

Snowflake GES-C01 Questions - Pass Exam and Get Career Benefits

Snowflake DEA-C01 Practice Questions

SnowPro Advanced Data Engineer Certification

Order our DEA-C01 Practice Questions Today and Get Ready to Pass with Flying Colors!



DEA-C01 Practice Exam Features | QuestionsTube

- Latest & Updated Exam Questions
- Subscribe to FREE Updates
- Both PDF & Exam Engine
- Download Directly Without Waiting

<https://www.questiontube.com/exam/dea-c01/>

At QuestionsTube, you can read DEA-C01 free demo questions in pdf file, so you can check the questions and answers before deciding to download the Snowflake DEA-C01 practice questions. These free demo questions are parts of the DEA-C01 exam questions. Download and read them carefully, you will find that the DEA-C01 test questions of QuestionsTube will be your great learning materials online. Share some DEA-C01 exam online questions below.

```
1. sum(ifnull(quantity * (s.price - m.item_price), 0)) as profit
```

Our Snowflake GES-C01 Practice Materials are compiled by first-rank experts and GES-C01 Study Guide offer whole package of considerate services and accessible content. Furthermore, SnowPro® Specialty: Gen AI Certification Exam GES-C01 Actual Test improves our efficiency in different aspects. Having a good command of professional knowledge will do a great help to your life.

If you find someone around has a nice life go wild, it is because that they may have favored the use of study & work method different from normal people. GES-C01 dumps torrent files may be the best method for candidates who are preparing for their IT exam and eager to clear exam as soon as possible. People's success lies in their good use of every change to self-improve. Our GES-C01 Dumps Torrent files will be the best resources for your real test. If you choose our products, we will choose efficient & high-passing preparation materials.

>> GES-C01 Actual Dump <<

Valid Snowflake Actual Dump – High-quality Reliable GES-C01 Exam Tutorial

You can learn GES-C01 quiz torrent skills and theory at your own pace, and you are not necessary to waste your time on some useless books or materials and you will save more time and energy that you can complete other thing. We also provide every candidate who wants to get certification with free Demo to check our materials. It is time for you to realize the importance of our

GES-C01 Test Prep, which can help you solve these annoyance and obtain a GES-C01 certificate in a more efficient and productive way.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q302-Q307):

NEW QUESTION # 302

A Snowflake account is located in the AWS US East 1 (N. Virginia) region. The 'ACCOUNTADMIN' has set the 'CORTEX MODELS ALLOWLIST' to 'mistral-7b' and 'CORTEX ENABLED CROSS REGION' to 'ANY REGION'. A data scientist, whose role has only the 'SNOWFLAKE.CORTEX USER' database role, performs several 'AI COMPLETE' calls. Which of the following statements correctly describe the behavior of these calls under the given configuration?

- ☐ An 'AI_COMPLETE('mistral-7b', ...)' call will execute successfully and natively, incurring standard latency for the US East 1 region.
- ☐ An 'AI_COMPLETE('claude-3-5-sonnet', ...)' call will fail with an error indicating an unauthorized model, as it is not in the 'CORTEX_MODELS_ALLOWLIST', despite 'ANY_REGION' being enabled.
- ☐ An 'AI_COMPLETE('llama3.1-8b', ...)' call would return a completion, potentially via cross-region inference from a supported region, but only if 'llama3.1-8b' was also added to the 'CORTEX_MODELS_ALLOWLIST'.
- ☐ The Cortex LLM Playground will display all models available across 'ANY_REGION', effectively bypassing the 'CORTEX_MODELS_ALLOWLIST' due to the broad cross-region setting.
- ☐ The 'CORTEX_ENABLED_CROSS_REGION = ANY_REGION' setting implies that all models available in any Snowflake region are automatically added to the allowlist, making them accessible to any user with 'CORTEX USER'.

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

Answer: A,D

Explanation:

The 'CORTEX MODELS_ALLOWLIST' parameter restricts which models can be used for Cortex LLM functions, including 'AI_COMPLETE'. The parameter allows inference requests to be processed in a different region from the default, which may incur increased latency. Option A is correct: 'mistral-7b' is in the list. It is natively available in AWS US East 1 (N. Virginia). Thus, an 'AI_COMPLETE' call using this model will execute successfully and locally, with standard latency. Option B is incorrect: 'claude-3-5-sonnet' is not in the list. The parameter applies to 'AI_COMPLETE'. Therefore, the call would return 'NULL' (if using 'TRY_COMPLETE') or raise an error (if using 'AI_COMPLETE') due to the allowlist restriction, not solely because of 'ANY_REGION'. Option C is incorrect: 'mistral-7b' is in the current 'ANY_REGION' allows models to be processed in other regions. However, this parameter does not bypass the 'CORTEX MODELS_ALLOWLIST'. Therefore, for 'llama3.1-8b' to be used successfully (even via cross-region inference), it 'must' be included in the 'CORTEX MODELS_ALLOWLIST'. If it were included, it would execute via cross-region inference with potential latency, as it is available cross-cloud. The statement accurately describes the condition for its successful execution. Option D is incorrect: The Cortex LLM Playground respects the 'CORTEX MODELS_ALLOWLIST'. The 'ANY_REGION' setting does not bypass this governance control. Option E is incorrect: The parameter enables cross-region processing for 'allowed' features and models; it does not implicitly modify the 'CORTEX MODELS_ALLOWLIST' to include all models.

NEW QUESTION # 303

A data scientist is leveraging various Snowflake Cortex LLM functions to process extensive text data for an application. To effectively manage their budget, they need a clear understanding of how costs are incurred for each specific function. Which of the following statements accurately describe how costs are calculated for Snowflake Cortex LLM functions, with a particular focus on token usage?

The `SNOWFLAKE.CORTEX.EMBED_TEXT_768` function incurs compute costs based on both the input and output tokens processed.

For the `SNOWFLAKE.CORTEX.EXTRACT_ANSWER` function, billable tokens are determined by the sum of tokens present in the `source_document` and `question` fields.

When utilizing `AI_COMPLETE` with structured outputs, an additional compute cost is explicitly charged for the overhead of validating each generated token against the provided JSON schema.

The `SNOWFLAKE.CORTEX.CLASSIFY_TEXT` function includes the token count from category descriptions and examples as part of the input tokens for each record processed, thereby increasing the overall cost.

The `AI_PARSE_DOCUMENT` function is billed based on the total number of individual documents processed, regardless of the number of pages within each document.

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

Answer: A,B

Explanation:

Option B is correct because for the 'EXTRACT ANSWER' function, the number of billable tokens is the sum of the tokens in the 'from_text' (source_document) and 'question' fields. Option D is correct as for 'CLASSIFY_TEXT' (or 'AI_CLASSIFY'), labels, descriptions, and examples provided in the categories are counted as input tokens for each record processed, which directly increases the cost. Option A is incorrect because functions only count 'input tokens' towards the billable total, not both input and output tokens. Option C is incorrect because Cortex COMPLETE Structured Outputs does not incur additional compute cost for the overhead of verifying tokens against the supplied JSON schema, although schema complexity can increase total token consumption. Option E is incorrect because 'AI_PARSE_DOCUMENT' (and 'SNOWFLAKE.CORTEX.PARSE_DOCUMENT') billing is based on the 'number of document pages processed' (e.g., 3.33 Credits per 1,000 pages for Layout mode), not just the number of documents.

NEW QUESTION # 304

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. Ensure the internal stage is configured with 'ENCRYPTION = (TYPE = 'SNOWFLAKE_SSE')'.
- C. Increase the 'max_tokens' parameter within the '!PREDICT function options to accommodate longer document processing.
- D. Split any PDF documents exceeding 125 pages into smaller, compliant files, or reject them if splitting is not feasible.
- E. Grant the 'SNOWFLAKE_CORTX_USER' database role to the role executing the '!PREDICT function.

Answer: B,D

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 # 305

A data engineer is integrating SNOWFLAKE. CORTEX. CLASSIFY_TEXT into an automated data pipeline that uses dynamic tables to process and transform streaming text data. They have ensured that the service account used has been granted the necessary SNOWFLAKE. CORTEX_USER database role. After deploying the pipeline, they consistently receive an error whenever CLASSIFY_TEXT is invoked. Which of the following is the most likely cause of the error encountered by the data engineer?

- A. The role used by the data engineer, despite having 'SNOWFLAKE.CORTEX_USER', lacks the fundamental 'USAGE' privilege on the database where the text data is stored.
- B. The input text being processed by 'CLASSIFY_TEXT' includes extensive non-plain English content, such as code blocks, which causes the function to fail with an error.
- C. The 'task_description' provided in the optional arguments for 'CLASSIFY_TEXT' exceeds the recommended length of approximately 50 words, leading to a validation error.
- D. The array contains more than 100 unique categories, exceeding the maximum allowed limit for the function.
- E. Snowflake Cortex functions, including 'CLASSIFY_TEXT', currently do not support integration with dynamic tables within data pipelines.

Answer: E

Explanation:

Option A is plausible for a data-specific error, but the question describes a 'consistent error' during pipeline integration. The maximum number of categories is 100. Option B is incorrect because if the text contains non-plain English content like code snippets, the function 'won't return an error, but the results may not be what you expect'. This would lead to inaccurate results, not a consistent error preventing the function's execution. Option C is less likely to be the 'most' likely cause of an error specific to the 'CLASSIFY_TEXT' function's invocation, especially since the 'SNOWFLAKE.CORTEX_USER' role, which grants access to Cortex AI functions, has already been granted. Missing 'USAGE' on the data's database would typically manifest as a more general SQL access error. Option D is correct because a known limitation for Snowflake Cortex functions, including 'CLASSIFY_TEXT', is that they do not support dynamic tables. This is a fundamental incompatibility that would cause consistent errors when integrating into a dynamic table pipeline. Option E is incorrect. While a 'task_description' should be 'no more than about 50 words', this is a recommendation for optimal performance, not a strict limit that is explicitly stated to cause an error when exceeded.

NEW QUESTION # 306

A Gen AI specialist is preparing to upload a large volume of diverse documents to an internal stage for Document AI processing. The objective is to extract detailed information, including lists of items and potentially classifying document types, and then automate this process. Which of the following statements represent 'best practices or important considerations/limitations' when preparing documents and setting up the Document AI workflow in Snowflake? (Select ALL that apply.)

- A. When defining data values for extraction, especially for nonstandard formats or combinations of values, fine-tuning the model with annotations is generally more effective than relying solely on complex prompt engineering.
- B. If the Document AI model does not find an answer for a specific field, the '!PREDICT method will omit the 'value' key but will still return a 'score' key to indicate confidence that the answer is not present.
- C. To improve model training, documents uploaded should represent a real use case, and the dataset should consist of diverse documents in terms of both layout and data.
- D. Documents with a page count exceeding 125 pages or a file size greater than 50 MB will be processed, but with a potential reduction in extraction accuracy.
- E. For continuous processing of new documents, it is best practice to create a stream on the internal stage and a task to automate the '!PREDICT method execution.

Answer: A,B,C,E

Explanation:

Several best practices and limitations apply to Document AI, particularly during the preparation and training phases: - **Option A: Correct.** To enhance model training, it is crucial that the uploaded documents accurately reflect a real use case and that the dataset exhibits diversity in both layout and content. - **Option B: Correct.** If the Document AI model cannot find an answer for a specified field, it will not return a 'value' key, but it will still provide a 'score' key to indicate its confidence that the answer is absent from the document. This behaviour is shown in the example for 'buyer_name' field. - **Option C: Correct.** Creating an automated pipeline for continuous document processing involves setting up a stream on the internal stage and a task to trigger the 'IPREDICT' method. This approach is explicitly mentioned as a capability and best practice. - **Option D: Incorrect.** Documents exceeding 125 pages or 50 MB in size are strict limitations for Document AI and will result in processing errors, not merely reduced accuracy. - **Option E: Correct.** For complex extraction scenarios, such as nonstandard formats, lists, or classification, 'showing' the model what is expected through annotations and fine-tuning is often more effective than 'telling' it via elaborate prompt engineering. This is described as 'Show, don't tell' in the best practices, recommending annotations for combinations of values, arrays, nonstandard formats, normalization, and classification tasks instead of spending too much time on prompt engineering alone.

NEW QUESTION # 307

.....

By these three versions we have many repeat orders in a long run. The PDF version helps you read content easier at your process of studying with clear arrangement, and the PC Test Engine version allows you to take stimulation exam to check your process of exam preparing, which support windows system only. Moreover, there is the APP version, you can learn anywhere at any time with it at your cellphones without the limits of installation. As long as you are willing to exercise on a regular basis, the exam will be a piece of cake, because what our GES-C01 practice materials include are quintessential points about the exam.

Reliable GES-C01 Exam Tutorial: https://www.bootcamppdf.com/GES-C01_exam-dumps.html

Appropriate price, Snowflake GES-C01 Actual Dump If you want to have a general review of what you have learned, you can choose us, Each candidate takes only a few days can attend to the GES-C01 exam, In consideration of the various requirements of our customers we develop three different versions of Snowflake GES-C01 practice test questions (PDF version, PC test engine and APP test engine) for you reference, To assist you in the objective of cracking the Snowflake GES-C01 Exam, Snowflake GES-C01 Dumps is offering a study material which comes in three versions and meets all needs of your exam preparation.

What is the maximum number of Root Ports GES-C01 that a Catalyst switch can have, New First-Step book Provides an Easy Start to Computer Networking, Appropriate price, If GES-C01 Actual Dump you want to have a general review of what you have learned, you can choose us.

Providing You Marvelous GES-C01 Actual Dump with 100% Passing Guarantee

Each candidate takes only a few days can attend to the GES-C01 Exam, In consideration of the various requirements of our customers we develop three different versions of Snowflake GES-C01 practice test questions (PDF version, PC test engine and APP test engine) for you reference.

To assist you in the objective of cracking the Snowflake GES-C01 Exam, Snowflake GES-C01 Dumps is offering a study material which comes in three versions and meets all needs of your exam preparation.

- GES-C01 Advanced Testing Engine □ GES-C01 Advanced Testing Engine □ Latest GES-C01 Exam Format □ Search for > GES-C01 □ and download exam materials for free through ➡ www.prep4pass.com □ □GES-C01 Latest Practice Questions
- Online Snowflake GES-C01 Practice Test □ Copy URL ☀ www.pdfvce.com □☀□ open and search for □ GES-C01 □ to download for free □GES-C01 Certification Materials
- GES-C01 latest dumps □ Search for □ GES-C01 □ and easily obtain a free download on ⇒ www.passcollection.com ⇐ □GES-C01 Reliable Test Camp
- Reliable GES-C01 Guide Files □ GES-C01 Latest Practice Questions □ GES-C01 Advanced Testing Engine □ Open { www.pdfvce.com } and search for ➡ GES-C01 □ to download exam materials for free □Exam GES-C01 Fees
- Knowledge GES-C01 Points □ GES-C01 Exam Simulator Fee □ Exam GES-C01 Fees □ Download ➡ GES-C01 □ for free by simply entering ☀ www.prep4pass.com □☀□ website □New GES-C01 Test Tutorial
- Unparalleled Snowflake GES-C01 Actual Dump: SnowPro® Specialty: Gen AI Certification Exam Pass Guaranteed □ Go to website > www.pdfvce.com < open and search for ➡ GES-C01 □ to download for free □Knowledge GES-C01 Points
- Online Snowflake GES-C01 Practice Test □ Search on ➡ www.real4dumps.com □ for ✓ GES-C01 □✓□ to obtain exam materials for free download □GES-C01 Pdf Pass Leader
- Useful GES-C01 Actual Dump Supply you Realistic Reliable Exam Tutorial for GES-C01: SnowPro® Specialty: Gen AI Certification Exam to Prepare casually □ Search for □ GES-C01 □ and download it for free on 《 www.pdfvce.com 》

website □ Knowledge GES-C01 Points

- Free PDF Quiz 2025 GES-C01: Efficient SnowPro® Specialty: Gen AI Certification Exam Actual Dump □ Easily obtain { GES-C01 } for free download through ☀ www.torrentvalid.com □ ☀ □ Reliable GES-C01 Guide Files
- GES-C01 Latest Practice Questions □ GES-C01 Advanced Testing Engine □ GES-C01 Latest Practice Questions □ Open (www.pdfvce.com) and search for ➡ GES-C01 □ to download exam materials for free □ Latest GES-C01 Exam Format
- Trustworthy GES-C01 Exam Torrent □ GES-C01 Valid Exam Materials □ GES-C01 Pdf Pass Leader □ Search for ➡ GES-C01 □ and easily obtain a free download on □ www.passtestking.com □ □ New GES-C01 Test Tutorial
- www.stes.tyc.edu.tw, cursos.hongency.com, hindi.sachpress.com, thesocraticmethod.in, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, ispausa.org, goaanforex.com, benward394.onzeblog.com, msadvisory.co.zw, tedcole945.iyublog.com, Disposable vapes