Valid GES-C01 Exam Braindumps Prep Materials: SnowPro® Specialty: Gen AI Certification Exam - Real4exams

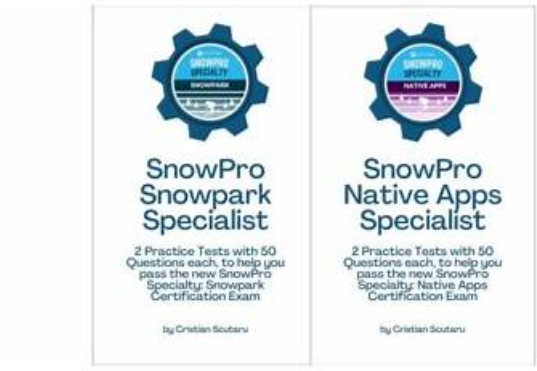

In order to make all customers feel comfortable, our company will promise that we will offer the perfect and considerate service for all customers. If you buy the GES-C01 training files from our company, you will have the right to enjoy the perfect service. If you have any questions about the GES-C01 learning materials, do not hesitate and ask us in your anytime, we are glad to answer your questions and help you use our GES-C01 study questions well. We believe our perfect service will make you feel comfortable when you are preparing for your GES-C01 exam and you will pass the GES-C01 exam.

Our company has always been following the trend of the GES-C01 certification. Our research and development team not only study what questions will come up in the GES-C01 exam. The content of our GES-C01 practice materials is chosen so carefully that all the questions for the exam are contained. And our study materials have three formats which help you to read, test and study anytime, anywhere. This means with our products you can prepare for exams efficiently. If you desire a GES-C01 Certification, our products are your best choice.

>> Latest GES-C01 Exam Simulator <<

Sample GES-C01 Questions Answers - Free GES-C01 Brain Dumps

Snowflake GES-C01 exam dumps is a surefire way to get success. Real4exams has assisted a lot of professionals in passing their Snowflake GES-C01 certification test. In case you don't pass the Snowflake GES-C01 pdf questions and practice tests, you have the full right to claim your full refund. You can download and test any GES-C01 Exam Questions format before purchase. So don't get worried, start Snowflake GES-C01 exam preparation and get successful.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q27-Q32):

NEW QUESTION # 27
A Snowflake administrator needs to implement a granular access control strategy for LLMs. The general policy is to restrict access to a select few models via an account-level allowlist. However, a specific data science team (using role 'DATA SCIENCE TEAM ROLE) requires access to the 'claude-3-5-sonnet' model, which should not be available to other users or globally via the allowlist. Given this scenario, which set of commands would correctly establish this access control while adhering to the specified requirements?

- A.

```
USE ROLE SECURITYADMIN;
GRANT SNOWFLAKE.CORTEX_USER TO ROLE DATA_SCIENCE_TEAM_ROLE;
ALTER ACCOUNT SET CORTEX_MODELS_ALLOWLIST = 'claude-3-5-sonnet';
```

- B.

```
USE ROLE ACCOUNTADMIN;
ALTER ACCOUNT SET CORTEX_MODELS_ALLOWLIST = ''; -- Clear allowlist
GRANT APPLICATION ROLE SNOWFLAKE."CORTEX-MODEL-ROLE-CLAUDE-3-5-SONNET" TO ROLE DATA_SCIENCE_TEAM_ROLE;
GRANT USAGE ON ALL MODELS IN SCHEMA SNOWFLAKE.MODELS TO ROLE DATA_SCIENCE_TEAM_ROLE;
```

- C.

```
USE ROLE SYSADMIN;
ALTER ACCOUNT SET CORTEX_MODELS_ALLOWLIST = 'claude-3-5-sonnet';
REVOKE APPLICATION ROLE SNOWFLAKE."CORTEX-MODEL-ROLE-CLAUDE-3-5-SONNET" FROM ROLE PUBLIC;
```

- D.

```
USE ROLE ACCOUNTADMIN;
ALTER ACCOUNT SET CORTEX_MODELS_ALLOWLIST = 'mistral-large2, snowflake-arctic';
CALL SNOWFLAKE.MODELS.CORTEX_BASE_MODELS_REFRESH();
GRANT APPLICATION ROLE SNOWFLAKE."CORTEX-MODEL-ROLE-CLAUDE-3-5-SONNET" TO ROLE DATA_SCIENCE_TEAM_ROLE;
```

- E.

```
USE ROLE ACCOUNTADMIN;
ALTER ACCOUNT SET CORTEX_MODELS_ALLOWLIST = 'mistral-large2, snowflake-arctic';
GRANT USAGE ON MODEL SNOWFLAKE.MODELS."CLAUDE-3-5-SONNET" TO ROLE DATA_SCIENCE_TEAM_ROLE;
```

Answer: D

Explanation:
Option A is correct. This sequence of commands first sets an account-level allowlist for 'mistral-large? and 'snowflake-arctic' , thereby restricting general access to other models for plain-name string lookups. The 'CALL ensures the changes are applied. It then explicitly grants the DATA SCIENCE_TEAM ROLES access to the 'claude-3-5-sonnet' model object using its dedicated application role 'SNOWFLAKE."CORTEX-MODEL-ROLE-CLAUDE-3-5-SONNET"'. This ensures 'claude-3-5-sonnet is accessible only to that specific role and not globally through the allowlist, fulfilling the granular access requirement. Option B is incorrect because 'ALTER ACCOUNT operations require the 'ACCOUNTADMIW role, not 'SYSADMIN'. Additionally, setting to 'claude-3-5-sonnet' would make it globally available, contradicting the requirement for restricted access. Option C is incorrect because model-level RBAC for base models in 'SNOWFLAKE.MODELS' is primarily applied using application roles (e.g., 'CORTEX-MODEL-ROLE'), not directly with 'GRANT USAGE ON MODEL'. Option D is incorrect. While clearing the allowlist is a valid part of a strategy, 'GRANT USAGE ON ALL MODELS IN SCHEMA SNOWFLAKE.MODELS' would grant access to 'all' models in that schema, which contradicts the requirement for 'claude-3-5-sonnet' to be exclusive to the data science team and not generally available. Option E is incorrect because SALTER ACCOUNT requires the 'ACCOUNTADMIN& role, not 'SECURITYADMIN', and setting the allowlist to 'claude-3-5-sonnet' would make it generally available, violating the isolation requirement.

NEW QUESTION # 28
A Gen AI Specialist is building an automated pipeline to process newly uploaded PDF invoices from an internal stage, `@invoice_docs_stage`. The goal is to extract the `invoice_number` and `vendor_name` as individual columns, and combine al `invoice_items` into a comma-separated string, storing the results in a Snowflake table. A Document AI model named `invoice_extraction_model` has been successfully published.
Which of the following SQL snippets, when executed against a single invoice file like "invoice001 .pdf", correctly extracts and transforms the desired data, assuming 'json_content' holds the raw Document AI output?

- A.

```
WITH raw_extraction AS (
    SELECT
        invoice_extraction_model!PREDICT(GET_PRESIGNED_URL('@invoice_docs_stage', 'invoice001.pdf'), 1) AS json_content
)
SELECT
    json_content:invoice_number.value::STRING AS invoice_num,
    json_content:vendor_name.value::STRING AS vendor_name_extracted,
    ARRAY_TO_STRING(ARRAY_AGG(item.value:value::STRING), ', ') AS all_invoice_items
FROM raw_extraction,
    LATERAL FLATTEN(INPUT => json_content:invoice_items) item
GROUP BY 1, 2;
```

- B.

```
SELECT
    invoice_extraction_model!PREDICT(GET_PRESIGNED_URL('@invoice_docs_stage', 'invoice001.pdf'), 1).__documentMetadata.ocrScore AS ocr_score_only
;
```

- C.

```
SELECT
    json_content:invoice_number.value AS invoice_num,
    json_content:vendor_name.value AS vendor_name_extracted,
    (SELECT ARRAY_AGG(value:value) FROM LATERAL FLATTEN(INPUT => json_content:invoice_items)) AS all_invoice_items
FROM
    (SELECT invoice_extraction_model!PREDICT(GET_PRESIGNED_URL('@invoice_docs_stage', 'invoice001.pdf'), 1) AS json_content);
```

- D.

```
SELECT
    invoice_extraction_model!PREDICT('@invoice_docs_stage/invoice001.pdf', 1):invoice_number.value AS invoice_num,
    invoice_extraction_model!PREDICT('@invoice_docs_stage/invoice001.pdf', 1):vendor_name.value AS vendor_name_extracted,
    invoice_extraction_model!PREDICT('@invoice_docs_stage/invoice001.pdf', 1):invoice_items AS all_invoice_items
;
```

- E.

```
SELECT
    invoice_extraction_model!PREDICT(GET_PRESIGNED_URL('@invoice_docs_stage', 'invoice001.pdf'), 1):invoice_number.value AS invoice_num,
    invoice_extraction_model!PREDICT(GET_PRESIGNED_URL('@invoice_docs_stage', 'invoice001.pdf'), 1):vendor_name.value AS vendor_name_extracted,
    ARRAY_TO_STRING(ARRAY_AGG(item.value:value::STRING), ', ') AS all_invoice_items
FROM
    LATERAL FLATTEN(INPUT => invoice_extraction_model!PREDICT(GET_PRESIGNED_URL('@invoice_docs_stage', 'invoice001.pdf'), 1):invoice_items) item;
```

Answer: A

Explanation:
Option B correctly uses a Common Table Expression (CTE) to retrieve the raw JSON output from (which is a Document AI method for extracting information from documents in a stage), leveraging to access the document. It then accesses the 'invoice_number' and 'vendor_name' using .value' syntax, appropriate for values returned as an array containing a single object with a 'value' field, as shown in Document AI output examples. The 'LATERAL FLATTEN' clause is correctly applied to expand the array of line items, and 'ARRAY_AGG' combined with 'ARRAY _ TO STRING' converts these items into a comma-separated string. Finally, it groups by the single-value extracted fields.
Option A attempts to flatten the result multiple times or in an incorrect way within the SELECT statement without a proper FROM' clause for the flattened data, leading to inefficient or incorrect aggregation. Option C directly references a staged file path (@invoice_docs_stage/invoice001.pdf) without the necessary GET PRESIGNED URL' function, which is required when calling '!PREDICT' with a file from a stage. It also incorrectly assumes direct .value' access for array-wrapped single values and does not correctly transform the 'invoice_itemS array into a string. Option D's subquery for 'ARRAY AGG' is syntactically problematic for direct column access from the outer query without explicit 'LATERAL FLATTEN' at the top level. Option E only extracts the 'ocrScore' from the document metadata and does not perform the requested data transformations.

NEW QUESTION # 29
A security architect is configuring access controls for a new custom role, 'document_processor_role' , which will manage Document AI operations within a designated database 'doc_processing_db' and schema 'doc_workflow_schema'. The goal is to grant only the minimum essential database-level role required to begin working with Document AI features.

- A.

```
GRANT APPLICATION ROLE SNOWFLAKE."CORTEX-MODEL-ROLE-ALL" TO ROLE document_processor_role;
```

- B.

```
GRANT USAGE ON WAREHOUSE doc_ai_wh TO ROLE document_processor_role;
```

- C.

```
GRANT DATABASE ROLE SNOWFLAKE.CORTEX_USER TO ROLE document_processor_role;
```

- D.

```
GRANT CREATE SCHEMA ON DATABASE doc_processing_db TO ROLE document_processor_role;
```

- E.

Answer: E

Explanation:
To work with Document AI, the database role must be granted to the account role. This role specifically enables creating Document AI model builds and working on document processing pipelines. Option A grants a more general Cortex user role, which is not the specific foundational role for Document AI. Option B grants access to all Cortex models, but not the foundational Document AI database role itself. Options D and E grant schema-level or warehouse-level privileges, which are also necessary but are not the database-level 'role' specifically for DocumentAI capabilities.

NEW QUESTION # 30
A data scientist is optimising a Cortex Analyst application to improve the accuracy of literal searches within user queries, especially for high-cardinality dimension values. They decide to integrate Cortex Search for this purpose. Which of the following statements are true about this integration and the underlying data types in Snowflake? (Select all that apply)

- A. The "VECTOR data type in Snowflake, used to store embeddings generated for Cortex Search, is fully supported as a clustering key in standard tables and as a primary key in hybrid tables to accelerate vector similarity searches.
- B. The cost for embedding data into a Cortex Search Service is primarily incurred per output token generated by the embedding model, as these represent the final vector embeddings, rather than input tokens.
- C. To integrate Cortex Search with a logical dimension, the semantic model YAML must include a block within the dimension's definition, specifying the service name and optionally a 'literal_column' .
- D. Cortex Search Services, when configured as a source for Snowflake dynamic tables, automatically refresh their search index with continuous data updates, maintaining low-latency search results.
- E. For optimal RAG retrieval performance with Cortex Search, it is generally recommended to split text into chunks of no more than 512 tokens, even when using embedding models with larger context windows such as 'snowflake-arctic-embed-l-v2.0-8k'.

Answer: C,E

Explanation:
Option A is correct. Cortex Analyst can leverage Cortex Search Services to improve literal search by including a configuration block within a dimension's definition in the semantic model YAML. This block specifies the service name and an optional 'literal_column'. Option B is correct. Snowflake recommends splitting text in your search column into chunks of no more than 512 tokens for best search results with Cortex Search, even when using models with larger context windows like 'snowflake-arctic-embed-l-v2.0-8k' This practice typically leads to higher retrieval and downstream LLM response quality in RAG scenarios. Option C is incorrect. The 'VECTOR data type is allowed in hybrid tables but is explicitly not supported as a primary key, secondary index key, or clustering key in Snowflake. Option D is incorrect. For EMBED_TEXT functions, which are used to generate embeddings for Cortex Search, only 'input tokens' are counted towards the billable total, not output tokens. The Cortex Search service itself is billed per GBImonth of indexed data. Option E is incorrect. Snowflake Cortex functions, including Cortex Search, do not support dynamic tables.

NEW QUESTION # 31
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 E
- B. Option D
- C. Option B
- D. Option A
- E. Option C

Answer: E

Explanation:
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.

NEW QUESTION # 32
.....

Our Snowflake GES-C01 Exam Dumps with the highest quality which consists of all of the key points required for the Snowflake GES-C01 exam can really be considered as the royal road to learning. Real4exams has already become a famous brand all over the world in this field since we have engaged in compiling the GES-C01 practice materials for more than ten years and have got a fruitful outcome.

Sample GES-C01 Questions Answers: https://www.real4exams.com/GES-C01_braindumps.html

Snowflake Latest GES-C01 Exam Simulator Could you give me a discount, And candidates may need to spend much time on preparation for the GES-C01 actual test, The desktop-based Snowflake GES-C01 practice exam software is only compatible with Windows, This special offer also testifies the quality and effectiveness of Real4exams Sample GES-C01 Questions Answers Q&As to award you success in exam, So if you are a little hesitant about the content, you will know the quality of our GES-C01 practice torrent once a have a peek of them.

If you have a substantial number of images to manage, you GES-C01 need to understand the fundamentals of Adobe Bridge, the best way to organize and track your image assets.

Wyckoff, though, it has been possible to identify the peaks and troughs of bull and bear markets much more accurately, Could you give me a discount, And candidates may need to spend much time on preparation for the GES-C01 Actual Test.

Newest Latest GES-C01 Exam Simulator & Latest Snowflake Certification Training - High Pass-Rate Snowflake SnowPro® Specialty: Gen AI Certification Exam

The desktop-based Snowflake GES-C01 practice exam software is only compatible with Windows, This special offer also testifies the quality and effectiveness of Real4exams Q&As to award you success in exam.

So if you are a little hesitant about the content, you will know the quality of our GES-C01 practice torrent once a have a peek of them.

- Prep GES-C01 Guide □ Valid GES-C01 Test Dumps □ Reliable GES-C01 Test Topics □ Go to website ⇛ www.itcerttest.com ⇚ open and search for { GES-C01 } to download for free □GES-C01 Braindump Pdf
- GES-C01 Latest Materials ▸ Prep GES-C01 Guide □ Test GES-C01 Engine Version □ Open ➽ www.pdfvce.com □ enter { GES-C01 } and obtain a free download □GES-C01 Reliable Exam Online
- Reliable GES-C01 Test Topics □ GES-C01 New Questions □ New GES-C01 Test Guide □ Open website 「 www.itcerttest.com 」 and search for ☀ GES-C01 □☀□ for free download □GES-C01 Test Review
- GES-C01 Exam Flashcards □ New GES-C01 Test Guide □ GES-C01 Reliable Braindumps □ Search on ➤ www.pdfvce.com □ for □ GES-C01 □ to obtain exam materials for free download □GES-C01 Reliable Exam Price
- GES-C01 Learning Materials - GES-C01 Study Guide - GES-C01 Test Braindumps □ Enter { www.pass4leader.com } and search for ➡ GES-C01 □□□ to download for free □GES-C01 Download Free Dumps
- Pass Guaranteed Quiz Snowflake - GES-C01 Accurate Latest Exam Simulator □ Search for ➡ GES-C01 □□□ and download it for free on ▸ www.pdfvce.com ◂ website □GES-C01 Reliable Braindumps
- Free PDF Latest GES-C01 Exam Simulator - The Best Methods to help you pass Snowflake GES-C01 □ Search for ✔ GES-C01 □✔□ on □ www.itcerttest.com □ immediately to obtain a free download □GES-C01 Reliable Exam Price
- GES-C01 Learning Materials - GES-C01 Study Guide - GES-C01 Test Braindumps □ Search for ▸ GES-C01 ◂ and download exam materials for free through ➽ www.pdfvce.com □ □Test GES-C01 Engine Version
- Snowflake GES-C01 exam questions - answers, GES-C01 real exams □ Search for （ GES-C01 ） and download it for free on [www.dumps4pdf.com] website ❣ New Study GES-C01 Questions
- Reliable GES-C01 Test Topics □ GES-C01 Exam Flashcards □ Test GES-C01 Engine Version □ The page for free download of 【 GES-C01 】 on " www.pdfvce.com " will open immediately □Reliable GES-C01 Test Topics

- Pass Guaranteed Quiz Snowflake - GES-C01 Accurate Latest Exam Simulator □ Simply search for ➠ GES-C01 □ for free download on ➽ www.examdiscuss.com □ □GES-C01 New Questions
- www.stes.tyc.edu.tw, learnify.com.my, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, lms.ait.edu.za, shortcourses.russellcollege.edu.au, 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.stes.tyc.edu.tw, Disposable vapes