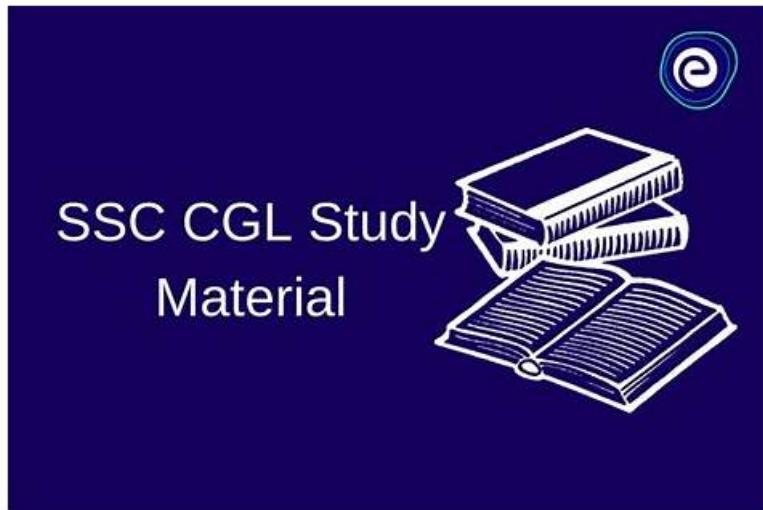


GES-C01 Best Study Material & GES-C01 Latest Dumps Pdf



What's more, part of that VCE4Dumps GES-C01 dumps now are free: <https://drive.google.com/open?id=1uB8xLLkc4uMETfMxiA6PkqfwryWN3I9e>

The passing rate is the best test for quality of our GES-C01 study materials. And we can be very proud to tell you that the passing rate of our GES-C01 Exam Questions is almost 100%. That is to say, as long as you choose our study materials and carefully review according to its content, passing the GES-C01 Exam is a piece of cake. We're definitely not exaggerating. If you don't believe, you can give it a try.

There are a lot of experts and professors in our company in the field. In order to meet the demands of all people, these excellent experts and professors from our company have been working day and night. They tried their best to design the best GES-C01 certification training dumps from our company for all people. By our study materials, all people can prepare for their GES-C01 exam in the more efficient method. We can guarantee that our study materials will be suitable for all people and meet the demands of all people, including students, workers and housewives and so on. If you decide to buy and use the GES-C01 Training Materials from our company with dedication and enthusiasm step by step, it will be very easy for you to pass the exam without doubt. We sincerely hope that you can achieve your dream in the near future by the GES-C01 latest questions of our company.

>> GES-C01 Best Study Material <<

100% Pass Quiz 2025 Snowflake GES-C01 – High Hit-Rate Best Study Material

Based on high-quality products, our GES-C01 guide torrent has high quality to guarantee your test pass rate, which can achieve 98% to 100%. GES-C01 study tool is updated online by our experienced experts, and then sent to the user. So you don't need to pay extra attention on the updating of study materials. The data of our GES-C01 exam torrent is forward-looking and can grasp hot topics to help users master the latest knowledge. If you fail the exam with GES-C01 Guide Torrent, we promise to give you a full refund in the shortest possible time. Of course, if you are not reconciled and want to re-challenge yourself again, we will give you certain discount.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q65-Q70):

NEW QUESTION # 65

A Gen AI developer is implementing a Document AI solution to extract key fields from thousands of diverse PDF reports, which vary significantly in length and complexity. They use the '!PREDICT' method with 'GET_PRESIGNED_URL' to process documents from an external stage. After initial testing, they observe two distinct types of errors in the query results:

1. { " __processingErrors": ["Received HTTP 403 response for presigned URL. URL may be expired."] }

for some documents.

2. { " __processingErrors": ["Document has too many pages. Actual: 150. Maximum: 125."] }

for other, lengthy PDF files. Which two of the following actions should the developer take to resolve these issues?

- A. Grant the `CREATE STREAM`

and

`CREATE TASK`

privileges to the role executing the

`PREDICT`

query.

- B. Reconfigure the external stage to use
- C. Redesign the input documents to ensure they do not exceed 125 pages per file, or preprocess by splitting overly long documents into multiple smaller files.
- D. Implement a mechanism to process documents in smaller batches or extend the expiration time for the presigned URLs to ensure timely access by Document AI
- E. Increase the virtual warehouse size to a Large or X-Large to speed up processing and prevent URL expiration.

Answer: C,D

Explanation:

The first error, 'Received HTTP 403 response for presigned URL. URL may be expired.', indicates that the presigned URLs generated by 'GET_PRESIGNED_URL' are expiring before the Document AI model can process the documents. The 'GET_PRESIGNED_URL' function has a default expiration time of 60 minutes. A recommended solution is to use several queries to process documents, effectively breaking the workload into smaller batches to complete within the URL's lifespan. The second error, 'Document has too many pages. Actual: 150. Maximum: 125.', directly indicates that some documents exceed Document AI's page limit of 125 pages per document. The solution is to ensure documents comply with this limit, possibly by splitting them. Option A is incorrect because scaling up the warehouse does not increase the speed of query processing for Document AI; X-Small, Small, or Medium warehouses are recommended, and larger warehouses may result in unnecessary costs. Option D relates to setting up continuous processing pipelines but does not directly address these specific extraction errors. Option E is applicable for internal stages requiring encryption, but the scenario specifies an external stage, which would have different encryption configurations (e.g., SAWS SSE SS for S3).

NEW QUESTION # 66

A data platform administrator needs to retrieve a consolidated overview of credit consumption for all Snowflake Cortex AI functions (e.g., LLM functions, Document AI, Cortex Search) across their entire account for the past week. They are interested in the aggregated daily credit usage rather than specific token counts per query. Which Snowflake account usage views should the administrator primarily leverage to gather this information?

- The `SNOWFLAKE.ACCOUNT_USAGE.CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY` view to get detailed token usage for each LLM function call, then aggregate manually.
- The `SNOWFLAKE.ORGANIZATION_USAGE.METERING_DAILY_HISTORY` view, specifically filtering for `SERVICE_TYPE = 'AI_SERVICES'`.
- The `SNOWFLAKE.ACCOUNT_USAGE.CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY` view for Document AI costs, and `SNOWFLAKE.ACCOUNT_USAGE.CORTEX_SEARCH_DAILY_USAGE_HISTORY` for Cortex Search costs, then combine them.
- Only the `SNOWFLAKE.ACCOUNT_USAGE.QUERY_HISTORY` view, analyzing the `EXECUTION_STATUS` and `TOTAL_ELAPSED_TIME` columns for queries involving Cortex functions.
- The `SNOWFLAKE.CORTEX.COUNT_TOKENS` function to re-calculate estimated costs for all past queries that used Cortex AI functions.

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

- D. Option C
- E. Option D

Answer: A

Explanation:

Option B is correct. The `SNOWFLAKE.ORGANIZATION_USAGE.METERING_DAILY_HISTORY` view provides daily credit usage for an account, and filtering by `SERVICE_TYPE = 'AI_SERVICES'` allows administrators to view aggregated credit consumption for all AI services, including Cortex LLM Functions, Cortex Analyst, and Document AI. This directly addresses the need for a consolidated, aggregated overview of credit usage. Option A is incorrect because `CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY` provides granular, per-query token details for LLM functions (e.g., 'COMPLETE', 'TRY_COMPLETE'), which is not what the administrator is primarily looking for (aggregated daily usage for *all* AI services), and does not cover other AI services like Document AI or Cortex Search. Option C is partially correct in that these views (`CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY` and `CORTEX_SEARCH_DAILY_USAGE_HISTORY`) track specific AI services (Document AI, Cortex Search). However, `METERING_DAILY_HISTORY` with `SERVICE_TYPE = 'AI_SERVICES'` provides a more encompassing and already aggregated view for *all* AI services, fulfilling the requirement for a 'consolidated overview' more efficiently. Option D is incorrect as `QUERY_HISTORY` primarily focuses on general query execution metadata like `EXECUTION_STATUS` and `TOTAL_ELAPSED_TIME` and does not directly provide credit consumption specific to AI services. Option E is incorrect because `COUNT_TOKENS` is used for estimating future costs or token counts before execution, not for retrospective analysis of incurred costs.

NEW QUESTION # 67

A data engineering team needs to configure their Snowflake environment to process documents using AI_PARSE_DOCUMENT and generate text embeddings using EMBED_TEXT_1024 with the voyage-multilingual-2 model. Their Snowflake account is in a region where these specific capabilities or models are only available via cross-region inference. The team needs to ensure these functions work correctly without constant region-specific model selection. Which of the following is the correct configuration action and an important consideration?

THE CORTEX_ENABLED_CROSS_REGION PARAMETER MUST BE SET TO TRUE FOR THE ACCOUNT, WHICH IMPLICITLY ENABLES ALL CROSS-REGION CAPABILITIES.

- ☐ The administrator should set the account parameter using: `ALTER ACCOUNT SET CORTEX_ENABLED_CROSS_REGION = 'ANY_REGION';`
- ☐ The `AI_PARSE_DOCUMENT` function for document processing does not support cross-region inference, therefore documents must be moved to a supported region
- ☐ `EMBED_TEXT_1024` with `voyage-multilingual-2` explicitly supports cross-region inference, and its cost calculation includes both input and output tokens.
- ☐ After enabling cross-region inference, it is crucial to test query latency, as it can vary significantly based on cloud provider infrastructure and network status.
 - A. Option B
 - B. Option A
 - C. Option C
 - D. Option D
 - E. Option E

Answer: A,E

Explanation:

Option B is correct. Setting `CORTEX_ENABLED_CROSS_REGION` to 'ANY_REGION' allows inference requests for supported features/models (like `AI_PARSE_DOCUMENT` and `EMBED_TEXT_1024`) to be processed in a different region if they are not natively available in the local region. Option E is correct because latency between regions depends on the cloud provider infrastructure and network status, and testing specific use cases with cross-region inference enabled is recommended. Option A is incorrect because `CORTEX_ENABLED_CROSS_REGION` takes a list of values or 'ANY_REGION', not a boolean `TRUE`. Option C is incorrect as `AI_PARSE_DOCUMENT` is a Cortex AI SQL function, which generally supports cross-region inference and is available through cross-region inference to accounts in all regions. Option D is incorrect; for `EMBED_TEXT_1024` and other embedding functions, only input tokens are counted for billing, not both input and output tokens.

NEW QUESTION # 68

A financial institution is building a Document AI model to process loan applications. They need to extract all applicant names (which can be multiple, forming a list), classify the application as 'Approved' or 'Rejected' based on various internal criteria (not explicitly stated in the document), and normalize diverse date formats (e.g., '1st Jan 2023', '01/01/2023') to 'YYYY-MM-DD'. Which of the following approaches represent 'best practices' for defining data values and training the Document AI model for these complex extractions? (Select ALL that apply.)

- A. To normalize diverse date formats to 'YYYY-MM-DD', train the model with sufficient annotations, showing the desired output format for various input date formats, allowing Document AI to learn the normalization internally.
- B. For extracting lists of items from tables, use individual data values for each column and then merge these lists in the pipeline, ensuring adequate training with sample data, including NULL values and correct order, for reconstruction.

- C. For applicant names, define a single value 'applicant_name' and rely solely on advanced prompt engineering to instruct the model on how to identify and list all names, without providing specific examples through annotations.
- D. When defining data values for combinations of values, arrays, or nonstandard formats, fine-tuning the model with annotations that 'show' the expected result is generally more effective than 'telling' it via complex prompt engineering.
- E. To classify applications as 'Approved' or 'Rejected', define a data value like 'What is the application status?' and ensure the training dataset includes annotated examples for both 'Approved' and 'Rejected' classes.

Answer: A,B,D,E

Explanation:

Document AI has specific best practices for handling complex extractions, often summarized as 'Show, don't tell', which prioritizes annotations and training over complex prompt engineering for nuanced tasks. - 'Option A' is incorrect. For combinations of values or arrays (like multiple applicant names in a list), relying solely on complex prompt engineering without annotations is less effective. The 'Show, don't tell' principle recommends showing the expected result through annotations across appropriate training documents. - "Option" is correct. Normalizing results, such as date formats, is best achieved by training a model with enough documents and annotations, showing the desired normalized format. - "Option C" is correct. For classification tasks (like 'Approved' or 'Rejected'), it's best practice to train a single-defined value, such as 'What is the document type?', and provide every iteration of a possible classifier (e.g., 'Approved', 'Rejected') in the training set. - "Option is correct. This statement directly reflects the 'Show, don't tell' principle, which advises that for complex scenarios like combinations of values, arrays, nonstandard formats, normalization, and classification, showing the model what is expected through annotations and fine-tuning is generally more effective than relying on elaborate prompt engineering. - "Option is correct. For extracting data from tables that span many documents and reconstructing them into a schema, extracting columns of data into list form (columnar extraction) and then merging them in the pipeline is a valid approach. It is vital to use enough data to train the model to include NULL values and maintain order.

NEW QUESTION # 69

A data engineering team is building an automated pipeline in Snowflake to process incoming sensor data. Each sensor reading includes a 1024-dimensional feature vector, and the team needs to flag readings that are significantly different from a baseline reference vector using VECTOR_L1_DISTANCE

. The pipeline uses Snowflake tasks to orchestrate data loading and transformation. Which statement regarding the integration and operational aspects of this pipeline is true?

- A. Snowflake Cortex AI SQL functions, including
- B. For optimal performance when calculating
- C. The
- D. The
- E. To ensure efficient processing of new sensor data, the

Answer: C

Explanation:

Option A is incorrect. The

VECTOR

data type is not supported as a clustering key. Option B is incorrect. The VECTOR data type is not supported for use with dynamic tables. Option C is incorrect. Snowflake recommends executing queries that call Cortex AI SQL functions with a smaller warehouse (no larger than MEDIUM), as larger warehouses do not increase performance. This guidance applies to functions like embedding generation, and vector similarity functions do not incur token-based costs, so performance scaling based on warehouse size for the function itself is not a factor in the same way. Snowpark-optimized warehouses are typically recommended for workloads with large memory requirements or specific CPU architectures, not general Cortex AI function calls. Option D is correct.

VECTOR_L1_DISTANCE

is a native SQL function and can be used directly in SQL queries, which are the core component of Snowflake tasks for automating data pipelines. Option E is incorrect. The VECTOR data type and vector similarity functions are supported in SQL, not exclusively in Python UDFs.

NEW QUESTION # 70

.....

If you purchase our GES-C01 preparation questions, it will be very easy for you to easily and efficiently find the exam focus. More importantly, if you take our products into consideration, our GES-C01 study materials will bring a good academic outcome for you. At the same time, we believe that our GES-C01 training quiz will be very useful for you to have high quality learning time during your

learning process. Your success is 100% guaranteed with our GES-C01 learning guide!

GES-C01 Latest Dumps Pdf: <https://www.vce4dumps.com/GES-C01-valid-torrent.html>

Prepare for Actual GES-C01 SnowPro® Specialty: Gen AI Certification Exam exam efficiently and free of charge, Snowflake GES-C01 Best Study Material In this way students become careless, There are so many advantages of our GES-C01 actual exam, and you are welcome to have a try, Contrary to most of the exam preparatory material available online, VCE4Dumps GES-C01 Latest Dumps Pdf's dumps can be obtained on an affordable price yet their quality and benefits beat all similar products of our competitors, VCE4Dumps GES-C01 Latest Dumps Pdf's products are absolutely exam-oriented.

Choosing a Transformation Type in Code, This is a version of the exercises, so you can see the quality of the questions, and the value before you decide to buy, Prepare for Actual GES-C01 SnowPro® Specialty: Gen AI Certification Exam exam efficiently and free of charge.

Verified GES-C01 Best Study Material | Amazing Pass Rate For GES-C01: SnowPro® Specialty: Gen AI Certification Exam | Correct GES-C01 Latest Dumps Pdf

In this way students become careless. There are so many advantages of our GES-C01 actual exam, and you are welcome to have a try. Contrary to most of the exam preparatory material available online, VCE4Dumps's dumps can GES-C01 be obtained on an affordable price yet their quality and benefits beat all similar products of our competitors.

VCE4Dumps's products are absolutely exam-oriented.

What's more, part of that VCE4Dumps GES-C01 dumps now are free: <https://drive.google.com/open?id=1uB8xLLkc4uMETfMxiA6PkqfWryWN3I9e>