

# GES-C01 Practice Braindumps & GES-C01 Guide Torrent



BONUS!!! Download part of Exams4sures GES-C01 dumps for free: <https://drive.google.com/open?id=1oidJ34Uno-xwR2Z5MdY6w9Xzonv7O1zX>

As a market leader, our company is able to attract quality staffs, it actively seeks out those who are energetic, persistent, and professional to various GES-C01 certificate and good communicator. And we strongly believe that the key of our company's success is its people, skills, knowledge and experience. Over 50% of the account executives and directors have been with the Group for more than ten years. The successful selection, development and GES-C01 training of personnel are critical to our company's ability to provide a high pass rate of GES-C01 exam questions for you to pass the GES-C01 exam.

We are committed to help you pass the exam just one time, so that your energy and time on practicing GES-C01 exam braindumps will be paid off. GES-C01 learning materials are high-quality, and they will help you pass the exam. Moreover, GES-C01 exam braindumps contain both questions and answers, and it's convenient for you to check answers after training. We offer you free update for one year for GES-C01 Training Materials, and the update version will be sent to you automatically. We have online and offline service for GES-C01 exam materials, if you have any questions, don't hesitate to consult us.

>> GES-C01 Practice Braindumps <<

## GES-C01 Guide Torrent | Real GES-C01 Dumps

First and foremost, you can get the latest version of our GES-C01 study materials for free during the whole year. Second, our responsible after sale service staffs are available in twenty four hours a day, seven days a week, so if you have any problem after purchasing GES-C01 study materials, you can contact our after sale service staffs anywhere at any time. Finally, we have installed the most advanced operation machines in our website, so you can use credit for payment in the process of trading and register your personal information under a safe payment environment. Do not waver any more, the most effective and the Latest GES-C01 Study Materials is right here waiting for you.

## Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q307-Q312):

### NEW QUESTION # 307

A financial institution is implementing a Document AI pipeline to extract specific data from loan application forms. A custom role, 'loan\_processor\_role', has been created and granted the 'SNOWFLAKE.DOCUMENT\_INTELLIGENCE\_CREATOR' database role. The development team has successfully created and published a Document AI model build named 'loan\_model'. They are now trying to automate the process by creating a Snowflake task to run the 'loan\_model!PREDICT' function on new documents in an internal stage. However, the 'CREATE TASK' statement consistently fails, reporting a privilege error. Which of the following privileges is most likely missing from the 'loan\_processor\_role'?

- A.  The 'EXECUTE IMMEDIATE' privilege on the account.
- B.  The 'CREATE DATABASE' privilege on the account.
- C.  The 'CREATE SCHEMA' privilege on the account.
- D.  The 'CREATE TASK' privilege on the schema where the task is being created.
- E.  The 'ALTER STAGE' privilege on the stage containing the document.

**Answer: D**

Explanation:

While the database role provides broad access to Document AI features and the ability to work on model builds, creating pipeline components like streams and tasks requires specific object privileges on the relevant database objects. For creating a task, the 'CREATE TASK' privilege on the schema where the task is being created is explicitly required. Option A ('CREATE DATABASE') is a much higher-level privilege and not directly relevant to task creation within an existing schema. Option B ('USAGE' on warehouse) is necessary for the task to use that warehouse, but not for the 'CREATE TASK' statement itself to succeed. Option C ('EXECUTE IMMEDIATE') is not a standard privilege for creating tasks. Option E ('ALTER STAGE') is for modifying stage properties, not for creating tasks.

#### NEW QUESTION # 308

A data engineering manager needs to audit Cortex LLM function costs to identify specific SQL queries that are unexpectedly high in token consumption for the 'llama3.1-8b' model. They require granular analysis of prompt, completion, and guardrail token usage for these queries. Which of the following Snowflake methods or views would provide the necessary insights?

- Querying the `SNOWFLAKE.ORGANIZATION_USAGE.METERING_DAILY_HISTORY` view, filtered by `SERVICE_TYPE = 'AI_SERVICES'`, to identify daily aggregated credit usage for all AI services.
- Examining the `SNOWFLAKE.ACCOUNT_USAGE.CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY` view, specifically filtering by `MODEL_NAME = 'llama3.1-8b'` and analyzing the `PROMPT_TOKENS`, `COMPLETION_TOKENS`, `GUARD_TOKENS`, and `QUERY_ID` columns.
- Utilizing the `SNOWFLAKE.CORTEX.COUNT_TOKENS` function to estimate token counts for input prompts before any LLM function execution, comparing these estimates to actual spend.
- For SQL queries using `AI_COMPLETE` or `COMPLETE` with `show_details => TRUE`, inspecting the `usage` object within the returned JSON to retrieve `prompt_tokens`, `completion_tokens`, and `guard_tokens` for each individual call.
- Analyzing the `SNOWFLAKE.ACCOUNT_USAGE.CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY` view, filtering by the specific Document AI model builds associated with the 'llama3.1-8b' model.

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

**Answer: C,D**

Explanation:

Option B is correct because the `SNOWFLAKE.ACCOUNT_USAGE.CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY` view provides granular usage information for individual Cortex LLM function calls, including `'PROMPT_TOKENS'`, `'COMPLETION_TOKENS'`, `'GUARD_TOKENS'`, and the `'QUERY_ID'` for specific queries and `'MODEL_NAME'`. This directly addresses the need to audit token consumption for a specific model and identify high-usage queries. Option D is also correct as the `'COMPLETE'` and `'AI_COMPLETE'` functions, when called with `'show_details => TRUE'` (available via the `'options'` argument for `'COMPLETE'`), return a JSON object that includes a `'usage'` key with `'prompt_tokens'`, `'completion_tokens'`, and `'guard_tokens'` details for that specific invocation. This provides per-call details directly at the point of execution. Option A is incorrect because `'METERING_DAILY_HISTORY'` provides aggregated daily credit usage for all AI services, not granular token counts per query or model. Option C is incorrect as `'COUNT_TOKENS'` is used for estimating token counts `'before'` execution to avoid exceeding limits or for cost planning, not for tracking `*actual*` historical usage. Option E is incorrect because `CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY` tracks Document AI processing functions like `'!PREDICT'` and `'PARSE_DOCUMENT'`, and aggregates pages processed and credits used, not granular token counts for general LLM functions like `'COMPILE'`. The 'llama3.1-8b' model is an LLM available for `'AI_COMPILE'` / `'COMPILE'`.

### NEW QUESTION # 309

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 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;
```

- B.

```
 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;
```

```
 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;
```

- C. GRANT USAGE ON ALL MODELS IN SCHEMA SNOWFLAKE.MODELS TO ROLE DATA\_SCIENCE\_TEAM\_ROLE;

- D.

```
 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;
```

```
 USE ROLE SECURITYADMIN;
GRANT SNOWFLAKE.CORTEX_USER TO ROLE DATA_SCIENCE_TEAM_ROLE;
```

- E.  ALTER ACCOUNT SET CORTEX\_MODELS\_ALLOWLIST = 'claude-3-5-sonnet';

### Answer: B

Explanation:

Option A is correct. This sequence of commands sets an account-level allowlist for 'mistral-large2' and 'snowflake-arctic' (restricting general access), and then explicitly grants the access to the 'claude-3-5-sonnet' model object using its dedicated application role 'SNOWFLAKE.CORTEX-MODEL-ROLE-CLAUDE-3-5-SONNET'. The call is often necessary after 'ALTER ACCOUNT SET' to ensure the changes are applied. Option B is incorrect because 'ALTER ACCOUNT' requires the 'ACCOUNTADMIN' role, and setting to 'claude-3-5-sonnet' would make it globally available, contradicting the requirement. 'REVOKE APPLICATION ROLE' from 'PUBLIC' is not the correct mechanism to grant specific access. Option C is incorrect because model-level RBAC for base models in 'SNOWFLAKE.MODELS' is primarily applied using application roles like 'CORTEX-MODEL-ROLE-', not directly with 'GRANT USAGE ON MODEL'. Option D is incorrect. While clearing the allowlist is possible, '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 'ALTER ACCOUNT' requires the 'ACCOUNTADMIN' role, and setting the allowlist to 'claude-3-5-sonnet' would make it generally available, violating the isolation requirement.

### NEW QUESTION # 310

A security administrator is implementing strict model access controls for Snowflake Cortex LLM functions, including those accessed via the Cortex REST API. By default, the 'SNOWFLAKE.CORTEX USER' database role is granted to the 'PUBLIC' role, allowing all users to call Cortex AI functions. To enforce a more restrictive access policy, the administrator revokes 'SNOWFLAKE.CORTEX USER' from 'PUBLIC'. Which of the following actions must the administrator take to ensure specific roles can 'still' make Cortex REST API requests, and what are the implications?

- A. The 'SNOWFLAKE.CORTEX USER' database role must be granted to the specific account roles, and then these account roles must be granted to users. Additionally, the account parameter can be used to restrict which models are accessible.
- B. Access for Cortex REST API is managed independently of database roles; a separate REST API key must be provisioned for each user or application.

- C. The 'SNOWFLAKCORTEX USER' database role is only required for SQL functions, not for the Cortex REST API, so no further action is needed after revoking 'PUBLIC' for REST API access.
- D. Only the role can make cortex REST API calls after revoking 'SNOWFLAKE.CORTEX\_USER' from 'PUBLIC', as this role inherently bypasses all other access controls.
- E. The 'SNOWFLAKE.CORTEX USER' database role must be granted directly to individual users who need access, as it cannot be granted to other account roles.

**Answer: A**

Explanation:

To send a REST API request to Cortex, the default role of the calling user must be granted the 'SNOWFLAKE.CORTEX\_USER' database role. By default, this role is granted to 'PUBLIC', but it can be revoked. If revoked, the 'CORTEX USER' role must be explicitly granted to other account roles, which are then granted to users. The 'CORTEX\_USER' role cannot be granted directly to a user. The 'CORTEX MODELS\_ALLOWLIST' parameter can also be used to restrict which models are accessible at the account level for Cortex functions, including those accessed via the REST API. Therefore, option B correctly outlines the required actions and an additional control. Options A, C, D, and E are incorrect as they misrepresent the access control mechanisms or requirements for Cortex REST API.

**NEW QUESTION # 311**

An ML engineer is preparing a Docker image for a custom LLM application that will be deployed to Snowpark Container Services (SPCS). The application uses a mix of packages, some commonly found in the Snowflake Anaconda channel and others from general open-source repositories like PyPI. They have the following Dockerfile snippet and need to ensure the dependencies are correctly installed for the SPCS environment to support a GPU workload. Which of the following approaches for installing Python packages in the Dockerfile would ensure a robust and compatible setup for a custom LLM running in Snowpark Container Services, based on best practices for managing dependencies in this environment?

- A.

```
RUN conda install -n rapids -c defaults snowflake-ml-python snowflake-snowpark-python pandas jupyterlab && \
pip install transformers==4.34.0 tokenizers peft sentencepiece vllm==0.2.1.post1 bitsandbytes datasets absl-py==1.3.0
```

```
RUN apt-get update && apt-get install -y -o=Install-Recommends python3-pip && \
```

- B. pip install --no-cache-dir -r requirements.txt
- C.

```
RUN conda install -n rapids -c conda-forge snowflake-ml-python snowflake-snowpark-python pandas jupyterlab transformers==4.34.0 tokenizers && \
pip install peft sentencepiece vllm==0.2.1.post1 bitsandbytes datasets absl-py==1.3.0
```

- D.

```
RUN conda install -n rapids -c https://repo.anaconda.com/pkgs/snowflake snowflake-ml-python snowflake-snowpark-python pandas jupyterlab && \
pip install transformers==4.34.0 tokenizers peft sentencepiece vllm==0.2.1.post1 bitsandbytes datasets absl-py==1.3.0
```

- E.

```
RUN pip install snowflake-ml-python snowflake-snowpark-python pandas jupyterlab transformers==4.34.0 tokenizers peft sentencepiece vllm==0.2.1. \
post1 bitsandbytes datasets absl-py==1.3.0
```

**Answer: D**

Explanation:

Option B is correct. The provided Dockerfile example for deploying Llama 2 in Snowpark Container Services explicitly uses 'conda install -n rapids -c https://repo.anaconda.com/pkgs/snowflake' to install Snowflake-specific packages like 'snowflake-ml-python' and 'snowflake-snowpark-python' from the Snowflake Anaconda channel. It then uses 'pip install' for other open-source libraries that are not available or preferred from the Anaconda channels. Option A is incorrect because while pip can install many packages, the provided example demonstrates using 'conda' from the Snowflake Anaconda channel for certain foundational packages. Option C is incorrect because while 'conda-forge' is a common channel for open-source packages, the specific Snowflake-related packages in the example are pulled directly from the 'https://repo.anaconda.com/pkgs/snowflake' channel. Although Source notes that 'conda-forge' is assumed for 'conda\_dependencies' in when building container images, a Dockerfile explicitly defining 'RUN conda install' can specify the channel, which the example demonstrates. Option D is incorrect because the 'defaults' channel often requires user acceptance of Anaconda terms, which is not feasible in an automated build environment. Option E is a generic approach for pip dependencies but doesn't specifically address the recommended use of 'conda' from the Snowflake Anaconda channel for certain core Snowflake packages as shown in the practical example.

**NEW QUESTION # 312**

.....

We provide 1 year of free updates. In conclusion, Exams4sures guarantees that if you use the product, you will pass the GES-C01 exam on your first try. Its primary goal is to save students time and money, not just conduct a business transaction. Candidates can take advantage of the free trials to evaluate the quality and standard of the GES-C01 Dumps before making a purchase. With the right GES-C01 study material and support team passing the examination at first attempt is an achievable goal.

**GES-C01 Guide Torrent:** <https://www.exams4sures.com/Snowflake/GES-C01-practice-exam-dumps.html>

SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam dumps offers are categorized into several categories, so you can find the one that's right for you. In order to make you be rest assured to buy our GES-C01 exam software, we provide the safest payment method –PayPal payment, Snowflake GES-C01 Practice Braindumps We made the practice materials for conscience's sake to offer help, Snowflake GES-C01 Practice Braindumps How long can I use my product?

He is a center associate and advisory board member of the University of GES-C01 Pittsburgh Asia Studies Center and a member of the economics advisory board of the Duquesne University Palumbo Donahue School of Business.

## Quiz GES-C01 - SnowPro® Specialty: Gen AI Certification Exam High Hit-Rate Practice Braindumps

One question that comes up is, What if someone is on my blog and she clicks to share a post on Twitter, SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam dumps offers are categorized into several categories, so you can find the one that's right for you.

In order to make you be rest assured to buy our GES-C01 exam software, we provide the safest payment method –PayPal payment. We made the practice materials for conscience's sake to offer help.

How long can I use my product, But is it enough?

BONUS!!! Download part of Exams4sures GES-C01 dumps for free: <https://drive.google.com/open?id=1oidJ34Uno-xwR2Z5MdY6w9Xzonv7O1zX>