

Snowflake GES-C01 Valid Test Papers & Valid GES-C01 Test Preparation



Our Snowflake dumps torrent contains everything you need to pass GES-C01 actual test smoothly. We always adhere to the principle that provides our customers best quality GES-C01 Exam Prep with most comprehensive service. This is the reason why most people prefer to choose our GES-C01 vce dumps as their best preparation materials.

Choosing our GES-C01 real dumps as your study guide means you choose a smart and fast way to get succeed in the certification exam. There are accurate GES-C01 test answers and some explanations along with the exam questions that will boost your confidence to solve the difficulty of GES-C01 Practice Test. You will enjoy great benefits if you buy our GES-C01 braindumps now and free update your study materials one-year.

>> Snowflake GES-C01 Valid Test Papers <<

2025 GES-C01 Valid Test Papers: SnowPro® Specialty: Gen AI Certification Exam - Trustable Snowflake Valid GES-C01 Test Preparation

We can't forget the advantages and the conveniences that reliable GES-C01 study materials compiled by our companies bring to us. First, by telling our customers what the key points of learning, and which learning GES-C01 method is available, they may save our customers money and time. They guide our customers in finding suitable jobs and other information as well. Secondly, a wide range of practice types and different version of our GES-C01 Study Materials receive technological support through our expert team.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q23-Q28):

NEW QUESTION # 23

A team is building a critical Document AI pipeline for continuous processing of new financial reports. They've identified that occasionally, the 'GET_PREIGNED_URL' generated for specific documents expires before the 'PREDICT' function can successfully process them, resulting in HTTP 403 errors. To mitigate this, they plan to modify their SQL query logic. Which approach effectively addresses the presigned URL expiration issue without altering the stage definition or the model build itself, and adheres to recommended practices for handling batch processing as described in Snowflake's troubleshooting documentation?

- A.

```
SELECT model_name||PREDICT(GET_PREIGNED_URL(@stage_name, RELATIVE_PATH, 'expiration_time=3600')) FROM DIRECTORY(@stage_name);
```

- B.

```
CREATE TASK process_docs WAREHOUSE = my_wh SCHEDULE = '1 minute' AS MERGE INTO processed_table USING (SELECT RELATIVE_PATH, model_name||  
PREDICT(GET_PREIGNED_URL(@stage_name, RELATIVE_PATH)) AS extracted_data FROM DIRECTORY(@stage_name) WHERE NOT EXISTS  
(SELECT 1 FROM processed_table WHERE processed_table.file_path = RELATIVE_PATH) LIMIT 500) AS new_docs ON FALSE WHEN NOT MATCHED THEN INSERT  
(file_path, extracted_data) VALUES (new_docs.RELATIVE_PATH, new_docs.extracted_data);
```

- C.
- D.

```
ALTER STAGE @stage_name SET COPY_OPTIONS = (ON_ERROR = 'CONTINUE');
```

```
SELECT model_name!PREDICT(GET_PREIGNED_URL(@stage_name, RELATIVE_PATH)) FROM DIRECTORY(@stage_name) WHERE METADATA$FILE_LAST_MODIFIED > DATEADD(hour, -1, CURRENT_TIMESTAMP());
```

Answer: C

Explanation:

The core problem is the expiration of presigned URLs when processing documents with S!PREDICT, which defaults to 60 minutes. Snowflake's troubleshooting documentation specifically recommends 'Use several queries to process the documents' as a solution for 'Presigned URL has expired'. This implies breaking down the workload into smaller, more manageable batches to ensure that the processing for each batch completes within the URL's active lifespan. Option A is incorrect as when used with a stage and in the context of does not directly support an 'expiration_time' parameter within the provided syntax examples. Option B describes a task that processes new documents in batches, which is a practical implementation of the recommended solution (using several queries), but option D describes the underlying recommended strategy more broadly and accurately as per the documentation's troubleshooting guidance. Option C attempts to filter documents by modification time, which doesn't directly prevent a URL from expiring if the subsequent processing is slow. Option E modifies 'COPY OPTIONS' which is irrelevant to '!PREDICT query errors for Document AI.

NEW QUESTION # 24

A data platform administrator needs to retrieve a consolidated overview of credit consumption for all Snowflake Cortex AI functions (e.g., LLM functions, Document AI, Cortex Search) across their entire account for the past week. They are interested in the aggregated daily credit usage rather than specific token counts per query. Which Snowflake account usage views should the administrator primarily leverage to gather this information?

☐ The SNOWFLAKE.ACCOUNT_USAGE.CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY view to get detailed token usage for each LLM function call, then aggregate manually.

☐ The SNOWFLAKE.ORGANIZATION_USAGE.METERING_DAILY_HISTORY view, specifically filtering for SERVICE_TYPE = 'AI_SERVICES'.

☐ The SNOWFLAKE.ACCOUNT_USAGE.CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY view for Document AI costs, and SNOWFLAKE.ACCOUNT_USAGE.CORTEX_SEARCH_DAILY_USAGE_HISTORY for Cortex Search costs, then combine them.

☒ Only the SNOWFLAKE.ACCOUNT_USAGE.QUERY_HISTORY view, analyzing the EXECUTION_STATUS and TOTAL_ELAPSED_TIME columns for queries involving Cortex functions.

☐ The SNOWFLAKE.CORTEX.COUNT_TOKENS function to re-calculate estimated costs for all past queries that used Cortex AI functions.

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

Answer: A

Explanation:

For a consolidated overview of credit consumption for all AI services, including Cortex LLM Functions, Cortex Analyst, and Document AI, the SNOWFLAKE.ORGANIZATION_USAGE.METERING_DAILY_HISTORY view provides daily credit usage for an account, and filtering by 'SERVICE_TYPE = 'AI_SERVICES'' allows administrators to view aggregated credit consumption for all AI services. This directly addresses the need for a consolidated, aggregated view of credit usage. Option A is incorrect because 'CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY' provides granular, per-query token details for LLM functions (e.g., 'COMPLETE', 'TRY_COMPLETE'), which is not what the administrator is primarily looking for (aggregated daily usage for *all* AI services), and does not cover other AI services like Document AI or Cortex Search. Option C is partially correct in that these views ('CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY' and 'CORTEX_SEARCH_DAILY_USAGE_HISTORY') track specific AI services (Document AI, Cortex Search). However, 'METERING_DAILY_HISTORY' with 'SERVICE_TYPE = 'AI_SERVICES'' provides a more encompassing and already aggregated view for all AI services, fulfilling the requirement for a 'consolidated overview' more efficiently. Option D is incorrect as 'QUERY_HISTORY' primarily focuses on general query execution metadata like 'EXECUTION_STATUS' and 'TOTAL_ELAPSED_TIME' and does not directly provide credit consumption specific to AI services. Option E is incorrect because 'COUNT_TOKENS' is used for estimating future costs or token counts.

NEW QUESTION # 25

A data science team is using SNOWFLAKE.CORTEX.CLASSIFY_TEXT to categorize product reviews into detailed segments like 'Bug Report - Critical', 'Feature Request - UI/UX', 'General Praise', or 'Query - Billing Issue'. For highly nuanced reviews, they find the initial classifications lack precision, and they are also concerned about the associated compute costs for processing large volumes of data. Which strategies should they employ to optimize classification accuracy and manage costs effectively with this function?

- A. CLASSIFY_TEXT labels, descriptions, and examples are counted as input tokens only once per function call, regardless of the number of records processed in a batch, to optimize cost efficiency.
- B. To reduce input token costs for classifications, the input text should be pre-processed to remove common stop words and punctuation, as these characters are counted as billable tokens without contributing to classification accuracy.
- C. To improve accuracy for ambiguous classifications, they should augment the list_of_categories with explicit description and examples for each category, understanding that these additions will increase input token costs for each record processed.

```
SELECT SNOWFLAKE.CORTEX.CLASSIFY_TEXT(
  'The app freezes after login, making it unusable.',
  [
    { 'label': 'Bug Report - Critical', 'description': 'Software defect causing severe issues with application functionality.', 'examples': ['App crashes frequently', 'Login button unresponsive'] },
    { 'label': 'Feature Request - UI/UX', 'description': 'Suggestion for user interface or experience improvement.' }
  ]
);
```

- D. For complex scenarios where the relationship between review text and categories is not straightforward, including a concise task_description (e.g., 'Classify the product review focusing on technical support relevance') in the options argument is recommended to guide the model.

```
SELECT SNOWFLAKE.CORTEX.CLASSIFY_TEXT(
  'The new update introduced a confusing navigation bar that is hard to use.',
  [
    'Bug Report', 'Feature Request'
  ],
  { 'task_description': 'Classify the feedback regarding app usability issues.' }
);
```

- E. If classifying thousands of reviews, they can significantly reduce overall compute costs by setting the temperature option to 0.0 within CLASSIFY_TEXT to ensure deterministic and cheaper inference.

Answer: C,D

Explanation:

Option A is correct because adding label descriptions and examples can improve classification accuracy, especially when category definitions are ambiguous. The source explicitly states that each label, description, and example counts as input tokens for each record processed by a 'CLASSIFY_TEXT' function call, incurring costs accordingly. Option B is correct because adding a clear 'task_description' can improve accuracy when the relationship between the input text and categories is ambiguous or nuanced. Option C is incorrect; while token counts contribute to cost, the sources do not recommend removing stop words and punctuation for cost reduction or as a general best practice for SCLASSIFY TEXT. The focus is on using plain English input. Option D is incorrect because the 'temperature' option is available for 'COMPLETE' and functions to control output randomness, but it is not listed as an option for 'CLASSIFY_TEXT' in its syntax. Furthermore, while a lower temperature can make results more deterministic, the source does not link it to 'cheaper' inference cost for these task-specific functions, but rather to consistency for 'COMPLETE'. Option E is incorrect because 'AI_CLASSIFY' labels, descriptions, and examples are indeed counted as input tokens for 'each record processed', not just once per call, as clearly stated in the cost considerations.

NEW QUESTION # 26

An analytics team is preparing documents for a new Document AI model build to extract information from internal policy reviews. They have a variety of documents that they intend to upload to an internal stage for processing. The document list includes: (1) a 70 MB PDF with 100 pages, (2) a 45 MB DOCX with 150 pages, (3) a 30 MB PNG image, (4) a 60 MB TIFF image, and (5) a 20 MB HTML file. All documents are in English. Which of these documents would 'fail' to meet the direct input requirements for Document AI processing?

- A. The 70 MB PDF with 100 pages.
- B. The 45 MB DOCX with 150 pages.
- C. The 60 MB TIFF image.
- D. The 30 MB PNG image.
- E. The 20 MB HTML file.

Answer: A,B,C

Explanation:

Document AI has specific limitations for document processing. Documents must be no more than 125 pages long and 50 MB or less in size. Supported formats include PDF, PNG, DOCX, EML, JPEG/JPG, HTM/HTML, TEXT/TXT, and TIF/TIFF. - **Option A** (70 MB PDF, 100 pages): Fails on the maximum file size requirement (max 50 MB). - **Option B** (45 MB DOCX, 150 pages): Fails on the maximum page count requirement (max 125 pages). - **Option C** (30 MB PNG image): Meets requirements as it is within size limits and a supported format. - **Option D** (60 MB TIFF image): Fails on the maximum file size requirement (max 50 MB). - **Option E** (20 MB HTML file): Meets requirements as it is within size limits and a supported format.

NEW QUESTION # 27

A new ML Engineer, 'data_scientist_role', has been assigned to a project involving custom machine learning models in Snowflake. They need to gain the necessary permissions to perform the following actions related to Snowflake Model Registry and Snowpark Container Services: 1. Log a custom model into a specified schema. 2. Deploy that model to an existing Snowpark Container Service compute pool. 3. Call the deployed model for inference using SQL. Which of the following SQL commands grant the 'minimal' required privileges to the for these actions, assuming the compute pool and image repository already exist and are appropriately configured?

- A.

```
GRANT CREATE MODEL ON SCHEMA my_ml_schema TO ROLE data_scientist_role;
```
- B.

```
GRANT CREATE COMPUTE POOL ON ACCOUNT TO ROLE data_scientist_role;
```
- C.

```
GRANT USAGE ON DATABASE my_database TO ROLE data_scientist_role;
```
- D.

```
GRANT READ ON IMAGE REPOSITORY my_ml_images TO ROLE data_scientist_role;
```
- E.

```
GRANT READ ON IMAGE REPOSITORY my_ml_images TO ROLE data_scientist_role;
```

Answer: A,B

Explanation:

Option A is correct because the 'CREATE MODEL' privilege on the target schema is required to log a new model (which creates a model object) in the Snowflake Model Registry. Option D is correct because deploying a model to a Snowpark Container Service creates a service object within a schema, which requires the 'CREATE SERVICE' privilege on that schema. The role would also implicitly need 'USAGE' on the specified compute pool. Option B is incorrect. While 'USAGE ON DATABASE' is generally needed for accessing objects within a database, it's a broader prerequisite and not specifically a minimal privilege for the direct model registry actions of logging, deploying, and calling the model. Option C is incorrect because 'CREATE COMPUTE POOL' is for creating the compute pool itself, not for deploying a service 'to' an existing one. The role would need 'USAGE' on the existing compute pool, but not the right to create it from scratch for this scenario. Option E is incorrect because 'READ ON IMAGE REPOSITORY' is required for the 'service' to pull the image from the repository, but the question asks for privileges for the to perform the 'actions' of logging, deploying, and calling. While the role might need to manage or verify the image, this isn't a direct privilege for the user's interaction with the deployed model in the same way 'CREATE MODEL' or 'CREATE SERVICE' are.

NEW QUESTION # 28

.....

Due to lots of same products in the market, maybe you have difficulty in choosing the GES-C01 guide test. We can confidently tell you that our products are excellent in all aspects. You can directly select our products. Firstly, we have free trials of the GES-C01 exam study materials to help you know our products. One of the great advantages is that you will soon get a feedback after you finish the exercises. So you are able to adjust your learning plan of the GES-C01 Guide test flexibly. We hope that our new design can make study more interesting and colorful. You also can send us good suggestions about developing the study material.

Valid GES-C01 Test Preparation: <https://www.pass4suresvce.com/GES-C01-pass4sure-vce-dumps.html>

If you are still troubled for the Snowflake GES-C01 certification exam, then select the Pass4suresVCE's training materials please, It is the short version of our official GES-C01 dumps pdf, Like our innovative Valid GES-C01 Test Preparation - SnowPro® Specialty: Gen AI Certification Exam Practice Tests, they introduce you to the real exam scenario, Attention!

The ip routing global command is missing from the GES-C01 configuration, Close the Sharing Preferences panel, or continue configuring other sharing options, If you are still troubled for the Snowflake GES-C01 Certification Exam, then select the Pass4suresVCE's training materials please.

Snowflake's Exam Questions for GES-C01 Ensure 100% Success on Your First Attempt

It is the short version of our official GES-C01 dumps pdf, Like our innovative SnowPro® Specialty: Gen AI Certification Exam Practice Tests, they introduce you to the real exam scenario, Attention!

Some candidates like to take help of their friends or tutors, while some simply rely on GES-C01 books.

- Free PDF 2025 Snowflake GES-C01 Newest Valid Test Papers □ Download 【GES-C01】 for free by simply searching on ➡ www.testsdumps.com □□□ □Latest GES-C01 Test Voucher
- GES-C01 PDF Cram Exam □ Study Materials GES-C01 Review □ New GES-C01 Exam Dumps □ Download [GES-C01] for free by simply searching on “www.pdfvce.com” □New GES-C01 Exam Dumps
- Free PDF 2025 Snowflake GES-C01 Newest Valid Test Papers □ Open website “www.examcollectionpass.com” and search for ➡ GES-C01 □ for free download □Latest GES-C01 Test Voucher
- SnowPro® Specialty: Gen AI Certification Exam valid training collection - GES-C01 study prep torrent - SnowPro® Specialty: Gen AI Certification Exam exam practice pdf □ Search for □ GES-C01 □ and easily obtain a free download on 《www.pdfvce.com》 □GES-C01 Test Sample Questions
- GES-C01 PDF Cram Exam ➔ GES-C01 Exam Price ↘ New GES-C01 Test Prep □ The page for free download of ➡ GES-C01 □ on [www.prep4pass.com] will open immediately □New GES-C01 Test Prep
- Free PDF 2025 Snowflake GES-C01 Newest Valid Test Papers □ Search for ✓ GES-C01 □✓□ and download it for free on 【www.pdfvce.com】 website □New Exam GES-C01 Braindumps
- New GES-C01 Test Prep □ Testking GES-C01 Exam Questions □ Exam GES-C01 Blueprint □ Search for ▷ GES-C01 ◁ and obtain a free download on 《www.pass4leader.com》 □Testking GES-C01 Exam Questions
- 100% Pass Quiz Snowflake - Fantastic GES-C01 Valid Test Papers ↗ Open ▶ www.pdfvce.com ◀ and search for 「GES-C01」 to download exam materials for free □GES-C01 Detailed Study Plan
- New GES-C01 Exam Sample □ New GES-C01 Exam Dumps □ GES-C01 PDF Cram Exam □ Search for （GES-C01） on “www.vceengine.com” immediately to obtain a free download □GES-C01 Detailed Study Plan
- GES-C01 Dumps Collection □ Exam GES-C01 Blueprint □ GES-C01 Dumps Reviews □ Search on [www.pdfvce.com] for ➡ GES-C01 □ to obtain exam materials for free download □GES-C01 Dumps Collection
- New GES-C01 Test Prep □ GES-C01 Reliable Exam Prep □ GES-C01 Dumps Collection □ Open （www.pdfdumps.com） enter ☀ GES-C01 □☀□ and obtain a free download □Exam GES-C01 Blueprint
- myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, ncon.edu.sa, johnlee994.nizarblog.com, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.mygradeapro.com, daotao.wisebusiness.edu.vn, Disposable vapes