

Reliable Snowflake GES-C01 Exam Simulations & Latest Real GES-C01 Exam



DOWNLOAD the newest TestsDumps GES-C01 PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=14uBNy3Hy8G_fdYi2Nlboysj1xGhes

In order to help customers, who are willing to buy our GES-C01 test torrent, make good use of time and accumulate the knowledge, Our company have been trying our best to reform and update our SnowPro® Specialty: Gen AI Certification Exam exam tool. “Quality First, Credibility First, and Service First” is our company’s purpose, we deeply hope our GES-C01 study materials can bring benefits and profits for our customers. So we have been persisting in updating our GES-C01 Test Torrent and trying our best to provide customers with the latest study materials. More importantly, the updating system we provide is free for all customers. If you decide to buy our GES-C01 study materials, we can guarantee that you will have the opportunity to use the updating system for free.

As we all know, we are now facing more and more competition. The GES-C01 exam is an important way to improve our competitiveness. The certification can show others whether we have a certain skill, whether we meet the requirements of others, for us. Get approved at work to increase your chips. For different needs, our GES-C01 Certification Exam questions are flexible and changeable. On the one hand, GES-C01 pdf files allow you to make full use of fragmented time, and you will be able to pass the GES-C01 exam with the least time and effort with our GES-C01 training materials.

>> Reliable Snowflake GES-C01 Exam Simulations <<

Latest Real GES-C01 Exam - GES-C01 Latest Dumps Ppt

Our GES-C01 questions pdf is up to date, and we provide user-friendly GES-C01 practice test software for the SnowPro® Specialty: Gen AI Certification Exam exam. Moreover, we are also providing money back guarantee on all of SnowPro® Specialty: Gen AI Certification Exam test products. If the GES-C01 braindumps products fail to deliver as promised, then you can get your money back. The GES-C01 Sample Questions include all the files you need to prepare for the Snowflake GES-C01 exam. With the help of the GES-C01 practice exam questions and test software, you will be able to feel the real GES-C01 exam scenario, and it will allow you to assess your skills.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q228-Q233):

NEW QUESTION # 228

A data processing team is using Snowflake Document AI to extract data from incoming supplier invoices. They observe that many documents are failing to process, and successful extractions are taking longer than expected, leading to increased costs. Upon investigation, they find error messages such as

```
[ "__processingErrors": [ "Document has too many pages. Actual: 130. Maximum: 125." ] ]  
[ "__processingErrors": [ "File exceeds maximum size. Actual: 54096026 bytes. Maximum: 50000000 bytes." ] ]  
and  
[ "processingErrors": [ "cannot identify image file < io.BytesIO object at 0x...>" ] ]  
Additionally, their 'X-LARGE' virtual warehouse is constantly active, contributing to higher-than-anticipated bills. Which two of the following actions are essential steps to troubleshoot and address the root causes of these processing errors and optimize their Document AI pipeline?
```

- A. Redefine extraction questions to be more generic and encompassing, reducing the number of distinct questions needed per document.
- B. Scale down the virtual warehouse to 'X-SMALC' or 'SMALL' size, as larger warehouses do not increase Document AI query processing speed and incur unnecessary costs.
- C. Increase the 'max_tokens' parameter within the 'PREDICT' function options to accommodate longer document responses from the model.
- D. **Configure the internal stage used for storing invoices with 'ENCRYPTION = (TYPE = 'SNOWFLAKE_SSE')'.**
- E. **Implement a pre-processing step to split documents exceeding 125 pages or 50 MB into smaller, compliant files before loading to the stage.**

Answer: D,E

Explanation:

The error messages 'Document has too many pages. Actual: 130. Maximum: 125.' and 'File exceeds maximum size. Actual: 54096026 bytes. Maximum: 50000000 bytes.' directly indicate that the documents do not meet Document AI's input requirements, which specify a maximum of 125 pages and 50 MB file size. Therefore, implementing a pre-processing step to split or resize these documents is an essential solution (Option B). The error 'cannot identify image file < io.BytesIO object at 0x...>' is a known issue that occurs when an internal stage used for Document AI is not configured with 'SNOWFLAKE_SSE' encryption. Correctly configuring the stage with this encryption type is crucial for resolving this processing error (Option D). Option A, while addressing cost optimization, is not a root cause of the 'processing errors' themselves, although it is a best practice for cost governance as larger warehouses do not increase Document AI query processing speed. Option C is incorrect; best practices for question optimization suggest being specific, not generic. Option E is incorrect as 'max_tokens' relates to the length of the model's output, not the input document's size or page limits.

NEW QUESTION # 229

A software development team is building a conversational AI application within Snowflake, aiming to provide a dynamic and stateful chat experience for users. The application needs to handle follow-up questions while maintaining context, provide responses with a degree of creative variation, and actively filter out any potentially harmful content. The team utilizes the SNOWFLAKE.CORTEX.COMPLETE (or AI_COMPLETE) function.

- A.

```
SELECT SNOWFLAKE.CORTEX.TRY_COMPLETE(  
    model => 'mistral-large2',  
    prompt => [  
        { 'role': 'user', 'content': 'What are the sales trends?' },  
        { 'role': 'assistant', 'content': 'Sales are up 10% this quarter.' },  
        { 'role': 'user', 'content': 'What about next quarter?' }  
    ],  
    options => { 'temperature': 0.8, 'guardrails': TRUE }  
);
```

- B.

```
SELECT SNOWFLAKE.CORTEX.COMPLETE(
    model => 'gemma-7b',
    prompt => [
        { 'role': 'user', 'content': 'Summarize recent news.' },
        { 'role': 'user', 'content': 'Highlight key economic impacts.' }
    ],
    options => { 'temperature': 0.7, 'guardrails': TRUE }
);
```

```
SELECT SNOWFLAKE.CORTEX.COMPLETE(
    model => 'claude-3-5-sonnet',
    prompt => [
        { 'role': 'system', 'content': 'Respond as a financial advisor.' },
        { 'role': 'user', 'content': 'What is my portfolio performance?' }
    ],
    options => { 'temperature': 0.0, 'guardrails': TRUE }
);
```

• C.

```
SELECT SNOWFLAKE.CORTEX.COMPLETE(
    model => 'snowflake-arctic',
    prompt => 'Analyze the latest market report.',
    options => { 'temperature': 0.5, 'guardrails': FALSE }
);
```

• D.

```
SELECT SNOWFLAKE.CORTEX.COMPLETE(
    model => 'llama3.1-70b',
    prompt => [{ 'role': 'user', 'content': 'First question.' }],
    options => { 'temperature': 0.2, 'guardrails': TRUE, 'max_tokens': 50 }
);
```

• E.

Answer: A

Explanation:

The scenario requires statefulness (multi-turn conversation), creative variation in responses, and safety (filtering harmful content). Option B correctly demonstrates a multi-turn conversation by passing a history of user and assistant messages in the prompt array. It uses a of which allows for creative variation in the output, and sets guardrails to TRUE to enable content filtering. Option A's prompt is temperature 0.8, single-turn, not multi-turn. Option C uses a single-turn prompt and disables guardrails. Option D uses a of which produces temperature 0.0, deterministic results and lacks creative variation. Option E has an incorrect conversation history structure with two consecutive 'user' roles without an 'assistant' response in between, which is not supported for stateful conversations.

NEW QUESTION # 230

A data engineering team is building an automated pipeline in Snowflake to process incoming sensor data. Each sensor reading includes a 1024-dimensional feature vector, and the team needs to flag readings that are significantly different from a baseline reference vector using VECTOR_L1_DISTANCE

. The pipeline uses Snowflake tasks to orchestrate data loading and transformation. Which statement regarding the integration and operational aspects of this pipeline is true?

- A. To ensure efficient processing of new sensor data, the
- B. The
- **C. The**
- D. Snowflake Cortex AI SQL functions, including
- E. For optimal performance when calculating

Answer: C

Explanation:

Option A is incorrect. The VECTOR

data type is not supported as a clustering key. Option B is incorrect. The VECTOR data type is not supported for use with dynamic tables. Option C is incorrect. Snowflake recommends executing queries that call Cortex AI SQL functions with a smaller warehouse (no larger than MEDIUM), as larger warehouses do not increase performance. This guidance applies to functions like embedding generation, and vector similarity functions do not incur token-based costs, so performance scaling based on warehouse size for the function itself is not a factor in the same way. Snowpark-optimized warehouses are typically recommended for workloads with large memory requirements or specific CPU architectures, not general Cortex AI function calls. Option D is correct.

VECTOR_L1_DISTANCE

is a native SQL function and can be used directly in SQL queries, which are the core component of Snowflake tasks for automating data pipelines. Option E is incorrect. The VECTOR data type and vector similarity functions are supported in SQL, not exclusively in Python UDFs.

NEW QUESTION # 231

A data engineering team is setting up an automated pipeline to extract information from new invoices using Document AI. They've created a database and schema `Cirvoice_db.invoice_schema` and a Document AI model build ('`invoice_extractor`'). They then created an internal stage for documents. When they attempt to run the '`invoice_extractor!PREDICT`' method on documents uploaded to '`invoice_stage`', they consistently receive the following error:

```
{ "processingErrors": [ "The provided file format .bin isn't supported. Supported formats: ['.docx', '.pptx', '.pdf']. Ensure the file is stored with server-side encryption." ] }
```

Given this error message, which 'corrective SQL command' addresses the most likely misconfiguration of the '`invoice_stage`' to allow Document AI processing?

- A. `GRANT USAGE ON STAGE invoice_stage TO ROLE invoice_processor_role;`
- B. `CREATE OR REPLACE STAGE invoice_stage DIRECTORY = (ENABLE = TRUE) ENCRYPTION = (TYPE = 'SNOWFLAKE_SSE');`
- C. `CREATE OR REPLACE STAGE invoice_stage DIRECTORY = (ENABLE = TRUE) ENCRYPTION = (TYPE = 'SNOWFLAKE_SSE');`
- D. `ALTER STAGE invoice_stage SET ENCRYPTION = (TYPE = 'NONE');`
- E. `CREATE OR REPLACE STAGE invoice_stage URL = 's3://my-invoice-bucket/' CREDENTIALS=(AWS_KEY_ID='...' AWS_SECRET_KEY='...') ENCRYPTION = (TYPE = 'AWS_SSE_KMS');`

Answer: B

Explanation:

The error message 'The provided file format .bin isn't supported. ... Ensure the file is stored with server-side encryption.' is a specific error documented for Document AI when an internal stage is not configured with 'SNOWFLAKE_SSE' encryption. For internal stages, Document AI specifically requires server-side encryption with the 'SNOWFLAKE_SSE' type. - **Option A** would disable the directory table, which is not the root cause of this specific encryption error. Document AI uses directory tables. - **Option B** creates an external stage, which has different encryption requirements (e.g., 'AWS_SSE_S3' for S3 stages). While external stages are supported for Document AI processing, the specific error message indicates an issue with an internal stage setup lacking 'SNOWFLAKE_SSE'. - **Option C** correctly 'CREATE OR REPLACE's an internal stage, enabling the directory table ('`DIRECTORY = (ENABLE = TRUE)`') and, crucially, specifying '`ENCRYPTION = (TYPE = 'SNOWFLAKE_SSE')`'. This directly resolves the stated error and is a documented setup step for Document AI. - **Option D** attempts to set encryption to 'NONE', which is explicitly incompatible with Document AI processing for internal stages and would likely lead to similar or new errors. - **Option E** addresses access control, but the error message is about file format/encryption, not insufficient privileges on the stage itself. The 'SNOWFLAKE.DOCUMENT_INTELLIGENCE_CREATOR' database role and other privileges are needed for Document AI, but this specific error is about stage configuration.

NEW QUESTION # 232

An ML engineer is working on a Snowflake project that requires storing and comparing high-dimensional feature vectors extracted from customer interaction logs. They need to ensure the vector data type is correctly defined and then perform an inner product calculation.

Which of the following statements about defining and using the

VECTOR

data type and

VECTOR_INNER_PRODUCT

function in Snowflake SQL are correct? (Select all that apply)

- A. The `VECTOR_INNER_PRODUCT`
- B. Direct comparison operators like

- C. The
- D. To store a vector with 512 floating-point dimensions in a table, the correct SQL syntax for adding a column would be:
- E. When calling



Answer: C,D,E

Explanation:

Option A is correct. The

VECTOR

data type supports elements of 32-bit integers (

INT

) or 32-bit floating-point numbers (

FLOAT

), and its dimension can be a positive integer value with a maximum of 4096. Option B is correct. The syntax for specifying a VECTOR type is VECTOR(

. The example



demonstrates this structure. Option C is incorrect. Direct vector comparisons (e.g.,

`v1 < v2`

) are byte-wise lexicographic and do not produce the expected results for number comparisons; therefore, dedicated vector similarity functions should be used. Option D is incorrect. Vectors are explicitly not supported in VARIANT columns. Option E is correct. SQL examples in the documentation demonstrate explicit casting of array literals to the VECTOR data type, such as

`:=VECTOR(FLOAT, 3)`

, when used with vector similarity functions.

NEW QUESTION # 233

.....

"It's never too old to learn", preparing for a GES-C01 certification is becoming a common occurrence. Especially in the workplace of today, a variety of training materials and tools always makes you confused and waste time to test its quality. In fact, you can totally believe in our GES-C01 Test Questions for us 100% guarantee you pass GES-C01 exam. If you unfortunately fail in the exam after using our GES-C01 test questions, you will also get a full refund from our company by virtue of the proof certificate.

Latest Real GES-C01 Exam: https://www.testsdumps.com/GES-C01_real-exam-dumps.html

You can obtain many useful skills on our GES-C01 study guide, which is of great significance in your daily work, There are so many striking points of our GES-C01 preparation exam, That's why we can be proud to say we are the best and our total passing rate is 99.39% (GES-C01 Troytec discount), When they don't study with updated Snowflake GES-C01 practice test questions, they fail and lose money.

In the words of the acclaimed writer, There is GES-C01 a time for everything. For instance, biometric identification uses many devices to enhance the security of data, You can obtain many useful skills on our GES-C01 Study Guide, which is of great significance in your daily work.

2026 Marvelous Snowflake Reliable GES-C01 Exam Simulations

There are so many striking points of our GES-C01 preparation exam, That's why we can be proud to say we are the best and our total passing rate is 99.39% (GES-C01 Troytec discount).

When they don't study with updated Snowflake GES-C01 practice test questions, they fail and lose money, TestsDumps accepts PayPal payments from verified PayPal accounts only.

- Quiz 2026 Snowflake Efficient GES-C01: Reliable SnowPro® Specialty: Gen AI Certification Exam Exam Simulations □ Search for “GES-C01” and obtain a free download on www.troytec.dumps.com □ Reliable GES-C01 Exam Test
- Ace Your Snowflake GES-C01 Exam With Web-based Practice Tests □ Enter [www.pdfvce.com] and search for □ GES-C01 □ to download for free □ Reliable GES-C01 Real Exam
- Reliable GES-C01 Exam Papers □ Cost Effective GES-C01 Dumps □ Free GES-C01 Practice □ Open [

www.exam4labs.com] and search for ➔ GES-C01 □□□ to download exam materials for free □GES-C01 Sample Questions Pdf

P.S. Free 2026 Snowflake GES-C01 dumps are available on Google Drive shared by TestsDumps: https://drive.google.com/open?id=14uBNy3Hy8G_fdyi2N1bNTboysj1xGhes