


# Latest GES-C01 Exam Preparation & Test GES-C01 Score Report



BTW, DOWNLOAD part of ExamCost GES-C01 dumps from Cloud Storage: <https://drive.google.com/open?id=1J9R9GJI7BOa73HZuhhWuoDy9VBgWQZDR>

The solution is closer to you than you can imagine, just contact the support team and continue enjoying your study with the SnowPro® Specialty: Gen AI Certification Exam preparation material. ExamCost offers affordable SnowPro® Specialty: Gen AI Certification Exam exam preparation material. You don't have to go beyond your budget to buy updated Snowflake GES-C01 Dumps. To make your GES-C01 exam preparation material smooth, a bundle pack is also available that includes all the 3 formats of dumps questions.

To make an open entrance and cash, everybody should gather themselves with the right and built up base on material for Snowflake GES-C01. The top-notch highlights are given to clients to affect the essential undertaking in certification. Every one of you can test your course of action with Snowflake GES-C01 Dumps by giving the phony test. Mock tests are outstandingly worked for you to make heads or tails of your goofs while giving Snowflake GES-C01.

>> Latest GES-C01 Exam Preparation <<

## Latest GES-C01 Exam Preparation - Snowflake Test GES-C01 Score Report: SnowPro® Specialty: Gen AI Certification Exam Pass Success

At the same time, our service guidelines have always been customer first. As long as you choose GES-C01 real exam, we will be responsible for you in the end. Every GES-C01 exam practice's staff member is your family they will accompany you to achieve your dream! Our company's service aim is to make every customer satisfied! GES-C01 Training Materials are looking forward to being able to accompany you on such an important journey.

## Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample

## Questions (Q315-Q320):

### NEW QUESTION # 315

An ML engineering team is preparing to log a custom Python model to the Snowflake Model Registry. This model has several Python package dependencies. The team wants to ensure the model can be deployed optimally, either in a Snowflake warehouse or to Snowpark Container Services (SPCS), depending on future needs. They are particularly concerned with how dependency specification impacts deployment eligibility. Which statements accurately describe how Snowflake handles model dependencies and determines deployment eligibility for custom Python models logged in the Model Registry, particularly when considering both Snowflake warehouse and Snowpark Container Services (SPCS) environments? (Select all that apply.)

- A. For models intended for SPCS, 'pip\_requirements' are always preferred over 'conda\_dependencies' because SPCS strictly prohibits the use of any conda packages from 'conda-forge'
- B. Specifying both 'conda\_dependencies' and 'pip\_requirements' for a model is recommended to cover all possible deployment scenarios, and Snowflake's build process ensures compatibility between them.
- C. When 'conda\_dependencies' are specified for a model to be deployed to SPCS, these dependencies are by default obtained from 'conda-forge' rather than the Snowflake conda channel.
- D. The function will fail if 'WAREHOUSE' is specified in 'target\_platforms' but the model's size or GPU requirements make it ineligible for warehouse deployment.
- E. If all of a model's 'conda\_dependencies' are available in the Snowflake conda channel, the model is automatically deemed eligible to run in a warehouse.

**Answer: C,D,E**

Explanation:

Option A is correct. When a model version is logged using 'reg\_log\_model', its 'conda\_dependencies' are validated against the Snowflake conda channel. If all dependencies are found there, the model is considered eligible to run in a warehouse. Option B is incorrect. Snowpark Container Services models, by default, obtain their 'conda\_dependencies' from 'conda-forge'. Therefore, SPCS does not prohibit conda packages from 'conda-forge'. Option C is correct. The Snowflake documentation explicitly states that for models running on Snowpark Container Services (SPCS), 'conda-forge' is the assumed channel for 'conda\_dependencies', while the Snowflake conda channel is for warehouse deployments only. Option D is correct. If the 'WAREHOUSE' platform is specified in the 'target\_platforms' argument of, and the model is ineligible for warehouse deployment (e.g., due to its size, dependencies, or GPU requirements), the call will fail. Option E is incorrect. Snowflake recommends using 'either' 'conda\_dependencies' or, but not both simultaneously. This is because combining both can lead to package conflicts, causing the container image to build successfully but potentially resulting in an unexpected or non-functional container image.

### NEW QUESTION # 316

A Gen AI Specialist is developing a conversational analytics application using Cortex Analyst, aiming to provide a seamless multi-turn conversation experience for business users querying structured data. The team observes that follow-up questions are sometimes misinterpreted, especially when the conversation history is long. Which of the following statements accurately describe how Cortex Analyst handles multi-turn conversations and key considerations for optimizing this functionality?

- A. Cortex Analyst incorporates an additional LLM summarization agent before its original workflow to rewrite current-turn questions based on conversation history, with Llama 3.1 70B being a recommended model for this task due to its performance in evaluating summarization quality.
- B. When a user shifts intent frequently in a multi-turn conversation, Cortex Analyst automatically resets the conversation history to prevent misinterpretations and improve accuracy.
- C. Multi-turn conversation in Cortex Analyst is primarily handled by the CORTEX\_ANALYST\_MODEL\_AZURE\_OPENAI parameter, which, when enabled, allows Azure OpenAI models to manage conversational context more effectively.
- D. To address misinterpretation in long conversations, the max\_tokens parameter for the Cortex Analyst REST API should be significantly increased to ensure the LLM receives the complete historical context without truncation.
- E. Cortex Analyst simply passes the entire conversation history to all subsequent LLM calls, and optimizing this requires manually truncating the array in messages the REST API request.

**Answer: A**

Explanation:

Cortex Analyst supports multi-turn conversations for data-related questions by incorporating an additional LLM summarization agent. This agent processes the conversation history and rewrites the current-turn question to include relevant context from previous turns, thereby providing a more coherent and accurate query for subsequent processing. This approach avoids passing arbitrarily

long conversation histories to every LLM agent, which would lead to longer inference times and non-determinism. 'Llama 3.1 70B' has been identified as a sufficient model for this summarization task, achieving high accuracy in rewriting questions. Option A is incorrect because Cortex Analyst specifically uses a summarization agent to avoid simply passing the entire, potentially long, conversation history to all LLM calls. Option C is incorrect. The parameter controls the option to use Azure OpenAI models with Cortex Analyst (a legacy path that is discouraged), but it does not describe the mechanism for handling multi-turn conversational context. Option D is incorrect. While 'max\_tokens' influences the length of LLM outputs, the strategy for handling long conversation history in Cortex Analyst is to use a summarization agent to create a concise, relevant context, not simply to increase the token limit to send all historical data. Increasing 'max\_tokens' for entire conversation histories would lead to higher costs and potentially longer latencies. Option E is incorrect. The documentation suggests that if a conversation becomes too long or the user's intent shifts frequently, users \*might need to reset\* the conversation, but it does not state that Cortex Analyst \*automatically\* performs this reset.

### NEW QUESTION # 317

A data engineering team is onboarding a new client whose workflow involves extracting critical financial data from thousands of daily scanned PDF receipts. They decide to use Snowflake Document AI and store all incoming PDFs in an internal stage name. After deploying their pipeline, they observe intermittent failures and varying error messages in the output, specifically:

```
1.
{ "__processingErrors": [ "cannot identify image file <_io.BytesIO object at 0x...>" ] }

2.
{ "__processingErrors": [ "Document has too many pages. Actual: 130. Maximum: 125." ] }
```

Which two of the following actions are most likely required to resolve these processing errors?

- A. Change the virtual warehouse size from an X-Small to a Large to improve Document AI processing speed.
- B. Grant the 'SNOWFLAKE\_CORTEX\_USER' database role to the role executing the '!PREDICT' function.
- C. Increase the 'max\_tokens' parameter within the '!PREDICT' function options to accommodate longer document processing.
- D. Ensure the internal stage is configured with 'ENCRYPTION = (TYPE = 'SNOWFLAKE\_SSE')'.
- E. Split any PDF documents exceeding 125 pages into smaller, compliant files, or reject them if splitting is not feasible.

**Answer: D,E**

Explanation:

The first error message, 'cannot identify image file', is a known error that occurs when an internal stage used for Document AI is not configured with 'SNOWFLAKE\_SSE' encryption. Therefore, option A is a direct solution. The second error message, 'Document has too many pages. Actual: 130. Maximum: 125.', indicates that some documents exceed Document AI's page limit of 125 pages per document. Option B directly addresses this limitation. Option C is incorrect because 'max\_tokens' is relevant for LLM output length, not document input page/size limits. Option D is incorrect because scaling up the warehouse for Document AI does not increase query processing speed and is not recommended for cost efficiency; X-Small, Small, or Medium warehouses are typically sufficient for Document AI. Option E is incorrect because is the required database role for Document AI, not 'SNOWFLAKE\_CORTEX\_USER'.

### NEW QUESTION # 318

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 ('invoice\_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.

```
CREATE OR REPLACE STAGE invoice_stage URL = 's3://my-invoice-bucket/' CREDENTIALS=(AWS_KEY_ID='...' AWS_SECRET_KEY='...') ENCRYPTION = (TYPE = 'AWS_SSE_KMS');
```

- B.

- C. `ALTER STAGE invoice_stage SET DIRECTORY = (ENABLE = FALSE);`
- D. `ALTER STAGE invoice_stage SET ENCRYPTION = (TYPE = 'NONE');`
- E. `CREATE OR REPLACE STAGE invoice_stage DIRECTORY = (ENABLE = TRUE) ENCRYPTION = (TYPE = 'SNOWFLAKE_SSE');`

**Answer: E**

**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 # 319

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 DATABASE ROLE SNOWFLAKE.CORTEX_USER TO ROLE document_processor_role;`
- B. `GRANT DATABASE ROLE SNOWFLAKE.CORTEX_USER TO ROLE document_processor_role;`
- C. `GRANT CREATE SCHEMA ON DATABASE doc_processing_db TO ROLE document_processor_role;`
- D. `GRANT DATABASE ROLE SNOWFLAKE.DOCUMENT_INTELLIGENCE_CREATOR TO ROLE document_processor_role;`
- E. `GRANT USAGE ON WAREHOUSE doc_processing_db TO ROLE document_processor_role;`

**Answer: D**

**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 Document AI capabilities.

### NEW QUESTION # 320

.....

A free trial service is provided for all customers by GES-C01 study materials, whose purpose is to allow customers to understand our products in depth before purchase. Many students often complain that they cannot purchase counseling materials suitable for themselves. A lot of that stuff was thrown away as soon as it came back. However, you will definitely not encounter such a problem when you purchase GES-C01 Study Materials. All consumers who are interested in GES-C01 study materials can download our free trial database at any time by visiting our platform.

**Test GES-C01 Score Report:** <https://www.examcost.com/GES-C01-practice-exam.html>

The Desktop GES-C01 practice exam software is created and updated in a timely by a team of experts in this field, Snowflake Latest GES-C01 Exam Preparation and as the people say, "whatever is meant to be will find its way eventually." i got passed in my exam, ♦ Based on GES-C01 Real Test, Snowflake Latest GES-C01 Exam Preparation Did you have bad purchase experience that after your payment your emails get no reply, your contacts with the site become useless, ExamCost will help you to take your Snowflake GES-C01 exam with confidence and pass it will good marks.

Once all the data has been parsed, the 'Read()' method will GES-C01 return a 'null', Teaches solid practices for establishing an

The Desktop GES-C01 Practice Exam software is created and updated in a timely by a team of experts in this field, and as the people say, “whatever is meant to be will find its way eventually.” i got passed in my exam.

**Quiz 2026 Latest GES-C01: Latest SnowPro® Specialty: Gen AI Certification**  
Exam Exam Preparation

◆ Based on GES-C01 Real Test, Did you have bad purchase experience that after your payment your emails get no reply, your contacts with the site become useless?

ExamCost will help you to take your Snowflake GES-C01 exam with confidence and pass it with good marks.

- [illegible]

BONUS!!! Download part of ExamCost GES-C01 dumps for free: <https://drive.google.com/open?id=1J9R9GJI7BOa73HZuhhWuoDy9VBgWQZDR>