir versprechen, dass Sie 100% die Prüfung bestehen können. Wenn Sie lange denken, ist es besser entschlossen eine Entscheidung zu treffen, die Schulungsunterlagen zur Snowflake GES-C01 Zertifizierungsprüfung von ExamFragen zu kaufen.

Die Fragenkataloge zur Snowflake GES-C01 Prüfung von ExamFragen sind die besten im Vergleich zu den anderen Materialien. Wenn Sie Fragenkataloge suchen, wählen Sie doch die Fragenkataloge zur Snowflake GES-C01 Prüfung von ExamFragen. Und Sie würden viel davon profitieren. Sonst würden Sie bereuen.

>> GES-C01 Probesfragen <<

Echte und neueste GES-C01 Fragen und Antworten der Snowflake GES-C01 Zertifizierungsprüfung

Die Schulungsunterlagen zur Snowflake GES-C01 Zertifizierungsprüfung von unserem ExamFragen haben präzise und flächendeckende Inhalte. Diese Lernhilfe sind geeignet für Sie und werden die notwendigsten Ausbildungsmaterialien sein, wenn Sie die Zertifizierungsprüfung bestehen möchten. Hier versprechen wir, dass Sie einjährige Aktualisierung kostenlos genießen können, nachdem Sie unsere Schulungsunterlagen zur Snowflake GES-C01 Zertifizierungsprüfung gekauft haben. Wenn Sie die GES-C01 Prüfung nicht bestehen oder unsere Fragenkataloge irgend ein Qualitätsproblem haben, geben wir Ihnen eine bedingungslose volle Rückerstattung.

Snowflake SnowPro® Specialty: Gen AI Certification Exam GES-C01 Prüfungsfragen mit Lösungen (Q279-Q284):

279. Frage

A data engineering team needs to establish an automated pipeline in Snowflake to continuously extract 'contract_id' and effective_date' from new PDF contract documents uploaded to an internal stage named They have a pre-trained Document AI model named 'contract_processor'. Which of the following sets of SQL commands correctly configures the necessary Snowflake objects for this automated processing pipeline, including handling file access and initial data loading?

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

Antwort: C

Begründung:

280. Frage

A data engineering team is developing a Cortex Analyst semantic model in YAML for an e-commerce platform. They need to ensure high accuracy for common queries, improve literal matching for product names, and understand limitations for supported data types. Which of the following statements correctly describe aspects of semantic model configuration or capabilities for these requirements?

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

Antwort: B,C

Begründung:

Option A is correct. Cortex Search Services can be integrated into a dimension's definition (using the field with 'service' and fields) to improve literal matching by performing semantic search over the underlying column. This is specifically useful for 'fuzzy' searches of product names. Option C is correct. The 'verified_queries' section allows pre-defining accurate SQL queries for specific natural language questions. Setting 'use_as_onboarding_question true' for entries ensures these queries are used when relevant and presented as suggested questions to users, or as onboarding questions. Option B is incorrect because 'VARIANT', 'OBJECT', 'GEOGRAPHY', and 'ARRAY' data types are currently not supported for dimension, fact, or metric columns in a semantic model. Option D is incorrect; the 'sample_values' field is recommended for dimensions with relatively low-cardinality (approximately 1-10 distinct values) to aid in semantic search for literals, not for high-cardinality dimensions like millions of customer IDs. Option E is incorrect because a 'base_table' in a semantic model must refer to a physical database table or a view, not directly to a stage location.

281. Frage

A developer is building a prototype Gen AI application using the Snowflake Cortex LLM REST API. During testing, they frequently encounter HTTP 429 'Too Many Requests' errors, particularly when rapidly invoking the 'COMPLETE' endpoint. Which of the following statements accurately describes the usage quotas and troubleshooting steps for such rate limit events in the Cortex REST API?

- A. The HTTP 429 response indicates that the daily credit consumption for Snowflake Cortex AISQL usage has exceeded the 10-credit limit for on-demand accounts without a valid payment method, requiring conversion to a paid account.
- B. If the usage is below the RPM limit but still triggers a 429, the developer should review the token usage rate as TPM limits might be exceeded.
- C. Rate limits are model-specific, and exceeding either the Tokens Processed per Minute (TPM) or Requests per Minute (RPM) limits will result in an HTTP 429 response code.
- D. Cortex REST API rate limits are primarily based on Requests per Minute (RPM) and do not consider the total Tokens Processed per Minute (TPM).
- E. To resolve rate limiting, the developer should increase the virtual warehouse size for their API requests, as larger warehouses automatically increase API quotas.

Antwort: B,C

Begründung:

Option C is correct because the Cortex REST API implements rate limits for both Tokens Processed per Minute (TPM) and Requests per Minute (RPM), and offending either of these limits will result in a 429 response code. Option D is correct as the sources explicitly state that if REST API usage is below the RPM limit but still receives a 429, the user should double-check the token usage rate, implying that TPM limits could be the cause. Option A is incorrect because while on-demand Snowflake accounts without a valid payment method are limited to 10 credits per day, an HTTP 429 specifically refers to rate limiting due to high usage, not a daily credit budget exceeding for the account type. Option B is incorrect as Cortex REST API rate limits consider both TPM and RPM. Option E is incorrect because Cortex REST API requests do not require a warehouse, and increasing warehouse size does not directly increase API quotas or resolve rate limiting for the REST API.

282. Frage

A data analytics team aims to enhance their understanding of customer feedback stored in a Snowflake table called CUSTOMER_FEEDBACK. This table has a REVIEW_TEXT column containing raw customer comments and a CUSTOMER_SEGMENT column. The team wants to classify each review into predefined categories and then generate a concise summary of all reviews for each customer segment. Which of the following Snowflake Cortex AI functions and approaches should they use?

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

Antwort: A,B,C

Begründung:

283. Frage

A data operations team is attempting to scale up their Document AI automated pipeline. They are using a Snowflake Task to process a large volume of daily scanned invoices and receipts, which are stored in an internal stage 'financial_docs stage'. The current processing involves documents that are frequently around 75 MB each, and often there are batches exceeding 1,200 documents in a single day. The pipeline is failing consistently. Which of the following factors could be contributing to the failures in this Document AI automated pipeline? (Select all that apply.)

- A. The account role executing the task lacks the 'SNOWFLAKCORTEX USER database role, which is a prerequisite for Document AI functions.
- B. The individual document size (75 MB) exceeds the maximum supported file size for Document AI.
- C. The was created without specifying 'ENCRYPTION = (TYPE = 'SNOWFLAKE SSE')'.
- D. The Snowflake Task is configured as a 'SERVERLESS TASK, which is not supported by Document AI.
- E. The number of documents in a single daily batch (exceeding 1,200) surpasses Document AI's processing limit per query.

Antwort: B,C,D,E

Begründung:

