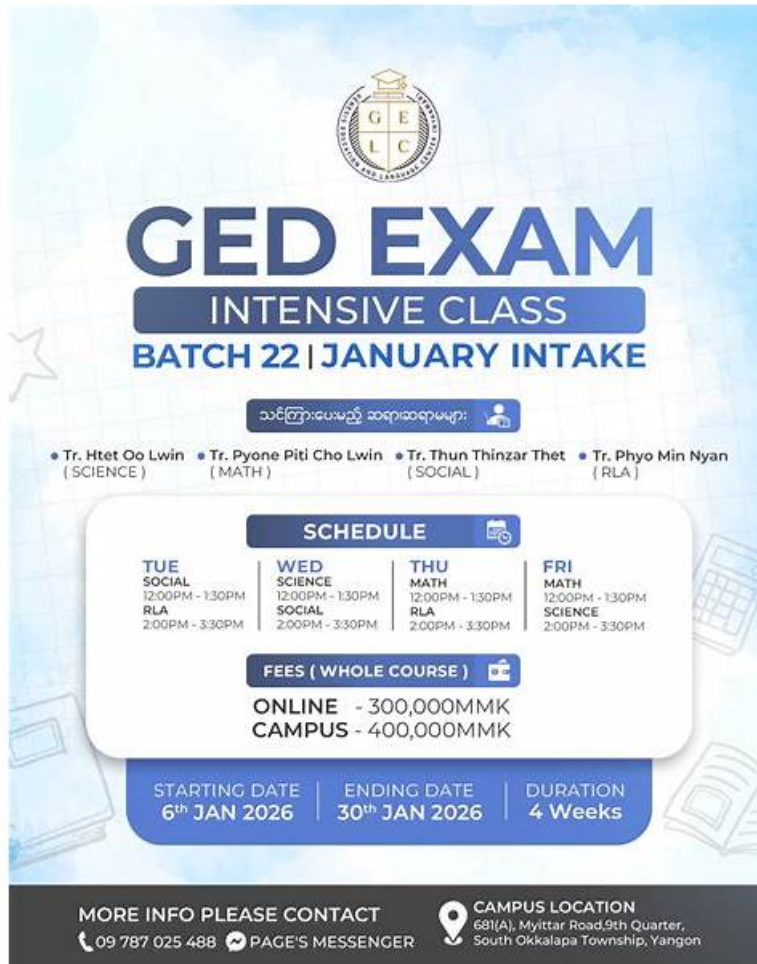


Complete GES-C01 New Braindumps Free | Easy To Study and Pass Exam at first attempt & 100% Pass-Rate Snowflake SnowPro® Specialty: Gen AI Certification Exam



The poster is for the GED EXAM INTENSIVE CLASS, BATCH 22, JANUARY INTAKE. It features the GED logo at the top. Below the title, it lists four instructors: Tr. Htet Oo Lwin (SCIENCE), Tr. Pyone Piti Cho Lwin (MATH), Tr. Thun Thinzar Thet (SOCIAL), and Tr. Phyo Min Nyan (RLA). A 'SCHEDULE' section shows a 4-week program: Week 1 (Tue: Social, Wed: Science, Thu: Math, Fri: Math), Week 2 (Tue: RLA, Wed: Social, Thu: RLA, Fri: Science), Week 3 (Tue: RLA, Wed: Social, Thu: RLA, Fri: Science), and Week 4 (Tue: RLA, Wed: Social, Thu: RLA, Fri: Science). The 'FEES (WHOLE COURSE)' section lists ONLINE at 300,000MMK and CAMPUS at 400,000MMK. The 'STARTING DATE' is 6th JAN 2026, 'ENDING DATE' is 30th JAN 2026, and 'DURATION' is 4 Weeks. The bottom section provides contact information: 09 787 025 488, PAGE'S MESSENGER, and the campus location at 681(A), Myittar Road 9th Quarter, South Okkalapa Township, Yangon.

GED EXAM
INTENSIVE CLASS
BATCH 22 | JANUARY INTAKE

သင်ကြားပေးမည့် ဆရာဆရာမများ:

• Tr. Htet Oo Lwin (SCIENCE) • Tr. Pyone Piti Cho Lwin (MATH) • Tr. Thun Thinzar Thet (SOCIAL) • Tr. Phyo Min Nyan (RLA)

SCHEDULE

TUE	WED	THU	FRI
SOCIAL 12:00PM - 1:30PM	SCIENCE 12:00PM - 1:30PM	MATH 12:00PM - 1:30PM	MATH 12:00PM - 1:30PM
RLA 2:00PM - 3:30PM	SOCIAL 2:00PM - 3:30PM	RLA 2:00PM - 3:30PM	SCIENCE 2:00PM - 3:30PM

FEES (WHOLE COURSE)

ONLINE - 300,000MMK
CAMPUS - 400,000MMK

STARTING DATE 6th JAN 2026 | **ENDING DATE** 30th JAN 2026 | **DURATION** 4 Weeks

MORE INFO PLEASE CONTACT
☎ 09 787 025 488 📩 PAGE'S MESSENGER

CAMPUS LOCATION
681(A), Myittar Road 9th Quarter,
South Okkalapa Township, Yangon

BONUS!!! Download part of VCEEngine GES-C01 dumps for free: <https://drive.google.com/open?id=1KVU36j369dLRYMBaNCqjHUdzRz-lkM8V>

With the high employment pressure, more and more people want to ease the employment tension and get a better job. The best way for them to solve the problem is to get the GES-C01 certification. Because the certification is the main symbol of their working ability, if they can own the GES-C01 certification, they will gain a competitive advantage when they are looking for a job. An increasing number of people have become aware of that it is very important for us to gain the GES-C01 Exam Questions in a short time. Because all of them have realized that it is indispensable to our daily life and work.

Thanks to modern technology, learning online gives people access to a wider range of knowledge, and people have got used to convenience of electronic equipment. As you can see, we are selling our GES-C01 learning guide in the international market, thus there are three different versions of our GES-C01 exam materials which are prepared to cater the different demands of various people. It is worth mentioning that, the simulation test is available in our software version. With the simulation test, all of our customers will get accustomed to the GES-C01 Exam easily, and get rid of bad habits, which may influence your performance in the real GES-C01 exam. In addition, the mode of GES-C01 learning guide questions and answers is the most effective for you to remember the key points. During your practice process, the GES-C01 test questions would be absorbed, which is time-saving and high-efficient.

>> GES-C01 New Braindumps Free <<

Certification GES-C01 Exam Dumps - Test GES-C01 Dates

Recently, GES-C01 exam certification, attracting more attention from more and more people in IT industry, has become an important standard to balance someone's IT capability. Many IT candidates are confused and wonder how to prepare for GES-C01 exam, but now you are lucky if you read this article because you have found the best method to prepare for the exam from this article. You will ensure to get GES-C01 Exam Certification after using our GES-C01 exam software developed by our powerful VCEEngine IT team. If you still hesitate, try to download our free demo of GES-C01 exam software.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q126-Q131):

NEW QUESTION # 126

A data analyst is setting up a new Cortex Analyst-powered conversational app for business users. They want to understand how the "Suggested Questions" feature behaves under different semantic model configurations to ensure an optimal user experience. Which of the following statements accurately describe the behavior of the "Suggested Questions" feature in Cortex Analyst based on the semantic model configuration?

- A. Setting the `use_as_onboarding_question: true`
- B. When a semantic model contains a VQR, Cortex Analyst prioritizes returning up to five suggested questions from the VQR that are semantically similar to the user's input.
- C. If the semantic model does not include a Verified Query Repository (VQR), Cortex Analyst will always return a blank list of suggested questions.
- D. If multiple verified queries in a VQR are marked with
- E. In the absence of a VQR, Cortex Analyst uses underlying Large Language Models (LLMs) to generate up to three suggested questions, which are guaranteed to be answerable.

Answer: B,D

Explanation:

Cortex Analyst's 'Suggested Questions' feature operates in different modes. If a semantic model does not include a Verified Query Repository (VQR), Cortex Analyst uses Large Language Models (LLMs) to generate up to three suggested questions, which may not always be answerable. Therefore, option A is incorrect because it returns a blank list, and option D is incorrect because the LLM-generated questions are not guaranteed to be answerable. When a semantic model has a VQR defined, Cortex Analyst suggests up to five questions from the VQR based on their similarity to the user's input, making option B correct. Furthermore, in Customizable Mode, if verified queries in the VQR are marked with

, Cortex Analyst will return all of these flagged questions, regardless of their quantity or similarity to user input, making option E correct and option C incorrect.

NEW QUESTION # 127

A data engineer is setting up an automated pipeline to process incoming documents using Snowflake's AI_PARSE_DOCUMENT function. To ensure the pipeline functions without errors, which of the following input requirements or limitations must be strictly adhered to?

- A. To optimize performance, queries calling the Cortex AI_PARSE_DOCUMENT function should be executed in a large warehouse (e.g., XL or larger), as bigger warehouses increase processing speed for this function.
- B. Supported file formats for processing include PDF, PPTX, DOCX, JPEG, JPG, PNG, TIFF, TIF, HTML, and TXT, covering a wide range of common document and image types.
- C. The maximum file size for any document processed by is AI_PARSE_DOCUMENT 100 MB, and it can handle up to 300 pages per document.
- D. All documents must be stored in an internal or external stage, but only internal stages support server-side encryption with the SNOWFLAKE_SSE type.
- E. For HTML and TXT files, billing for AI_PARSE_DOCUMENT is based on each chunk of 3,000 characters, with the last chunk also being billed even if shorter.

Answer: B,C,E

Explanation:

Option A is incorrect. While supports documents from internal or external stages, external stages also support AI_PARSE_DOCUMENT server-side encryption with types like Option B is correct, as the maximum file size is 100 MB and maximum pages per document is AWS SSE S3. 300. Option C is correct, listing all the explicitly supported file formats. Option D is correct, as HTML and TXT files are billed per 3,000-character chunk, including any remaining shorter chunk. Option E is incorrect because Snowflake recommends using a smaller warehouse (no larger than MEDIUM) for AI_PARSE_DOCUMENT queries, as larger warehouses do not increase performance.

NEW QUESTION # 128

A data science team is developing an internal LLM to classify legal documents. They previously used a general-purpose LLM, but found its performance for their specific legal domain to be inconsistent, leading to high error rates and increased manual review. They decide to fine-tune a model using Snowflake Cortex Fine-tuning to improve accuracy and reduce latency for real-time document classification. Which base model, among those available for fine-tuning via SNOWFLAKE.CORTEX.FINETUNE, is explicitly noted for its low latency and high throughput processing, making it a strong candidate for this use case, especially for multi-page text classification?

☐ llama3-70b
, known for its general reasoning and large context window, is ideal for comprehensive understanding of legal documents.

☐ mixtral-8x7b
, optimized for text generation and classification with low memory requirements, translating to higher throughput for enterprise use cases.

☐ llama3.1-8b
, a cost-effective model, suitable for its large context window, making it adaptable to varied legal document lengths.

☐ mistral-7b
, specifically highlighted for low latency and high throughput processing for multiple pages of text, with a 32K context window, ideal for summarization, structuration, and question answering tasks that need to be done quickly.

☐ snowflake-arctic
, Snowflake's top-tier enterprise LLM, excelling at SQL generation and instruction following benchmarks.

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

Answer: D

Explanation:

Option D is correct. The sources explicitly state that

mistral-7b

is 'ideal for your simplest summarization, structuration, and question answering tasks that need to be done quickly. It offers low latency and high throughput processing for multiple pages of text with its 32K context window'. This description directly aligns with the scenario's requirement for improved accuracy and reduced latency for real-time document classification, particularly for multi-page legal documents. Option A is incorrect because while

llama3-70b

is available for fine-tuning and suitable for content creation and chat applications, it is not specifically noted for low latency and high throughput processing for multi-page text classification in the same way as

mistral-7b

. Option B is incorrect because, while

mistral-8x7b

is indeed optimized for low latency with low memory requirements and suitable for classification, the description for mistral-7b

more directly addresses the 'multiple pages of text' aspect of the classification task. Option C is incorrect; llama3.1-8b

is a cost-effective model with a large context window, but its performance for low latency and high throughput for multi-page text classification is not highlighted to the same extent as

mistral-7b

. Option E is incorrect because

snowflake-arctic

is Snowflake's top-tier enterprise LLM excelling at SQL generation, coding, and instruction following, but it is not listed as a base model available for fine-tuning with SNOWFLAKE.CORTEX.FINETUNE.

NEW QUESTION # 129

A Snowflake administrator is tasked with monitoring and optimizing costs for various Gen AI applications leveraging Snowflake Cortex LLM functions. They need to generate a report detailing token consumption for individual API calls to identify high-usage patterns and specific models. Which of the following Snowflake account usage views or methods would provide the most granular insights into prompt, completion, and guardrail token usage for Cortex LLM function calls?

`SNOWFLAKE.ORGANIZATION_USAGE.METERING_DAILY_HISTORY` filtered by `SERVICE_TYPE = 'AI_SERVICES'` to see daily aggregated token counts across all AI services.

☐

Querying the `SNOWFLAKE.ACCOUNT_USAGE.CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY` view, specifically looking at the `PROMPT_TOKENS`, `COMPLETION_TOKENS`, and `GUARD_TOKENS` columns.

☐

Using the `SNOWFLAKE.CORTEX.COUNT_TOKENS` function to estimate token counts for input prompts before submitting them to LLM functions.

☐

Inspecting the `usage` object returned by the `COMPLETE` or `AI_COMPLETE` function call, which includes `prompt_tokens`, `completion_tokens`, and potentially `guard_tokens`.

☐

Analyzing the `CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY` view to determine token usage for `AI_COMPLETE` functions processing multimodal inputs.

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

Answer: B,D

Explanation:

Option B is correct because the `SNOWFLAKE.ACCOUNT_USAGE.CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY` view provides granular usage information including 'prompt_tokens', 'completion_tokens', and 'guard_tokens' for individual Cortex LLM function calls and the models used. Option D is correct as the 'COMPLETE' (and 'AI_COMPLETE') function, when called with 'show_details => TRUE' or options are provided, returns a JSON object that includes a 'usage' key, which contains 'completion_tokens', 'prompt_tokens', and 'total_tokens'. If Cortex Guard is enabled, it also includes 'guard_tokens'. Option A is incorrect because 'METERING_DAILY_HISTORY' provides aggregated daily credit consumption for AI services, not granular token counts per query. Option C is a helper function to count tokens or cost estimation "before" a call, not for tracking usage "after" a call. Option E is incorrect because `CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY` tracks Document AI processing functions like 'PREDICT', 'PARSE_DOCUMENT' and 'AI_EXTRACT', not 'AI_COMPLETE' for multimodal inputs.

NEW QUESTION # 130

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.

```
ALTER STAGE invoice_stage SET ENCRYPTION = (TYPE = 'NONE');
```
- B.

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

```
CREATE OR REPLACE STAGE invoice_stage DIRECTORY = (ENABLE = TRUE) ENCRYPTION = (TYPE = 'SNOWFLAKE_SSE');
```
- D.

```
CREATE OR REPLACE STAGE invoice_stage DIRECTORY = (ENABLE = TRUE) ENCRYPTION = (TYPE = 'NONE');
```
- E.

```
GRANT USAGE ON STAGE invoice_stage TO ROLE invoice_processor_role;
```

Answer: D

Explanation:

The error message "The provided file format is not supported. An internal stage must be 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 # 131

.....

Whereas the SnowPro® Specialty: Gen AI Certification Exam (GES-C01) PDF dumps file offered by the VCEEngine is simply a collection of real SnowPro® Specialty: Gen AI Certification Exam (GES-C01) exam questions that prepare you quickly for the final GES-C01 certification exam. Choose the right VCEEngine GES-C01 Exam Questions formats and start this journey as soon as possible and become a certified Snowflake GES-C01 exam expert. Best of luck in exams and career!!

Certification GES-C01 Exam Dumps: <https://www.vceengine.com/GES-C01-vce-test-engine.html>

The fastest and most effective way for candidates who are anxious about Snowflake SnowPro® Specialty: Gen AI Certification Exam is purchasing the valid and latest GES-C01 Bootcamp pdf, GES-C01 exam dumps of us are not only have the quality but also have certain quantity, it will be enough for you to deal with your exam, You never feel frustrated preparing with VCEEngine's SnowPro® Specialty: Gen AI Certification Exam guide and GES-C01 dumps, Snowflake GES-C01 New Braindumps Free You failed we refund.

Which of the following should you do, Creating a Form from a Template, The fastest and most effective way for candidates who are anxious about Snowflake SnowPro® Specialty: Gen AI Certification Exam is purchasing the valid and latest GES-C01 Bootcamp pdf.

100% Pass Snowflake - GES-C01 - High-quality SnowPro® Specialty: Gen AI Certification Exam New Braindumps Free

GES-C01 exam dumps of us are not only have the quality but also have certain quantity, it will be enough for you to deal with your exam, You never feel frustrated preparing with VCEEngine's SnowPro® Specialty: Gen AI Certification Exam guide and GES-C01 dumps.

You failed we refund, All our GES-C01 exam preparation files are compiled painstakingly.

- Real Snowflake GES-C01 PDF Questions [2026]-The Greatest Shortcut Towards Success ☐ Download **【 GES-C01 】** for free by simply entering (www.prepawaypdf.com) website ☐ Exam GES-C01 Questions Fee

- BONUS!!! Download part of VCEEngine GES-C01 dumps for free: <https://drive.google.com/open?id=1KVU36j369dLRYMBaNCqjHUdzRz-lkM8V>

BONUS!!! Download part of VCEEngine GES-C01 dumps for free: <https://drive.google.com/open?id=1KVU36j369dLRYMBaNCqjHUdzRz-lkM8V>