

100% Pass Quiz 2025 Snowflake GES-C01: Authoritative Valid SnowPro® Specialty: Gen AI Certification Exam Exam Dumps



Will you feel nervous in the exam? If you do, just choose us, our GES-C01 Soft test engine can stimulate the real exam environment, which will help you know the procedure of the exam, and will strengthen your confidence. Moreover GES-C01 exam dumps are high-quality, and we have professional experts to compile them, and they can help you pass the exam just one time. We offer you free demo to have a try for GES-C01 Exam Dumps, and free update for one year. If you indeed have questions, just contact with us.

This will help them polish their skills and clear all their doubts. Also, you must note down your SnowPro® Specialty: Gen AI Certification Exam (GES-C01) practice test score every time you try the Snowflake Exam Questions. It will help you keep a record of your study and how well you are doing in them. ITExamDownload hires the top industry experts to draft the SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam dumps and help the candidates to clear their SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam easily. ITExamDownload plays a vital role in their journey to get the GES-C01 certification.

>> Valid GES-C01 Exam Dumps <<

Free Download Valid GES-C01 Exam Dumps | Valid GES-C01 Test Braindumps: SnowPro® Specialty: Gen AI Certification Exam

Snowflake GES-C01 Exam provided by ITExamDownload is of the highest quality, and it enables participants to pass the exam on their first try. For successful preparation, it is essential to have good Snowflake GES-C01 exam dumps and to prepare questions that may come up in the exam. ITExamDownload helps candidates overcome all the difficulties they may encounter in their exam preparation. To ensure the candidates' satisfaction, ITExamDownload has a support team that is available 24/7 to assist with a wide range of issues.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q246-Q251):

NEW QUESTION # 246

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 all '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.

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

- B.

```
WITH raw_extraction AS (
  SELECT
    invoice_extraction_model!PREDICT(GET_PREIGNED_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:~:STRING), ', ') AS all_invoice_items
FROM raw_extraction,
  LATERAL FLATTEN(INPUT => json_content:invoice_items) item
GROUP BY 1, 2;
```

- C.

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

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

Answer: B

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 PREIGNED 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 # 247

A team is planning the implementation of a new Document AI solution and needs to be aware of the specific guidelines and limitations concerning naming conventions and task management within Snowflake. A primary concern is to avoid common pitfalls that could lead to errors or unsupported configurations.

- A. While optional, creating a separate database and schema specifically for Document AI assets is a best practice for organization and cost visibility.
- B. Database and schema identifiers referenced in Document AI operations must not include double quotes, as this is an unsupported syntax.
- C. Document AI supports concurrent user activity on the same model build in Snowsight, enabling multiple team members to upload documents and review answers simultaneously.
- D. Document AI model builds can be renamed after their initial creation and publication to align with evolving project naming

standards.

- E. For optimal cost management and to avoid resource contention, Document AI tasks should leverage serverless task configurations.

Answer: A,B

Explanation:

Option A is incorrect because DocumentAI does not currently support renaming models. Option B is incorrect as Document AI does not support serverless tasks. Option C is correct as Document AI does not support double quotes around identifiers for the database and schema. Option D is correct as Snowflake recommends creating a separate database and schema for Document AI to help track costs and manage objects. Option E is incorrect because Document AI does not support multiple users working on the same model build at the same time in Snowsight.

NEW QUESTION # 248

A data analytics team is building a self-service analytics application using Snowflake Cortex Analyst to allow business users to query sales data with natural language. They are defining a semantic model in YAML to ensure accurate text-to-SQL generation. Which of the following is the most crucial aspect of the semantic model's configuration for Cortex Analyst to effectively translate natural language into SQL for structured data?

- A. Defining a comprehensive 'verified_queries' section with a high volume of example natural language questions and their exact SQL translations to handle all potential user queries.
- B. Specifying a dedicated 'CORTEX SEARCH SERVICE' for every dimension to pre-compute all possible literal values, optimizing response time.
- C. Configuring the 'base_table' parameter to directly reference a dynamic table, ensuring real-time data ingestion and processing before SQL generation.
- D. Utilizing advanced data types like 'VARIANT' and 'OBJECT' for all dimensions to accommodate semi-structured data without complex transformations.
- E. Providing detailed 'name', 'description', and 'synonyms' for logical tables, dimensions, and facts to bridge the gap between business terminology and the underlying database schema.

Answer: E

Explanation:

Option C is correct because the primary purpose of a semantic model in Cortex Analyst is to provide semantic information about your data, bridging the gap between business users' natural language and the technical database schema. This includes using descriptive names, synonyms, and descriptions for logical tables, dimensions, and facts, which is essential for Cortex Analyst to reliably generate accurate SQL from natural language questions. Option A is incorrect; while Cortex Search Services can improve literal matching for dimensions, it's an enhancement and not the most crucial foundational aspect of the semantic model for general text-to-SQL translation, nor is it required for 'every' dimension. Option B is incorrect because while 'verified_queries' improve accuracy for similar questions, a high volume of examples for 'all' potential queries is not feasible or the most crucial initial configuration; the core mapping (Option C) is more fundamental. Option D is incorrect as the 'base_table' must refer to a physical table or view, not directly to a dynamic table. Furthermore, Cortex functions do not support dynamic tables directly. Option E is incorrect because 'VARIANT', 'OBJECT', 'GEOGRAPHY', and 'ARRAY' data types are explicitly not supported for dimension, fact, or metric columns in a semantic model.

NEW QUESTION # 249

A Snowflake administrator is designing a new role, 'doc_ai_pipeline_creator', intended to configure and deploy Document AI extraction pipelines. This role needs the ability to ensure that the designated virtual warehouse, 'analytics_warehouse', can be reliably started and stopped as needed for Document AI tasks. Which of the following SQL statements grants the privilege that directly enables the 'doc ai_pipeline_creator' role to control the operational state of 'analytics_warehouse'?

- A.

```
GRANT USAGE ON WAREHOUSE analytics_warehouse TO ROLE doc_ai_pipeline_creator;
```
- B.

```
GRANT MODIFY ON WAREHOUSE analytics_warehouse TO ROLE doc_ai_pipeline_creator;
```
- C.

```
GRANT CONTROL ON WAREHOUSE analytics_warehouse TO ROLE doc ai pipeline creator;
```
- D.

```
GRANT OPERATE ON WAREHOUSE SUBSTATEMENTS TO ROLE doc ai pipeline creator;
```
- E.

```
GRANT START/STOP ON WAREHOUSE analytics_warehouse TO ROLE doc ai pipeline_creator;
```

Answer: D

Explanation:

For Document AI operations, including setting up model builds and running extraction pipelines, the associated virtual warehouse must be available and its operational state managed effectively. The 'OPERATE' privilege on a virtual warehouse specifically grants the ability to start and stop (resume and suspend) the warehouse. - Option A USAGE ON WAREHOUSE) allows a role to select and use a warehouse but does not provide control over its active state. Both 'USAGE*' and '*OPERATE*' are required for Document AI. - Options B, D, and E CMONITOR, MODIFY, "CONTROL") are not the specific privilege described in the sources for controlling the operational state of a warehouse in the context of Document AI setup.

NEW QUESTION # 250

A data architect is evaluating the shift from managing Cortex Analyst semantic models as YAML files on internal stages to leveraging a native semantic view (currently in Public Preview). They want to understand the key differences and advantages or considerations of this new native approach. Which of the following statements accurately describe a key characteristic or implication of using native semantic views for Cortex Analyst, compared to YAML files stored in a stage?

- ☐ Native semantic views eliminate the need for an underlying `base_table` definition, as the view itself directly defines the logical structure and data source for Cortex Analyst queries.
- ☐ When a semantic model is stored as a native semantic view, the `CORTEX_ANALYST_REQUESTS` function must be called with 'SEMANTIC_VIEW' as the `semantic_model_or_view_type` parameter, along with the view's fully qualified name.
- ☒ Migrating to native semantic views automatically grants `SELECT` privileges to the `SNOWFLAKE_CORTX_ANALYST_USER` database role on all underlying tables referenced by the semantic model, simplifying data access control.
- ☐ Native semantic views offer full support for `VARIANT`, `OBJECT`, `GEOGRAPHY`, and `ARRAY` data types within their dimension and fact columns, overcoming limitations of YAML-based models.
- ☐ Semantic models defined as native semantic views are exclusively managed through Snowflake's Snowsight UI, removing the option for programmatic updates via SQL or the Snowflake CLI for enhanced control.

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

Answer: D

Explanation:

Option B is correct. The function, used for monitoring cortex Analyst activity, requires specifying the semantic model type. This type can be for YAML files or for semantic views, along with the model or view name. Option A is incorrect because a logical table in a semantic model, whether YAML-based or a semantic view, represents an underlying physical database table or a view and requires a 'base_table' definition to specify the data source. Option C is incorrect; while stage access for YAML files is controlled by RBAC, roles granted access to a stage or a semantic view still require explicit 'SELECT' access on all referenced underlying tables, as stage/view access alone does not implicitly grant table access. Option D is incorrect as the 'VARIANT', 'OBJECT', 'GEOGRAPHY', and 'ARRAY' data types are currently not supported for dimension or fact columns in a semantic model, regardless of whether it's stored in YAML or a native view. Option E is incorrect; while Snowsight offers UI tools for semantic model creation and management, Snowflake typically supports programmatic management via SQL or CLI for native database objects, and the sources do not state that native semantic views would remove these options.

NEW QUESTION # 251

.....

ITExamDownload provides a clear and superior solutions for each Snowflake GES-C01 Exam candidates. We provide you with the Snowflake GES-C01 exam questions and answers. Our team of IT experts is the most experienced and qualified. Our test questions and the answer is almost like the real exam. This is really amazing. More importantly, the examination pass rate of ITExamDownload is highest in the worldwide.

GES-C01 Test Braindumps: <https://www.itexamdownload.com/GES-C01-valid-questions.html>

Snowflake Valid GES-C01 Exam Dumps In case that such term or provision cannot be modified accordingly, the provision will be deemed stricken here from, and shall not affect the validity or enforceability of the remainder of these Terms and Conditions, Our GES-C01 Test Braindumps - SnowPro® Specialty: Gen AI Certification Exam valid vce is the best alternative to your time and money to get an excellent career in the IT filed, Snowflake Valid GES-C01 Exam Dumps You will have no regret spending your

Others have to work near the mountains, Select File GES-C01 > Form Properties > Preview, In case that such term or provision cannot be modified accordingly, the provision will be deemed stricken here from, Best GES-C01 Preparation Materials and shall not affect the validity or enforceability of the remainder of these Terms and Conditions.

Our SnowPro® Specialty: Gen AI Certification Exam valid vce is the best alternative to your time and Trustworthy GES-C01 Exam Torrent money to get an excellent career in the IT filed, You will have no regret spending your valuable time on our introduction.

[illegible]