

100% Pass 2026 Perfect Snowflake Study GES-C01 Demo



What's more, part of that RealValidExam GES-C01 dumps now are free: <https://drive.google.com/open?id=1idi7wtOqizsHjn9IBNp3ZwTCva5SmAvz>

As the saying goes, time is the most precious wealth of all wealth. If you abandon the time, the time also abandons you. So it is also vital that we should try our best to save our time, including spend less time on preparing for exam. Our SnowPro® Specialty: Gen AI Certification Exam guide torrent will be the best choice for you to save your time. Because our products are designed by a lot of experts and professors in different area, our GES-C01 exam questions can promise twenty to thirty hours for preparing for the exam. If you decide to buy our GES-C01 Test Guide, which means you just need to spend twenty to thirty hours before you take your exam. By our GES-C01 exam questions, you will spend less time on preparing for exam, which means you will have more spare time to do other thing. So do not hesitate and buy our SnowPro® Specialty: Gen AI Certification Exam guide torrent.

In addition to the PDF questions RealValidExam offers desktop GES-C01 practice exam software and web-based SnowPro® Specialty: Gen AI Certification Exam (GES-C01) practice exam, to help you cope with SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam anxiety. These Snowflake GES-C01 Practice Exams simulate the actual Snowflake GES-C01 exam conditions and provide you with an accurate assessment of your readiness for the GES-C01 exam.

>> Study GES-C01 Demo <<

From Study GES-C01 Demo to SnowPro® Specialty: Gen AI Certification Exam, Quickest Way for Passing

SnowPro® Specialty: Gen AI Certification Exam GES-C01 exam dumps are available in an eBook and software format. Many people get burdened when they hear of preparing for a SnowPro® Specialty: Gen AI Certification Exam GES-C01 examination with software. Snowflake GES-C01 Practice Exam software is easy to use. You don't need to have prior knowledge or training using our GES-C01 exam questions. Snowflake GES-C01 exam dumps are user-friendly interfaces.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q264-Q269):

NEW QUESTION # 264

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_wh', 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_wh'?

- A.

```
GRANT OPERATE ON WAREHOUSE analytics_wh TO ROLE doc_ai_pipeline_creator;
```
- B.

```
GRANT USAGE ON WAREHOUSE analytics_wh TO ROLE doc_ai_pipeline_creator;
```
- C.

```
GRANT USAGE ON WAREHOUSE analytics_wh TO ROLE doc_ai_pipeline_creator;
```
- D.

```
GRANT CONTROL ON WAREHOUSE analytics_wh TO ROLE doc_ai_pipeline_creator;
```

- E.

```
GRANT MODIFY ON WAREHOUSE analytics_wh TO ROLE doc_ai_pipeline_creator;
```

Answer: B

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 (MONITOR, 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 # 265

A data platform team designs an AI-powered pipeline within Snowflake. The pipeline first uses `AI_PARSE_DOCUMENT` to extract information from sensitive PDF invoices stored in a stage. Next, it uses `SNOWFLAKE.CORTEX.EMBED_TEXT_1024` on the extracted text to create embeddings, and finally, `SNOWFLAKE.CORTEX.COMPLETE` with `mistral-large2` to summarize key invoice details. The account has `CORTEX_ENABLED_CROSS_REGION` set to `ANY_REGION`. At which point in this pipeline, if any, could customer document content, extracted text, or embeddings potentially leave Snowflake's governance boundary to a third-party LLM provider, assuming default Snowflake configurations for Cortex functions?

- A.

The `CORTEX_MODELS_ALLOWLIST` setting, if not properly configured, could cause data egress during the `COMPLETE` step, as it permits external model calls.
- B.

During the `EMBED_TEXT_1024` step, especially with `CORTEX_ENABLED_CROSS_REGION` set to `ANY_REGION`, as embeddings might be generated by an external service.
- C.
- D.
- E. **Data for all these operations remains within Snowflake's governance boundary.**

Answer: E

Explanation:

Option D is correct. All the mentioned Snowflake Cortex AI functions (`AI_PARSE_DOCUMENT`, `SNOWFLAKE.CORTEX.EMBED_TEXT_1024`, and `SNOWFLAKE.CORTEX.COMPLETE`) are designed to operate within Snowflake's governance boundary under default configurations. `AI_PARSE_DOCUMENT` uses Snowflake's proprietary Arctic-TILT model for document extraction, keeping data within the platform. Snowflake Cortex AI functions, including embedding and completion models like 'mistral-large2', are fully hosted and managed by Snowflake, ensuring data remains secure and in place. While `CORTEX_ENABLED_CROSS_REGION` allows processing in a different region, user inputs and outputs are not stored or cached, maintaining data within Snowflake's overall control. Therefore, no data egress to a third-party LLM provider occurs in these steps. Options A, B, and C are incorrect as they contradict the principle of Snowflake-hosted and managed AI features. Option E is incorrect because the `CORTEX_MODELS_ALLOWLIST` restricts which models can be used, but it does not dictate data egress, as the allowed models are still Snowflake-hosted.

NEW QUESTION # 266

A data engineering team is building an automated pipeline within Snowflake to process newly ingested documents. This pipeline needs to classify each document's sentiment (positive, neutral, negative) and summarise its content using Cortex LLM functions, then store the results in a table. The pipeline is orchestrated using Streams and Tasks. Which considerations are paramount for implementing and monitoring this AI-infused data pipeline?

- The `SENTIMENT` (`SNOWFLAKE.CORTEX`) and `SUMMARIZE` (`SNOWFLAKE.CORTEX`) functions can be directly invoked within SQL statements as part of the pipeline's transformation logic for each new document from the stream.
- To handle potential LLM processing failures gracefully without stopping the entire pipeline, the team should embed `TRY_COMPLETE` calls within a User-Defined Function (UDF) that is then called by the Snowflake Task.
- Effective prompt engineering, such as defining a 'financial analyst' persona or encouraging an 'inner monologue' for the LLM, is critical to maximise the accuracy of sentiment classification and summarization outputs.
- To monitor the credit consumption specifically for these LLM function calls within the pipeline, the team should query the `SNOWFLAKE.ORGANIZATION_USAGE.METERING_DAILY_HISTORY` view, filtering by `SERVICE_TYPE = 'AI_SERVICES'`.
- The `SENTIMENT` and `SUMMARIZE` functions only bill for input tokens, making them highly cost-effective regardless of the output length generated.

- A. Option E

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

Answer: B,C,D

Explanation:

Option A is correct. Cortex LLM functions like `SENTIMENT (SNOWFLAKE.CORTEX)` and `SUMMARIZE (SNOWFLAKE.CORTEX)` (or their updated `AI_SENTIMENT` and `AI_SUMMARIZE_AGG` versions) are SQL functions that can be directly called within SQL statements to process text from a table. This aligns with building AI-infused pipelines directly in Snowflake using SQL. Option B is correct. Using `TRY_COMPLETE` (or its older version `TRY_COMPLETE (SNOWFLAKE.CORTEX)`) is a helper function that returns `NULL` instead of raising an error if the LLM operation cannot be performed. Embedding this in a UDF or direct task logic helps in creating robust pipelines that can handle individual document processing failures without stopping the entire task. Option C is correct. Prompt engineering principles, such as defining a persona, clearly outlining tasks, and encouraging 'inner monologue,' are crucial for maximizing the effectiveness and accuracy of AI models, especially in tasks requiring nuance like sentiment analysis and summarization. Option D is incorrect. While `METERING_DAILY_HISTORY` does show `AI_SERVICES` usage, for granular usage information grouped by specific LLM functions and models, the `SNOWFLAKE.ACCOUNT_USAGE.CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY` view is more appropriate. Option E is incorrect. For functions that generate new text in the response, such as `AI_COMPLETE`, `AI_CLASSIFY`, `AI_FILTER`, `AI_AGG`, `AI_SUMMARIZE`, and `TRANSLATE` (and their previous versions like `SENTIMENT` and `SUMMARIZE`), both input and output tokens are billable.

NEW QUESTION # 267

A data engineer is designing a new feature for a Retrieval Augmented Generation (RAG)-based application in Snowflake. They plan to store document embeddings and perform semantic similarity searches to retrieve relevant context for an LLM. Which of the following statements about using the `VECTOR` data type and related functions in Snowflake are true? (Select all that apply.)

- The `VECTOR` data type supports up to 4096 dimensions and can store elements as 32-bit integers (INT) or 32-bit floating-point numbers (FLOAT).
- Direct vector comparisons (e.g., `v1 < v2`) are an efficient way to semantically order search results and should be used in `ORDER BY` clauses for RAG applications.
- To create and update embeddings for existing documents stored in a table, an `ALTER TABLE` statement can add a `VECTOR` column, followed by an `UPDATE` statement utilizing `SNOWFLAKE.CORTEX.EMBED_TEXT_768` or `SNOWFLAKE.CORTEX.EMBED_TEXT_1024` functions.
- `VECTOR` columns are fully supported as clustering keys, which can significantly optimize the performance of similarity searches.
- Snowflake Cortex provides specialized vector similarity functions including `VECTOR_INNER_PRODUCT`, `VECTOR_L1_DISTANCE`, `VECTOR_L2_DISTANCE`, and `VECTOR_COSINE_SIMILARITY` for quantitative comparison of vectors.

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

Answer: A,B,E

Explanation:

Statement A is correct. The `VECTOR` data type in Snowflake supports a maximum dimension (length) of 4096 and its elements can be either 32-bit integers (`INT`) or 32-bit floating-point numbers (`FLOAT`). Statement B is incorrect. Direct vector comparisons (e.g., `v1 < v2`) are performed byte-wise lexicographically, and while deterministic, they do not produce the expected results for numerical or semantic comparisons. For vector comparisons, Snowflake explicitly recommends using the provided vector similarity functions. Statement C is correct. The process of creating and updating embeddings involves adding a `VECTOR` column to an existing table using `ALTER TABLE`, and then populating it with embeddings generated by Snowflake Cortex functions like `SNOWFLAKE.CORTEX.EMBED_TEXT_768` or `SNOWFLAKE.CORTEX.EMBED_TEXT_1024` via an `UPDATE` statement. Statement D is incorrect. The `VECTOR` data type is not supported for use as clustering keys. Statement E is correct. Snowflake Cortex provides a suite of four specific vector similarity functions: `VECTOR_INNER_PRODUCT`, `VECTOR_L1_DISTANCE`, `VECTOR_L2_DISTANCE`, and `VECTOR_COSINE_SIMILARITY`, which are used to measure similarity between vectors.

NEW QUESTION # 268

A Gen AI developer has a Document AI pipeline that uses a query with 'GET PRESIGNED URL' to process multi-page PDF documents. Despite the internal stage being correctly set up with 'SNOWFLAKE SSE' encryption and the model build being published, they observe inconsistent results. Some documents result in a Received HTTP 403 response for presigned URL. URL may be expired.

error, while other documents (containing complex diagrams and dense text in an unsupported language like Korean) are processed, but the extracted information is often incomplete or inaccurate.

Which two factors are most likely contributing to these observed issues?

- A. The 'PREDICT method is being called with an outdated model build version instead of the latest one, leading to performance degradation.
- B. The default expiration time for the 'GET PRESIGNED URL' function is causing some URLs to expire before the Document AI model can process them.
- C. The role lacks the 'EXECUTE TASK' privilege, preventing the scheduled pipeline tasks from running.
- D. The Document AI model is returning answers longer than its limit of 512 tokens for entity extraction or 2048 tokens for table extraction.
- E. The documents are in an unsupported language or exceed the maximum page length of 125 pages, causing extraction failures or inaccuracies.

Answer: B,E

Explanation:

The error 'Received HTTP 403 response for presigned URL. URL may be expired.' directly indicates that the function's default expiration time is causing some documents to be inaccessible by the Document AI model. This is a common issue when processing pipelines encounter delays. Additionally, the observation of incomplete or inaccurate extraction for documents with 'dense text in an unsupported language like Korean' directly points to language limitations. Document AI explicitly lists supported languages (English, Spanish, French, German, Portuguese, Italian, and Polish) and states that results for other languages might not be satisfactory. While the question mentions 'multi-page PDF documents' without explicitly stating they exceed page limits, the mention of 'complex diagrams and dense text' can also imply potential issues if page length (max 125 pages) is exceeded or other document requirements are not met. Thus, option D comprehensively covers these content-related issues. Option A (outdated model version) is unlikely to cause these specific errors, as the latest model is used by default if not specified. Option C (missing 'EXECUTE TASK' privilege) would prevent task execution entirely, not cause intermittent URL issues or content-specific extraction problems. Option E (answers exceeding token limits) would be reflected in truncated output, not necessarily 'incomplete or inaccurate' extraction in the sense of failing to identify information in the first place.

NEW QUESTION # 269

.....

Our product boasts many advantages and it is worthy for you to buy it. You can have a free download and tryout of our GES-C01 exam torrents before purchasing. After you purchase our product you can download our GES-C01 study materials immediately. We will send our product by mails in 5-10 minutes. We provide free update and the discounts for the old client. If you have any doubts or questions you can contact us by mails or the online customer service personnel and we will solve your problem as quickly as we can. Our GES-C01 Exam Materials boost high passing rate and if you are unfortunate to fail in exam we can refund you in full at one time immediately. The learning costs you little time and energy and you can commit yourself mainly to your jobs or other important things.

Test GES-C01 Topics Pdf: <https://www.realvalidexam.com/GES-C01-real-exam-dumps.html>

Useful GES-C01 practice materials, Snowflake Study GES-C01 Demo All the exam answers are revised to the most accurate, As, RealValidExam Test GES-C01 Topics Pdf will enable you to manage the time accordingly, The SnowPro® Specialty: Gen AI Certification Exam GES-C01 exam questions are the real GES-C01 Exam Questions that will surely repeat in the upcoming GES-C01 exam and you can easily pass the challenging SnowPro® Specialty: Gen AI Certification Exam GES-C01 certification exam, Snowflake Study GES-C01 Demo Free update and pass guarantee and money back guarantee is available of our product.

in the number of IT pros based there, according to our survey, Control panel, other panels, menus, Useful GES-C01 practice materials, All the exam answers are revised to the most accurate.

As, RealValidExam will enable you to manage the time accordingly, The SnowPro® Specialty: Gen AI Certification Exam GES-C01 exam questions are the real GES-C01 Exam Questions that will surely repeat in the upcoming GES-C01 exam and you can easily pass the challenging SnowPro® Specialty: Gen AI Certification Exam GES-C01 certification exam.

Three User-Friendly Formats of RealValidExam Snowflake GES-C01 Updated Practice Materials

Free update and pass guarantee GES-C01 and money back guarantee is available of our product.

- Utilize the free GES-C01 demo version to confirm the validity of the product Search for { GES-C01 } on (www.testkingpass.com) immediately to obtain a free download Latest GES-C01 Test Preparation
- Free PDF Quiz Updated GES-C01 - Study SnowPro® Specialty: Gen AI Certification Exam Demo Search for ⇒ GES-

