

GES-C01 Testing Engine - GES-C01 Online Praxisprüfung



Das IT-Expertenteam von Zertprüfung haben eine kurzfristige Schulungsmethode nach ihren Kenntnissen und Erfahrungen bearbeitet. Diese Dumps könne Ihnen effektiv helfen, in kurzer Zeit den erwarteten Effekt zu erzielen, besonders für diejenigen, die arbeiten und zugleich lernen. Zertprüfung kann Ihnen viel Zeit und Energir ersparen. Wählen Sie Zertprüfung und Sie werden Ihre wünschsten Schulungsmaterialien zur Snowflake GES-C01 Zertifizierungsprüfung bekommen.

Manchmal muss man mit große Menge von Prüfungsaufgaben üben, um eine wichtige Prüfung zu bestehen. Die Snowflake GES-C01 von uns hat diese Forderung gut erfüllt. Und mit den fachlichen Erklärungen können Sie besser die Antworten verstehen. Die Demo der Snowflake GES-C01 von unterschiedlichen Versionen werden von uns gratis angeboten. Probieren Sie mal und wählen Sie die geeignete Version für Sie! Mit unserer gemeinsamen Arbeit werden Sie bestimmt die Snowflake GES-C01 Prüfung erfolgreich bestehen!

>> GES-C01 Testing Engine <<

Hilfsreiche Prüfungsunterlagen verwirklicht Ihren Wunsch nach der Zertifikat der SnowPro® Specialty: Gen AI Certification Exam

100% Garantie SnowPro® Specialty: Gen AI Certification Exam Prüfungserfolg, Wenn Sie Zertprüfung GES-C01 Prüfung wählen Snowflake Zertprüfung Test Engine ist das perfekte Werkzeug, um auf die Zertifizierungsprüfung vorbereiten. Erfolg kommt einfach,

wenn Sie bereiten mit Hilfe von Original bis zu SnowPro® Specialty: Gen AI Certification Exam Produkte mit Zertprüfung Datum. Wie ein seltener Fall, wenn Sie es versäumen, diese Prüfung geben wir Ihnen eine volle Rückerstattung Ihres Einkaufs passieren.

Snowflake SnowPro® Specialty: Gen AI Certification Exam GES-C01 Prüfungsfragen mit Lösungen (Q186-Q191):

186. Frage

A data platform architect is evaluating the integration of 'SNOWFLAKE.CORTEX.TRANSLATE' into several automated data pipelines. One pipeline involves real-time translation of messages for a chat application, while another is for batch processing of archived documents. The architect is considering various Snowflake features for orchestration and deployment. Which of the following considerations about 'SNOWFLAKE.CORTEX.TRANSLATE' is accurate?

- A. The Snowflake managed model used by the 'TRANSLATE' function has a context window of 4,096 tokens, meaning texts longer than this will be truncated before translation.
- B. When using 'TRANSLATE' within a Snowpark Python User-Defined Function (UDF), the raw text data must be explicitly moved out of Snowflake's network boundary to the underlying LLM service for translation.
- C. The 'TRANSLATE' function can be seamlessly integrated into a dynamic table's 'SELECT' statement to provide continuous, automated translation with minimal configuration.
- D. To manage potential failures in a production pipeline, 'SNOWFLAKE.CORTEX.TRANSLATE' should be wrapped in 'TRY COMPLETE' for robust error handling, returning NULL on failure instead of an error.
- E. If 'TRANSLATE' is not natively available in the account's primary Snowflake region, cross-region inference cannot be enabled, thus preventing its use.

Antwort: A

Begründung:

Option D is correct. The Snowflake managed model from the 'SUMMARIZE' function has a context window of 32,000 tokens, but for TRANSLATE the context window is 4,096 tokens. Input text exceeding this limit would be truncated. Option A is incorrect because Snowflake Cortex functions, including 'TRANSLATE', do not support dynamic tables. Option B is incorrect; when Snowflake Cortex LLM functions (like TRANSLATE) are called on Snowflake data, the data never actually leaves Snowflake's network boundary. Option C is incorrect; cross-region inference can be used if 'TRANSLATE' is not natively available in a region. Option E is incorrect because 'TRY_COMPLETE' is a specific function that performs the same operation as 'COMPLETE' but returns 'NULL' on error. While 'TRY_COMPLETE' handles errors for 'COMPLETE', it is not a general wrapper for any Cortex function like 'TRANSLATE'. There is no indication that 'TRANSLATE' should be wrapped in 'TRY_COMPLETE' for error handling; instead, standard error handling for SQL queries would apply, or if there were a 'TRY_TRANSLATE' function (which isn't mentioned). The prompt mentions using 'TRY_COMPLETE' with a 'COMPLETE' like function.

187. Frage

A financial institution is deploying a sentiment analysis application that uses Snowflake Cortex 'SENTIMENT' and 'COMPLETE' functions, with different LLMs, for processing customer feedback. They are using AI Observability (Public Preview) to compare the cost-efficiency of using 'mistral-7b' versus 'claude-3-5-sonnet' as LLM judges for evaluation metrics, and also tracking the overall cost of their AI Observability usage. Which statements accurately reflect the cost implications and monitoring tools for this scenario?

- ☐ AI Observability incurs charges for LLM judges (e.g., 'mistral-7b', 'claude-3-5-sonnet') invoked via 'COMPLETE (SNOWFLAKE.CORTEX)' calls to compute evaluation metrics, and these charges are based on 'tokens processed'.
- ☐ The 'CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY' view should be used to monitor the credit consumption of the LLM judges specifically, as it tracks all Cortex function calls.
- ☐ Comparing models with vastly different context windows (e.g., 'mistral-7b' at 32k tokens vs. 'claude-3-5-sonnet' at 200k tokens) using AI Observability will not impact the billed 'tokens processed' if the actual prompt sizes are small and similar.
- ☐ In addition to LLM judge costs, warehouse charges are incurred for tasks managing evaluation runs and for queries used to compute evaluation metrics within AI Observability.
- ☐ The 'METERING_DAILY_HISTORY' view, filtering by 'SERVICE_TYPE ILIKE '%ai_services%', can provide an overview of daily credit usage for all AI services, including AI Observability LLM judge activity and associated warehouse costs.

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

Antwort: A,C,E

Begründung:

Option A is correct because AI Observability utilizes LLM judges (such as 'mistral-7b' or 'claude-3-5-sonnet') through 'COMPLETE (SNOWFLAKE.CORTEX)' function calls to compute evaluation metrics, and these calls incur charges based on the 'tokens processed'. Option D is correct as, beyond LLM judge costs, AI Observability also incurs warehouse charges for managing evaluation runs and for queries that compute evaluation metrics. Option E is correct because the view, with a filter for 'SERVICE _ TYPE ILIKE', provides a comprehensive daily credit usage report for all AI services, which would include AI Observability's components. Option B is incorrect; the view is specifically for Document AI processing functions like '!PREDICT and 'AI_EXTRACT', not for general LLM judge usage in AI Observability. The view is more appropriate for tracking individual Cortex function calls. Option C is incorrect because while prompt sizes might be similar, the pricing for different LLMs (e.g., 'mistral-7b' at 0.12 credits per million tokens vs. 'claude-3-5-sonnet' at 2.55 credits per million tokens for AI Complete) will still result in different billed amounts due to varying per-token costs, even if the number of tokens is the same.

188. Frage

A marketing team is analyzing social media comments using Snowflake and wants to categorize them into predefined campaign sentiments (e.g., 'Positive Campaign Engagement', 'Negative Campaign Feedback', 'Neutral Discussion'). They decide to use the SNOWFLAKE.CORTEX.CLASSIFY TEXT function for this task. Which of the following statements about its usage are correct?

- A. The input string to CLASSIFY_TEXT is case-insensitive, meaning 'Great product!' and 'great product!' will yield identical classification results due to automatic normalization.
- B. The argument must contain exactly two string values for effective binary classification, otherwise an error is returned.
- C. If the input text exceeds a model-specific token limit, CLASSIFY_TEXT will automatically truncate the text before processing without raising an error.
- D. To provide more context and potentially improve classification accuracy, categories within the can be defined as SQL objects, including 'description' and 'examples' fields.
- E. CLASSIFY_TEXT can return a JSON object with a 'label' field, where the value of this field indicates the classified category of the input text.

Antwort: D,E

Begründung:

Option A is incorrect because both the input string and categories for 'CLASSIFY TEXT' are case sensitive, meaning different capitalizations can lead to different results. Option B is incorrect because the argument must contain at least two and at most 100 unique categories. Option C is incorrect as returns an OBJECT (VARIANT) whose 'label' field specifies the category to which the input prompt belongs. Option D is correct because categories can be simple strings or SQL objects, allowing for a description and examples to be provided, which can improve accuracy. Option E is incorrect because the documentation for 'CLASSIFY _ TEXT' does not mention automatic truncation of input text based on a token limit, although LLMs typically have context windows. The source only mentions that for non- plain English text, results may not be what you expect, not that the input would be truncated.

189. Frage

A Gen AI Engineer is configuring a new semantic model for Cortex Analyst to process customer feedback. The goal is to ensure that when a user asks for sentiment analysis, the generated SQL queries always include an aggregation by a dimension and present the results as a percentage. The engineer plans to use custom instructions for this purpose. Which of the following details about is true and crucial for successful implementation?

☐ The 'custom_instructions' field directly accepts SQL snippets, such as

```
GROUP BY customer_segment
```

and

```
AVG(sentiment) * 100
```

, to embed into the generated queries.

☐ The 'custom_instructions' are specified in a separate, linked Python file, allowing for complex procedural logic to modify the SQL generation process dynamically.

☐ The 'custom_instructions' are provided in natural language within the semantic model YAML to guide the LLM on how to interpret user intent and structure the SQL response (e.g., 'Always group sentiment analysis by customer segment and display results as a percentage').

☐ There is a strict character limit of 50 words for the 'custom_instructions' field, similar to the 'task_description' in 'CLASSIFY_TEXT'.

☐ The 'custom_instructions' are user-specific and are managed at the individual user level in Snowsight, rather than being part of the shared semantic model definition.

- A. Option C

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

Antwort: A

Begründung:

Custom instructions in Cortex Analyst provide unique business context to the LLM to control SQL query generation. These instructions are provided in natural language within the semantic model YAML. This means the engineer should describe the desired behavior (grouping by 'customer_segment' and presenting as a percentage) in plain English for the LLM to interpret and apply, making option C correct. Option A is incorrect because 'custom_instructions' guide the LLM's *generation* process, not directly inject SQL snippets. Option B is incorrect as custom instructions are part of the YAML, not a separate Python file. Option D is incorrect; while Copilot's custom instructions have a 2,000 character limit, the source does not specify such a limit for Cortex Analyst's semantic model 'custom_instructions', and the 'task_description' for 'CLASSIFY_TEXT' is limited to about 50 words. Option E is incorrect; the 'custom_instructions' in the semantic model are part of the shared model definition, not user-specific in the way Snowflake Copilot's custom instructions are.

190. Frage

An ML Engineer has developed a custom PyTorch model for image processing that requires GPU acceleration and specific PyPI packages ('torch', 'torchvision'). They want to deploy it as a service on Snowpark Container Services (SPCS) using the Snowflake Model Registry. Which of the following statements are true regarding the deployment of this model to SPCS and its requirements? (Select all that apply.)

- ☐ The model should be logged using `conda_dependencies` for PyTorch, explicitly specifying the 'conda-forge' channel if custom versions are needed, and `pip_requirements` should be entirely avoided to prevent potential conflicts. Example: `conda_dependencies=["conda-forge:pytorch", "torchvision"]`.
- ☐ A compute pool with an `INSTANCE_FAMILY` starting with 'GPU' must be created using `CREATE COMPUTE POOL` and specified via the 'service_compute_pool' argument during service creation to leverage GPU acceleration.
- ☐ Within the custom model's Python code, any internal memory directory for writing temporary files should be set to `/tmp/`, as it is a universally writeable and safe location on Snowflake's warehouse nodes and SPCS containers.
- ☐ When calling the 'create_service' method on the 'ModelVersion' object, the 'gpu_requests' argument must be set to a string specifying the number of GPUs (e.g., "1") to ensure the model runs on GPU resources.
- ☐ The 'relax_version' option must be explicitly set to 'True' during the 'log_model' call to ensure compatibility with various SPCS runtime environments, especially for common open-source packages.

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

Antwort: A,B,C

Begründung:

Statement A is incorrect. While Snowflake recommends using only 'conda_dependencies' or only 'pip_requirements' (not both) to avoid package conflicts, the scenario explicitly mentions PyPI packages. If using 'pip_requirements', all required packages should be listed there. The example incorrectly assumes 'torchvision' would necessarily be best sourced from Conda and dictates avoiding 'pip_requirements' entirely, which is an oversimplification of the recommendation. Statement B is correct. To utilize GPU acceleration in SPCS, a compute pool configured with a GPU instance family (e.g., *GPU must be created and then referenced by name in the 'service_compute_pool' argument when creating the service. Statement C is correct. Snowflake's warehouse nodes have restricted directory access, and '/tmp' is recommended as a safe and writeable location for models that need to write files during execution. This principle extends to SPCS containers. Statement D is correct. The 'create_service' method for deploying models to SPCS takes a `gpu_requests` argument, which specifies the number of GPUs to allocate to the service. Setting this (e.g., to '1') is crucial for ensuring the model runs on GPU hardware. Statement E is incorrect. The 'relax_version' option, which modifies version constraints, defaults to 'True' in 'log_model'. While often beneficial, it is not mandatory to explicitly set it to 'True' for every deployment scenario.

191. Frage

.....

Fühlen Sie sich schmerzvoll, wenn Sie so viele IT-Zertifizierungen und Zertifizierungsunterlagen sehen? Was sollen Sie machen? Welche Prüfung und welche Prüfungsunterlage sollen Sie wählen? Wir Zertpruefung können die geeignete Prüfungen für Sie wählen, wenn Sie wissen nicht, wie sich zu entscheiden. Sie können jetzt sehr populäre Snowflake GES-C01 Zertifizierungsprüfung wählen. Diese Zertifizierung hat viele Vorteile. Außerdem, wenn Sie sehr effektiv die Prüfung vorbereiten, können Sie sich für Snowflake GES-C01 Dumps von Zertpruefung entscheiden. Es ist die beste Methode für dich, diese Snowflake GES-C01 Prüfung einfach zu bestehen.

GES-C01 Online Praxisprüfung: https://www.zertpruefung.de/GES-C01_exam.html

Sie können in kurzer Zeit ihre IT-Fachkenntnisse ergänzen und sich gut auf die Snowflake GES-C01 Zertifizierungsprüfung vorbereiten, Sobald Sie die Materialien gekauft haben, brauchen Sie nur 20 bis 30 Stunden, mit GES-C01 Fragen & Antworten zu studieren, Snowflake GES-C01 Testing Engine Deshalb wollen wir Ihnen helfen, Ihre Angst und Stress zu beseitigen, Die GES-C01 Prüfung auf die geeignetste Weise vorbereiten.

Aber der andre glaubte, er weiche ihm aus, damit er ihm nicht sagen müsse, GES-C01 wie wenig ihm das Manuskript gefallen habe, Sie beteten im Morgengrauen, vor dem Frühstück, und knieten um das stille schwarze Becken.

Snowflake GES-C01 Fragen und Antworten, SnowPro® Specialty: Gen AI Certification Exam Prüfungsfragen

Sie können in kurzer Zeit ihre IT-Fachkenntnisse ergänzen und sich gut auf die Snowflake GES-C01 Zertifizierungsprüfung vorbereiten, Sobald Sie die Materialien gekauft haben, brauchen Sie nur 20 bis 30 Stunden, mit GES-C01 Fragen & Antworten zu studieren.

Deshalb wollen wir Ihnen helfen, Ihre Angst und Stress zu beseitigen, Die GES-C01 Prüfung auf die geeignetste Weise vorbereiten, Mit nur 20 bis 30 Stunden von konzentriertem Übungen können Sie die Snowflake GES-C01 bestehen.

- GES-C01 Schulungsangebot, GES-C01 Testing Engine, SnowPro® Specialty: Gen AI Certification Exam Trainingsunterlagen □ Suchen Sie jetzt auf ► www.zertsoft.com □ nach ➡ GES-C01 □ um den kostenlosen Download zu erhalten □GES-C01 Prüfung
- GES-C01 Torrent Anleitung - GES-C01 Studienführer - GES-C01 wirkliche Prüfung □ Öffnen Sie die Website ➡ www.itzert.com □ Suchen Sie { GES-C01 } Kostenloser Download □GES-C01 Demotesten
- GES-C01 Übungsfragen: SnowPro® Specialty: Gen AI Certification Exam - GES-C01 Dateien Prüfungsunterlagen □ Öffnen Sie ► www.deutschpruefung.com ◀ geben Sie ☀ GES-C01 □☀□ ein und erhalten Sie den kostenlosen Download □GES-C01 Zertifikatsfragen
- GES-C01 Testengine □ GES-C01 Tests □ GES-C01 Vorbereitung □ Suchen Sie auf der Webseite “ www.itzert.com ”nach □ GES-C01 □ und laden Sie es kostenlos herunter □GES-C01 Prüfung
- 100% Garantie GES-C01 Prüfungserfolg □ Geben Sie ► www.zertpruefung.de ◀ ein und suchen Sie nach kostenloser Download von ► GES-C01 ◀ □GES-C01 Prüfungsübungen
- GES-C01 Übungsfragen: SnowPro® Specialty: Gen AI Certification Exam - GES-C01 Dateien Prüfungsunterlagen □ Suchen Sie jetzt auf ☀ www.itzert.com □☀□ nach [GES-C01] und laden Sie es kostenlos herunter □GES-C01 Online Tests
- Die seit kurzem aktuellsten Snowflake GES-C01 Prüfungsinformationen, 100% Garantie für Ihren Erfolg in der Prüfungen! □ □ Suchen Sie auf► www.itzert.com ◀ nach kostenlosem Download von ☀ GES-C01 □☀□ □GES-C01 Prüfung
- GES-C01 Tests □ GES-C01 Prüfungsübungen □ GES-C01 Fragen Und Antworten □ Suchen Sie jetzt auf □ www.itzert.com □ nach ➡ GES-C01 □□□ um den kostenlosen Download zu erhalten □GES-C01 Zertifikatsfragen
- GES-C01 Fragen Und Antworten □ GES-C01 Zertifizierungsprüfung □ GES-C01 German □ Erhalten Sie den kostenlosen Download von ► GES-C01 □ mühelos über ➡ www.deutschpruefung.com □□□ □GES-C01 Fragen&Antworten
- GES-C01 Prüfungsfragen Prüfungsvorbereitungen 2025: SnowPro® Specialty: Gen AI Certification Exam - Zertifizierungsprüfung Snowflake GES-C01 in Deutsch Englisch pdf downloaden □ Suchen Sie jetzt auf► www.itzert.com ◀ nach □ GES-C01 □ und laden Sie es kostenlos herunter □GES-C01 Dumps Deutsch
- GES-C01 Übungsfragen: SnowPro® Specialty: Gen AI Certification Exam - GES-C01 Dateien Prüfungsunterlagen □ Suchen Sie jetzt auf 【 www.zertfragen.com 】 nach ⇒ GES-C01 ⇐ um den kostenlosen Download zu erhalten □GES-C01 Zertifikatsfragen
- learn.srkk.com, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, elearning.omegasystems.gr, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.excelentaapulum.ro, study.stcs.edu.np, lms.ait.edu.za, skillkaro.com, Disposable vapes