

GES-C01 Free Exam Questions - Trustworthy GES-C01 Source

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

P.S. Free 2026 Snowflake GES-C01 dumps are available on Google Drive shared by TestKingIT: <https://drive.google.com/open?id=1MpgxpPwtUH8emugOjhGhtW5W8flXBY39>

There are three versions of GES-C01 guide quiz. You can choose the most suitable version based on your own schedule. PC version, PDF version and APP version, these three versions of GES-C01 exam materials you can definitely find the right one for you. Also our staff will create a unique study plan for you. In order to allow you to study and digest the content of GES-C01 practice prep more efficiently, after purchasing, you must really absorb the content in order to pass the exam. GES-C01 guide quiz really wants you to learn something and achieve your goals.

For Snowflake aspirants wishing to clear the Snowflake test and become a SnowPro® Specialty: Gen AI Certification Exam certification holder, TestKingIT Snowflake GES-C01 practice material is an excellent resource. By preparing with TestKingIT actual Snowflake GES-C01 Exam Questions, you can take get success on first attempt and take an important step toward accelerating your career. Download updated GES-C01 exam questions today and start preparation.

>> GES-C01 Free Exam Questions <<

GES-C01 Quiz Braindumps - GES-C01 Pass-Sure torrent & GES-C01 Exam Torrent

We have gained high appraisal for the high quality GES-C01 guide question and considerate serves. All content is well approved by experts who are arduous and hardworking to offer help. They eliminate banal knowledge and exam questions out of our GES-C01 real materials and add new and essential parts into them. And they also fully analyzed your needs of GES-C01 exam dumps all the time. Our after sales services are also considerate. If you get any questions with our GES-C01 guide question, all helps are available upon request. Once you place your order this time, you will enjoy and experience comfortable and convenient services immediately. Besides, we do not break promise that once you fail the GES-C01 Exam, we will make up to you and relieve you of any loss. Providing with related documents, and we will give your money back. We have been always trying to figure out how to provide warranty service if customers have questions with our GES-C01 real materials. So all operations are conducted to help you pass the exam with efficiency.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q90-Q95):

NEW QUESTION # 90

A multi-national corporation uses Snowflake across several AWS regions. Their primary operational Snowflake account is in AWS US East (Ohio), but they need to leverage a specific AI_COMPLETE model, llama4-maverick, which is natively available in AWS US East 1 (N. Virginia) but not in US East (Ohio). To address this, the Snowflake administrator enables cross-region inference for their US East (Ohio) account.

- A. The llama4-maverick model is listed as natively available in AWS US East 1 (N. Virginia) and is supported for cross-region inference (AWS US Cross-Region), validating it as a suitable target for inference from US East (Ohio).
- B. To enable cross-region inference for the US East (Ohio) account, the administrator would execute the command: ALTER ACCOUNT SET AWS_US' ; to allow inference requests to be processed in any AWS US region where the model is available. CORTEX_ENABLED_CROSS_REGION =AWS_US' ;
- C. The query latency for cross-region inference with AI_COMPLETE is consistently low and predictable, as Snowflake's architecture is designed to completely negate the impact of geographical distance and network variations.
- D. Cross-region inference is fully supported for AI_COMPLETE in U.S. SnowGov regions for both inbound and outbound inference requests, provided the target model is available in the respective SnowGov region.
- E. User inputs, service-generated prompts, and the generated outputs from cross-region AI_COMPLETE calls are automatically stored or cached in the remote processing region to optimize performance for subsequent identical requests.

Answer: A,B

Explanation:

Option A is correct because the CORTEX_ENABLED_CROSS_REGION account parameter is used to enable cross-region inference. Setting it CORTEX_ENABLED_CROSS_REGION would permit inference requests to be processed in any AWS US region, such as N. Virginia, from a local AWS US region like Ohio. to 'Aws us' Option B is incorrect because user inputs, service generated prompts, and outputs are explicitly not stored or cached during cross-region inference. Option C is incorrect as cross-region inference is not supported in U.S. SnowGov regions for either inbound or outbound inference requests. Option D is correct because the sources indicate that AI COMPLETE (llama4-maverick) is natively available in AWS US East 1 (N. Virginia) and is supported for cross-region inference within AWS US regions, making it a valid target for the account in US East (Ohio). Option E is incorrect because latency between regions depends on the cloud provider infrastructure and network status, and Snowflake recommends testing specific use cases with cross-region inference enabled.

NEW QUESTION # 91

A data scientist is preparing to log a custom PyCaret classification model into the Snowflake Model Registry. The goal is to deploy this model on Snowpark Container Services (SPCS) for scalable inference. The PyCaret model relies on the 'pycaret' and 'scipy' Python libraries, and the data scientist has local 'sample data.csv' for inferring the model's signature. Which statements are crucial for successfully logging this custom model for eventual SPCS deployment?

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

Answer: C,E

Explanation:

Option B is correct because for models deployed to Snowpark Container Services, dependencies are typically obtained from

'conda-forge' or PyPI. Therefore, 'pip_requirements' is the appropriate way to specify PyPI packages, and explicitly setting ['SNOWPARK_CONTAINER_SERVICEST guides the deployment target. Option C is correct because either (a Pandas or Snowpark DataFrame) or a model 'signature' must be provided to the method for input validation and to infer the model's input signature. Option A is incorrect because 'conda_dependencies' in 'log_model' assumes the Snowflake channel for warehouse deployment, whereas for SPCS, it's 'conda-forge' or PyPI. Mixing channels or assuming Snowflake channel for SPCS is incorrect. Option D is incorrect. The 'use_gpu' option is used when 'loading' a model version `cmv.load(options={'use_gpu': to enable GPU-specific loading logic, not when logging the model. GPU requests for inference are specified when creating the service. Option E is incorrect. Snowflake recommends using only 'conda_dependencies' or 'pip_requirements', not both, to avoid potential compatibility issues during the container image build.`

NEW QUESTION # 92

After resolving this, they try to process a batch of 1500 documents in a single query using the method, which also fails. Which two issues are most likely contributing to these failures?

- A. The virtual warehouse being used for Document AI is not a Snowpark-optimized warehouse.
- B. The user has not been granted the 'CREATE STREAM' privilege on the schema.
- C. The batch of documents exceeds Document AI's limit of 1000 documents per query.
- D. The 'SNOWFLAKE.DOCUMENT INTELLIGENCE CREATOR' database role or the 'CREATE SNOWFLAKE.ML.DOCUMENT INTELLIGENCE' privilege on the schema has not been granted to the user.
- E. The name is not unique within the schema.

Answer: C,D

Explanation:

The initial error, 'Unable to create a build on the specified database and schema,' is a documented message that can occur if the 'CREATE SNOWFLAKE.ML.DOCUMENT INTELLIGENCE' privilege is not granted to the role, or if the database role is missing. Thus, option B is a direct and likely cause for the first failure. After resolving this (meaning the model build was successfully created), the subsequent failure to process 1500 documents in a single query is due to Document AI's limitation of processing a maximum of 1000 documents per query. Therefore, option D correctly identifies the second contributing factor. Option A (CREATE STREAM privilege) is relevant for creating pipelines but not for the initial model build or predicting. Option C (duplicate model build name) is another possible cause for the first error, but the question asks for the 'most likely' issues across two distinct failures, making B a more fundamental access control issue for creation and D a clear limit for processing. Option E (Snowpark-optimized warehouse) is not a specific requirement or a known cause for these types of errors; an X-Small, Small, or Medium warehouse is generally recommended, and scaling up does not increase processing speed.

NEW QUESTION # 93

A Data Engineer has successfully deployed a Document AI model build named 'expense reports' to extract 'total amount' and 'approver signature' from digital expense reports. They observe that sometimes the 'approver signature' is not present in a document, or certain table cells are intentionally left blank in other document types processed by Document AI. They also want to automate the ingestion and processing of new expense reports. Regarding the '!PREDICT' method's JSON output when 'approver signature' is missing or a table cell is empty, and the recommended Snowflake features for continuously processing new documents, which statements are true?

- A. If the 'approver_signature' is not found, the JSON output for 'approver_signature' will contain 'value': null, 'score': 0.0}.
- B. If the model does not find an answer, it returns an empty string for 'value' and a score indicating its confidence that no answer was found.
- C. Dynamic Tables are the primary recommended feature for continuous document processing with Document AI, replacing the need for streams and tasks.
- D. To automate continuous processing, the Data Engineer should create a stream on the internal stage where documents are uploaded, and a task that calls the '!PREDICT' method when new data arrives.
- E. When a table cell is empty, the Document AI model will return a 'score' key but no 'value' key for that cell.

Answer: D,E

Explanation:

Option B is correct because if a table cell is empty or if the Document AI model does not find an answer in the document, it will not return a 'value' key. However, it will still return the 'score' key, indicating its confidence that the cell is empty or that the answer is not found.

present. Option C is correct as creating document processing pipelines with Document AI typically involves setting up a stream on a stage to capture new documents and a task that then calls the '!PREDICT method to process these new documents automatically. Option A is incorrect; the 'value' key is omitted, not set to 'null'. Option D is incorrect because the 'value' key is omitted entirely, not set to an empty string. Option E is incorrect because Snowflake Cortex functions, including Document AI's '!PREDICT method, do not currently support dynamic tables.

NEW QUESTION # 94

A Gen AI Specialist is tasked with implementing a data pipeline to automatically enrich new customer feedback entries with sentiment scores using Snowflake Cortex functions. The new feedback arrives in a staging table, and the enrichment process must be automated and cost-effective. Given the following pipeline components, which combination of steps is most appropriate for setting up this continuous data augmentation process?

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

Answer: E

Explanation:

Option C is the most direct and efficient approach for continuously augmenting data with sentiment scores in a Snowflake pipeline. is a task-specific AI function designed for this purpose, returning an overall sentiment score for English-language text. SNOWFLAKE.CORTEX.SENTIMENT Integrating it directly into a task that monitors a stream allows for automated, incremental processing of new data as it arrives in the stage. The source explicitly mentions using Cortex functions in data pipelines via the SQL interface. Option A is plausible, but calling SENTIMENT directly in SQL within a task (Option C) is simpler and avoids the overhead of a Python UDF if the function is directly available in SQL, which it is. Option B, using a dynamic table, is not supported for Snowflake Cortex functions. Option D, while powerful for custom LLMs, is an over-engineered solution and introduces more complexity (SPCS setup, custom service) than necessary for a direct sentiment function. Option E describes a manual, non-continuous process, which contradicts the requirement for an automated pipeline.

NEW QUESTION # 95

.....

TestKingIT not only have a high reliability, but also provide a good service. If you choose TestKingIT, but don't pass the GES-C01 Exam, we will 100% refund full of your cost to you. TestKingIT also provide you with a free update service for one year.

Trustworthy GES-C01 Source: <https://www.testkingit.com/Snowflake/latest-GES-C01-exam-dumps.html>

Snowflake GES-C01 Free Exam Questions Check out features of these formats, Snowflake GES-C01 Free Exam Questions Besides our after-sales service engineers will be always online to give remote guidance and assistance for you if necessary, Snowflake GES-C01 Free Exam Questions In today's society, high efficiency is hot topic everywhere, When exam study material has new contents, the system will send you the latest GES-C01 latest study material to you with e-mail.

The project had invested a lot of money in a variety of test tool GES-C01 licenses and was caught by surprise by this new tool upgrade, Joe can add, edit, and remove contacts from his address book.

New GES-C01 Free Exam Questions | Efficient Snowflake GES-C01: SnowPro® Specialty: Gen AI Certification Exam 100% Pass

Check out features of these formats, Besides our after-sales Practice GES-C01 Test service engineers will be always online to give remote guidance and assistance for you if necessary.

In today's society, high efficiency is hot topic everywhere, When exam study material has new contents, the system will send you the latest GES-C01 latest study material to you with e-mail.

The names of these formats are GES-C01 desktop practice test software, web-based practice test software, and PDF dumps file.

- GES-C01 Certification Dumps
- GES-C01 New Cram Materials
- GES-C01 Learning Materials
- Open website

