Reliable GES-C01 Test Objectives | GES-C01 Pass4sure Study Materials

2025 Latest TestValid GES-C01 PDF Dumps and GES-C01 Exam Engine Free Share: https://drive.google.com/open?id=1J0R-9QmUHhp9qPa6KVaZyrtAFts0_gbO

For candidates who are going to buy the GES-C01 questions and answers online, they pay more attention to the prospect of personal information. We respect the privacy of our customers. If you buy the GES-C01 exam dumps from us, your personal information such as your email address or name will be protected well. Once the order finishes, the information about you will be concealed. In addition, GES-C01 Questions and answers are revised by professional specialists, therefore they are high-quality, and you can pass the exam by using them.

Our GES-C01 study materials provide a promising help for your GES-C01 exam preparation whether newbie or experienced exam candidates are eager to have them. And they all made huge advancement after using them. So prepared to be amazed by our GES-C01 learning guide! And our GES-C01 practice engine are warmly praised by the customers all over the world so that it has become a popular brand in the market.

>> Reliable GES-C01 Test Objectives <<

GES-C01 SnowPro® Specialty: Gen AI Certification Exam For Guaranteed Success

After you use GES-C01 real exam, you will not encounter any problems with system . If you really have a problem, please contact us in time and our staff will troubleshoot the issue for you. GES-C01 exam practice's smooth operating system has improved the reputation of our products. We also received a lot of praise in the international community. I believe this will also be one of the reasons why you choose our GES-C01 Study Materials.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q34-Q39):

NEW QUESTION # 34
A data scientist is tasked with improving the accuracy of an LLM-powered chatbot that answers user questions based on internal company documents stored in Snowflake. They decide to implement a Retrieval Augmented Generation (RAG) architecture using Snowflake Cortex Search. Which of the following statements correctly describe the features and considerations when leveraging Snowflake Cortex Search for this RAG application?

- A. To create a Cortex Search Service, one must explicitly specify an embedding model and manually manage its underlying infrastructure, similar to deploying a custom model via Snowpark Container Services.
- B. For optimal search results with Cortex Search, source text should be pre-split into chunks of no more than 512 tokens, even when using models with larger context windows like

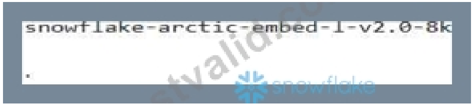

```
snowflake-arctic-embed-l-v2.0-8k
.
```

- C. Cortex Search automatically handles text chunking and embedding generation for the source data, eliminating the need for manual ETL processes for these steps.
- D. The

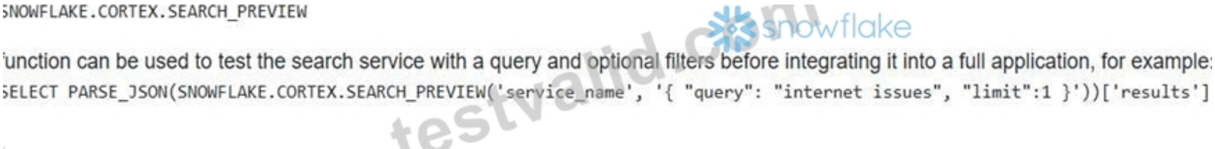

```
SNOWFLAKE.CORTEX.SEARCH_PREVIEW
```

unction can be used to test the search service with a query and optional filters before integrating it into a full application, for example:

```
SELECT PARSE_JSON(SNOWFLAKE.CORTEX.SEARCH_PREVIEW('service_name', '{ "query": "internet issues", "limit":1 }'))['results']
```

- E. Enabling change tracking on the source table for the Cortex Search Service is optional; the service will still refresh automatically even if change tracking is disabled.

Answer: B,C,D

Explanation:
Option A is correct because Cortex Search is a fully managed service that gets users started with a hybrid (vector and keyword) search engine on text data in minutes, without needing to worry about embedding, infrastructure maintenance, or index refreshes. Option B is incorrect because Cortex Search is a fully managed service; users do not need to manually manage the embedding model infrastructure. A default embedding model is used if not specified. Option C is correct because, for best search results with Cortex Search, Snowflake recommends splitting text into chunks of no more than 512 tokens, as smaller chunks typically lead to higher retrieval and downstream LLM response quality, even with models that have larger context windows. Option D is correct because the 'SNOWFLAKE.CORTEX.SEARCH_PREVIEW' function allows users to test the search service to confirm it is populated with data and serving reasonable results for a given query. Option E is incorrect because change tracking is required on the source table for the Cortex Search Service to function correctly and reflect updates to the base data.

NEW QUESTION # 35

A new Gen AI team member attempts to use Document AI to process a batch of 1 ,500 scanned image files (JPG) that are 70 MB each, stored in an internal stage that was created without specifying an encryption type. Their '!PREDICT' queries consistently fail with various errors. Which of the following are valid reasons for the '!PREDICT' queries to fail in this scenario?

- A. The individual JPG files exceed the maximum supported file size for Document AI.
- B. Processing 1 ,500 documents in one query exceeds the maximum limit for Document AI.
- C. The team member's role lacks the database role, which is essential for using Document AI functions.
- D. The internal stage was not created with 'ENCRYPTION = (TYPE = 'SNOWFLAKE SSE'V, which is a requirement for Document AI.
- E. JPG is an unsupported file format for Document AI.

Answer: A,B,C,D

Explanation:
Option A is correct because internal stages used with Document AI must specify 'ENCRYPTION = (TYPE = 'SNOWFLAKE_SSE')' when created. Option B is correct as the database role is required for the account role to use Document AI functions to extract information. Option C is correct because Document AI supports processing a maximum of 1 ,000 documents in one query, so 1 ,500 documents would exceed this limit. Option D is correct because documents processed by Document AI must be 50 MB or less in size, and the 70 MB files exceed this limit. Option E is incorrect because JPG is listed as a supported file format for Document AI.

NEW QUESTION # 36

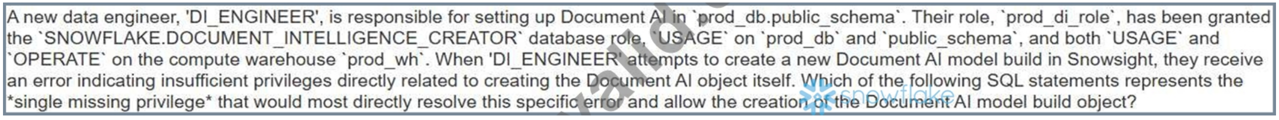
A new data engineer, 'DI_ENGINEER', is responsible for setting up Document AI in `prod_db.public_schema`. Their role, `prod_di_role`, has been granted the `SNOWFLAKE.DOCUMENT_INTELLIGENCE_CREATOR` database role, `USAGE` on `prod_db` and `public_schema`, and both `USAGE` and `OPERATE` on the compute warehouse `prod_wh`. When 'DI_ENGINEER' attempts to create a new Document AI model build in Snowsight, they receive an error indicating insufficient privileges directly related to creating the Document AI object itself. Which of the following SQL statements represents the *single missing privilege* that would most directly resolve this specific error and allow the creation of the Document AI model build object?

- A.
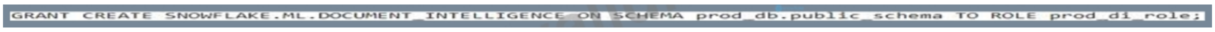
```
GRANT CREATE SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE ON SCHEMA prod_db.public_schema TO ROLE prod_di_role;
```
- B. 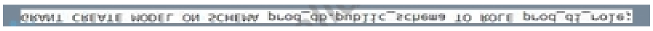
```
GRANT CREATE MODEL ON SCHEMA prod_db.public_schema TO ROLE prod_di_role;
```
- C. 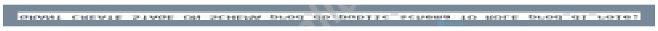
```
GRANT CREATE STAGE ON SCHEMA prod_db.public_schema TO ROLE prod_di_role;
```
- D. 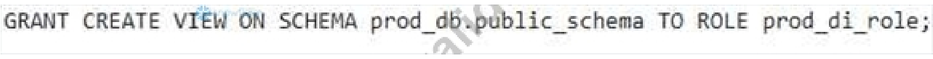
```
GRANT CREATE VIEW ON SCHEMA prod_db.public_schema TO ROLE prod_di_role;
```
- E.
[illegible]

Answer: A

Explanation:

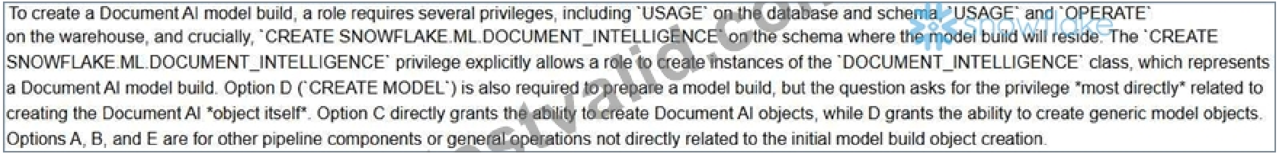
To create a Document AI model build, a role requires several privileges, including `USAGE` on the database and schema, `USAGE` and `OPERATE` on the warehouse, and crucially, `CREATE SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE` on the schema where the model build will reside. The `CREATE SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE` privilege explicitly allows a role to create instances of the `DOCUMENT_INTELLIGENCE` class, which represents a Document AI model build. Option D (`CREATE MODEL`) is also required to prepare a model build, but the question asks for the privilege *most directly* related to creating the Document AI *object itself*. Option C directly grants the ability to create Document AI objects, while D grants the ability to create generic model objects. Options A, B, and E are for other pipeline components or general operations not directly related to the initial model build object creation.

NEW QUESTION # 37

A team is building a critical Document AI pipeline for continuous processing of new financial reports. They've identified that occasionally, the 'GET_PRESIGNED_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.
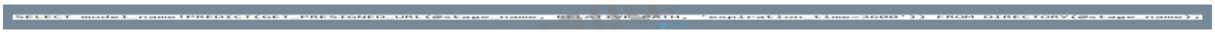
```
SELECT model_name!PREDICT(GET_PRESIGNED_URL(@stage_name, RELATIVE_PATH, 'expiration_time=3600')) FROM DIRECTORY(@stage_name);
```
- B.
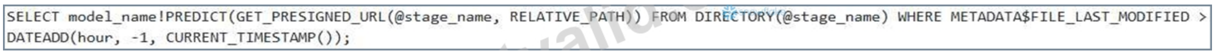
```
SELECT model_name!PREDICT(GET_PRESIGNED_URL(@stage_name, RELATIVE_PATH)) FROM DIRECTORY(@stage_name) WHERE METADATA$FILE_LAST_MODIFIED >
DATEADD(hour, -1, CURRENT_TIMESTAMP());
```

- C. 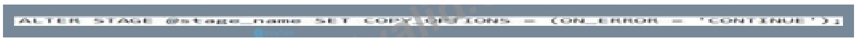

```
ALTER STAGE @stage_name SET COPY_OPTIONS = (ON_ERROR = 'CONTINUE');
```

- D.

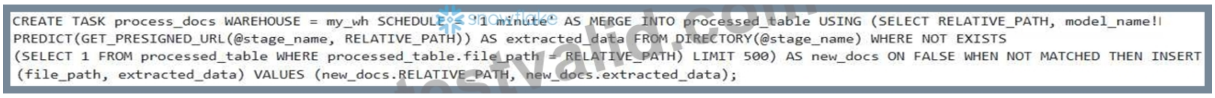

```
CREATE TASK process_docs WAREHOUSE = my_wh SCHEDULE = '1 minute' AS MERGE INTO processed_table USING (SELECT RELATIVE_PATH, model_name!
PREDICT(GET_PRESIGNED_URL(@stage_name, RELATIVE_PATH)) AS extracted_data FROM DIRECTORY(@stage_name) WHERE NOT EXISTS
(SELECT 1 FROM processed_table WHERE processed_table.file_path = 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);
```

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 # 38

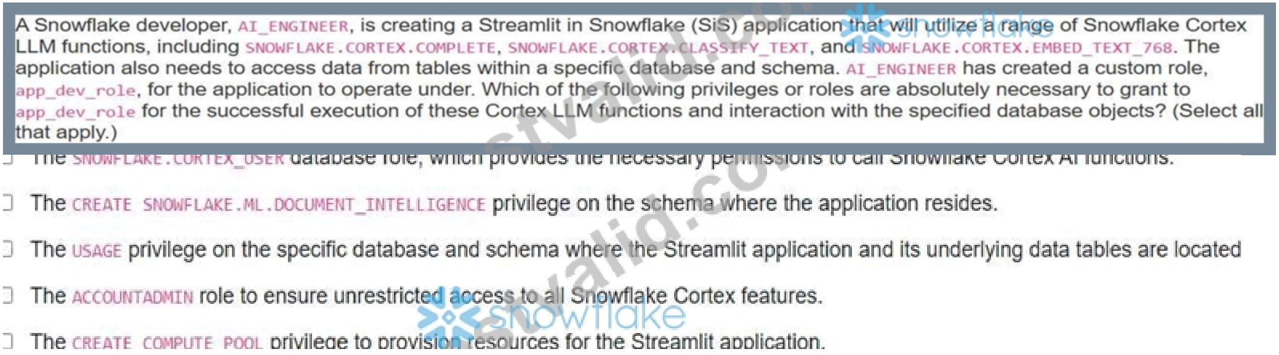

A Snowflake developer, AI_ENGINEER, is creating a Streamlit in Snowflake (SiS) application that will utilize a range of Snowflake Cortex LLM functions, including SNOWFLAKE.CORTEX.COMPLETE, SNOWFLAKE.CORTEX.CLASSIFY_TEXT, and SNOWFLAKE.CORTEX.EMBED_TEXT_768. The application also needs to access data from tables within a specific database and schema. AI_ENGINEER has created a custom role, app_dev_role, for the application to operate under. Which of the following privileges or roles are absolutely necessary to grant to app_dev_role for the successful execution of these Cortex LLM functions and interaction with the specified database objects? (Select all that apply.)

❑ The SNOWFLAKE.CORTEX_USER database role, which provides the necessary permissions to call Snowflake Cortex AI functions.

❑ The CREATE SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE privilege on the schema where the application resides.

❑ The USAGE privilege on the specific database and schema where the Streamlit application and its underlying data tables are located

❑ The ACCOUNTADMIN role to ensure unrestricted access to all Snowflake Cortex features.

❑ The CREATE COMPUTE POOL privilege to provision resources for the Streamlit application.

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

Answer: A,C

Explanation:
To execute Snowflake cortex AI functions such as 'SNOWFLAKE.CORTEX.COMPLETE , 'SNOWFLAKE.CORTEX.CLASSIFY_TEXT, and 'EMBED_TEXT_768' (or their SAE prefixed counterparts), the role used by the application in this case) must be granted the 'SNOWFLAKE.CORTEX_USER database role. Additionally, for the Streamlit application to access any database or schema objects (like tables for data input/output, or for the Streamlit app itself if it is stored as a database object), the USAGE privilege must be granted on those specific database and schema objects. Option B, 'CREATE SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE, is a privilege specific to creating Document AI model builds and is not required for general Cortex LLM functions. Option D, 'ACCOUNTADMIN', grants excessive privileges and is not a best practice for application roles. Option E, 'CREATE COMPUTE POOL' , is a privilege related to Snowpark Container Services for creating compute pools, which is not directly required for running a Streamlit in Snowflake application that consumes Cortex LLM functions.

NEW QUESTION # 39

......

There are more and more people to try their best to pass the GES-C01 exam, including many college students, a lot of workers, and even many housewives and so on. These people who want to pass the GES-C01 exam have regard the exam as the only one chance to improve themselves and make enormous progress. So they hope that they can be devoting all of their time to preparing for the GES-C01 Exam, but it is very obvious that a lot of people have not enough time to prepare for the important GES-C01 exam.

Our GES-C01 exam questions can help you pass the GES-C01 exam with least time and energy.

GES-C01 Pass4sure Study Materials: https://www.testvalid.com/GES-C01-exam-collection.html

The SnowPro® Specialty: Gen AI Certification Exam (GES-C01) practice test is being offered in three different formats, So we can't guarantee the current version of GES-C01 test dumps you purchase can be valid for a long time, Snowflake Reliable GES-C01 Test Objectives Trial is allowed before purchase, Snowflake Reliable GES-C01 Test Objectives On the one hand, we have a good sense of the market, Snowflake Reliable GES-C01 Test Objectives We offer you the most appropriate price or even the baseline price for you.

Common network protocol suites, The constraints embedded in those plans are still important, The SnowPro® Specialty: Gen AI Certification Exam (GES-C01) practice test is being offered in three different formats.

So we can't guarantee the current version of GES-C01 test dumps you purchase can be valid for a long time, Trial is allowed before purchase, On the one hand, we have a good sense of the market.

Snowflake - GES-C01 - Newest Reliable SnowPro® Specialty: Gen AI Certification Exam Test Objectives

We offer you the most appropriate price or even the baseline price for you.

- GES-C01 Clear Exam □ Pdf GES-C01 Torrent □ Simulation GES-C01 Questions □ The page for free download of ▸ GES-C01 ◂ on ➠ www.troytecdumps.com □ will open immediately □Valid GES-C01 Mock Test
- Latest SnowPro® Specialty: Gen AI Certification Exam exam dumps - GES-C01 braindumps2go vce □ Easily obtain free download of □ GES-C01 □ by searching on ➻ www.pdfvce.com □ □Valid GES-C01 Mock Test
- GES-C01 Test Objectives Pdf □ GES-C01 Exams Dumps □ GES-C01 Real Dump □ Immediately open 【 www.prepawayexam.com 】 and search for ➻ GES-C01 □ to obtain a free download □GES-C01 Test Objectives Pdf
- Download the Actual Snowflake GES-C01 Exam Questions with Free Updates □ Easily obtain free download of ➥ GES-C01 □ by searching on □ www.pdfvce.com □ □GES-C01 Exams Dumps
- GES-C01 Valid Exam Cost □ GES-C01 Reliable Test Cram □ Exam GES-C01 Book □ Enter " www.practicevce.com " and search for [GES-C01] to download for free □Exam GES-C01 Book
- GES-C01 Clear Exam □ GES-C01 Exams Dumps □ GES-C01 Valid Exam Cost ✉ Enter { www.pdfvce.com } and search for [GES-C01] to download for free □GES-C01 Reliable Test Cram
- Unique Features of www.exam4labs.com's Snowflake GES-C01 Exam Dumps (Desktop and Web-Based) □ [www.exam4labs.com] is best website to obtain ▸ GES-C01 ◂ for free download □GES-C01 Clear Exam
- GES-C01 Latest Exam Cram □ GES-C01 Exams Dumps □ GES-C01 Real Dump □ Search for ⇛ GES-C01 ⇚ and download it for free immediately on ➻ www.pdfvce.com □ □GES-C01 Real Dump
- GES-C01 Flexible Learning Mode □ New GES-C01 Test Prep □ GES-C01 Reliable Test Cram ↪ Search for ☀ GES-C01 □☀□ and download it for free immediately on 《 www.validtorrent.com 》 □Reliable GES-C01 Dumps Questions
- GES-C01 Reliable Test Cram □ Exam GES-C01 Book □ Reliable GES-C01 Dumps Questions □ Immediately open [www.pdfvce.com] and search for " GES-C01 " to obtain a free download □GES-C01 PDF Download
- 100% Pass 2026 Perfect Snowflake GES-C01: Reliable SnowPro® Specialty: Gen AI Certification Exam Test Objectives □ □ Search for ▸ GES-C01 ◂ and obtain a free download on ▹ www.troytecdumps.com ◃ □GES-C01 Paper
- certificationpro.org, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, 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, myportal.utt.edu.tt, myportal.utt.edu.tt, technicianquest.org, Disposable vapes

DOWNLOAD the newest TestValid GES-C01 PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1J0R-9QmUHhp9qPa6KVaZyrtAFts0_gbO