

Pass Guaranteed Quiz 2025 Professional Snowflake GES-C01: Reliable SnowPro® Specialty: Gen AI Certification Exam Test Notes



What's more, part of that TrainingQuiz GES-C01 dumps now are free: https://drive.google.com/open?id=1nDGFusT_xBodZeKDnS5iNNXV8r1DaopI

The TrainingQuiz SnowPro® Specialty: Gen AI Certification Exam (GES-C01) PDF dumps file is a collection of real, valid, and updated GES-C01 practice questions that are also easy to install and use. The GES-C01 PDF dumps file can be installed on a desktop computer, laptop, and even on your smartphone devices. Just download TrainingQuiz SnowPro® Specialty: Gen AI Certification Exam in GES-C01 PDF Questions on your desired device and start Snowflake GES-C01 exam dumps preparation today.

TrainingQuiz offers updated GES-C01 questions in a PDF document. These GES-C01 real exam questions come with accurate answers, ensuring reliability and authenticity. The PDF format provides portability, allowing you to study for the Snowflake GES-C01 examination without time and location constraints. You can access the PDF file on your laptop, tablet, or smartphone, making it incredibly convenient.

>> Reliable GES-C01 Test Notes <<

Use Valid Snowflake GES-C01 Dumps PDF And Overcome Your Fear Of Taking Exam

As you can see that on our website, we have free demos of the GES-C01 study materials are freebies for your information. In case you are tentative about their quality, we give these demos from which you could get the brief outline and questions closely related with the GES-C01 Exam Materials. And it is quite easy to free download the demos of the GES-C01 training guide, you can just click on the demos and input your email than you can download them in a second.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q317-Q322):

NEW QUESTION # 317

A data science team is fine-tuning a Snowflake Document AI model to improve the extraction accuracy of specific fields from a new type of complex legal document. They are consistently observing low confidence scores and inconsistent 'value' keys for extracted entities, even after initial training. Which two of the following best practices should the team follow to most effectively improve the model's extraction accuracy and confidence for this complex document type?

- A. Ensure the training dataset used for fine-tuning includes diverse documents representing various layouts, data variations, and explicit examples of values or empty cells where appropriate.
- B. Actively involve subject matter experts (SMEs) or document owners throughout the iterative process to help define data values, provide annotations, and evaluate the model's effectiveness.
- C. Set the 'temperature' parameter to a higher value (e.g., 0.7) during '!PREDICT' calls to encourage more creative and

diverse interpretations by the model.

- D. Limit the fine-tuning training data exclusively to perfectly formatted and clean documents to ensure the model learns from ideal examples without noise.
- E. Prioritize extensive prompt engineering by creating highly detailed and complex questions with intricate logic to guide the LLM's understanding of the extraction task.

Answer: A,B

Explanation:

To improve Document AI model training, it is crucial to ensure that the documents uploaded for training represent a real use case and that the dataset consists of diverse documents in terms of both layout and data. If all documents contain the same data or are always presented in the same form, the model might provide incorrect results. For table extraction, it is vital that enough data is used to train the model to include 'NULL' values and maintain order. Therefore, ensuring a diverse training dataset (Option B) is a key best practice. Additionally, Subject Matter Experts (SMEs) and document owners are crucial partners in understanding and evaluating the model's effectiveness in extracting the required information. Their involvement in defining data values, providing annotations, and evaluating results will significantly improve accuracy (Option C). Option A is not a best practice; it's recommended to keep questions as encompassing as possible and rely on training with annotations rather than complex prompt engineering especially for document variability. Option D is incorrect; a higher 'temperature' value increases the randomness and diversity of the model's output, which is generally undesirable for accurate data extraction where deterministic results are preferred. For most consistent results, 'temperature' should be set to 0. Option E is incorrect because training on a restricted set of perfectly formatted documents can lead to a model that performs poorly on real-world, varied documents; diversity in training data is essential.

NEW QUESTION # 318

An analytics team needs to use various Snowflake Cortex LLM functions and wants to understand the cost implications and governance mechanisms for controlling model access and usage. The "ACCOUNTADMIN" has restricted general LLM access using the 'CORTEX_MODELS_ALLOWLIST' parameter. The team is particularly interested in how model complexity and specific function calls impact billing and how access can be managed effectively for different user groups.

- When using 'AI_COMPLETE' with  structured outputs, complex JSON schemas result in higher token consumption and thus increased costs, but there is no additional charge for the schema validation overhead.
- The 'CORTEX_MODELS_ALLOWLIST' parameter, once set by 'ACCOUNTADMIN', applies uniformly to all Snowflake Cortex AI features, including Cortex Analyst, Cortex Search, and Document AI, restricting their use of LLMs to only those specified.
- For functions like 'AI_CLASSIFY' and 'AI_EXTRACT', both input and output tokens are billable, and the input token count may be higher than the text provided due to an added prompt by Snowflake.
- The 'SNOWFLAKE.CORTEX_USER' database role grants users the necessary privileges to call all Snowflake Cortex AI functions, and this role alone is sufficient to bypass any 'CORTEX_MODELS_ALLOWLIST' restrictions.
- Usage of Cortex LLM functions, including token consumption for different models, can be monitored through the 'SNOWFLAKE.ACOUNT_USAGE.CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY' view, which provides granular usage information grouped by the models used in a query.

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

Answer: A,B,E

Explanation:

Option A is correct because 'AI_COMPLETE' Structured Outputs incurs compute cost based on the number of tokens processed. The number of tokens increases with schema complexity, meaning larger and more complex schemas consume more input and output tokens, but there is no additional compute cost for the overhead of verifying each token against the supplied JSON schema. Option C is correct because for functions that generate new text, such as 'AI_CLASSIFY' (and its previous version 'CLASSIFY_TEXT') and 'AI_EXTRACT' (and its previous version 'EXTRACT_ANSWER'), both input and output tokens are billable. Additionally, these functions add a prompt to the input text, increasing the input token count beyond the user-provided text. Option E is correct as the 'SNOWFLAKE.ACOUNT_USAGE.CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY' view allows granular monitoring of Cortex LLM function usage, grouped by the models used in a query. Option B is incorrect because the 'CORTEX_MODELS_ALLOWLIST' specifically applies to Cortex LLM Functions, Cortex LLM Playground, and the Cortex REST API, not all Snowflake Cortex AI features universally (e.g., Cortex Analyst's model selection is managed differently, Document AI uses its proprietary model). Option D is incorrect because while 'SNOWFLAKE.CORTEX_USER' grants necessary privileges, it does not bypass the 'CORTEX_MODELS_ALLOWLIST' for plain model names; the allowlist provides an account-level restriction that applies unless a schema-level model object is explicitly used with RBAC.

NEW QUESTION # 319

A Snowflake team observes consistently high token costs from 'SNOWFLAKE.ACCOUNT_USAGE.CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY' for a summarization task using the 'mistral-large2' model. The task involves summarizing legal documents, which often exceed the context window of common LLMs. To optimize these token-based costs, which strategy should the team prioritize?

Increase the virtual warehouse size (e.g., from X-SMALL to MEDIUM) used for running the summarization queries to boost performance and reduce overall cost per query.

Switch from using the `COMPLETE` function to `TRY_COMPLETE` to automatically avoid billing for queries that fail due to context window limits, thereby reducing costs.

Implement a text splitting mechanism, potentially using `SPLIT_TEXT_RECURSIVE_CHARACTER`, to break down lengthy documents into smaller chunks before passing them to the summarization function, then aggregate the summaries.

Enable Cortex Guard for the `COMPLETE` function calls, as its filtering capabilities automatically reduce the number of tokens processed for unsafe content.

Set the `temperature` parameter to 0 in the `COMPIETE` function options to ensure more deterministic and thus more cost-efficient summarization outputs.

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

Answer: B

Explanation:

Option C is correct. For summarization of lengthy documents, exceeding the context window or using large inputs significantly increases token consumption. Text splitting, for example using `SPLIT_TEXT_RECURSIVE_CHARACTER`, can break documents into smaller, more manageable chunks. This reduces the number of input tokens per LLM call, directly leading to cost optimization, and is recommended for best search results and LLM response quality with Cortex Search. Option A is incorrect because for Cortex AISQL functions, Snowflake recommends using a smaller warehouse (no larger than MEDIUM) as larger warehouses do not increase performance but can result in unnecessary costs associated with keeping the warehouse active. The compute cost for Cortex LLM functions is based on tokens processed, not warehouse size performance. Option B is incorrect because 'TRY COMPLETE' only prevents costs for 'failed' operations by returning NULL instead of an error. It does not optimize the token consumption of 'successful' summarization tasks. Option D is incorrect; Cortex Guard processes additional tokens for its filtering, thus 'increasing' token consumption, not reducing it. Option E is incorrect because setting 'temperature' to 0 makes the output more deterministic, which might improve consistency but does not directly reduce the number of input or output tokens processed for a summarization task.

NEW QUESTION # 320

An operations team is investigating an issue with a generative AI application powered by Snowflake Cortex Analyst, where users reported unexpected behavior in generated SQL. To diagnose the problem, they examine the detailed event logs captured by Snowflake AI Observability. Which categories of information can they expect to find in these event tables to assist their investigation?

- A. Real-time CPU and memory usage statistics for the Snowflake virtual warehouse executing the LLM inference.
- B. The exact SQL queries that Cortex Analyst generated in response to user questions.
- C. The full text of the natural language questions submitted by the users.
- D. The complete request and response bodies associated with the application's execution steps.
- E. Any error messages or warnings that occurred during the processing of the request.

Answer: B,C,D,E

Explanation:

Cortex Analyst logs requests to an event table to aid in refining semantic models or views. These logs are comprehensive and include specific details crucial for debugging and monitoring. The captured information includes 'The user who asked the question', 'The question asked', 'Generated SQL', 'Errors and/or warnings', 'Request and response bodies', and 'Other metadata'. Therefore, options A, B, C, and D are all accurate descriptions of the data available in these event logs. Option E, real-time CPU and memory usage, refers to infrastructure monitoring metrics rather than the content specifically logged within the application's execution event table by Cortex Analyst itself.

NEW QUESTION # 321

A new Gen AI specialist is setting up Document AI. They have successfully created the necessary database, schema, and a custom role named 'doc_ai_specialist_role'. This custom role has been granted the 'SNOWFLAKE.DOCUMENT_INTELLIGENCE_CREATOR database role. However, when the specialist attempts to create a new Document AI model build in Snowsight, they receive the error: Unable to create a build on the specified database and schema. Please check the documentation to learn more. What is the most likely missing privilege for the that is preventing the model build creation?

- A.
- B.
The 'doc_ai_specialist_role' is missing 'CREATE SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE' privilege on the schema where the model build is being created.
- C. The virtual warehouse assigned to the Document AI environment is not large enough to support model build creation.
- D.
- E.
The 'doc_ai_specialist_role' is missing 'USAGE' privilege on the database where the model build is being created.

Answer: B

Explanation:

The error message 'Unable to create a build on the specified database and schema' points directly to insufficient privileges for creating the model build itself. While the database role enables working on Document AI models and using SQL for extraction, specific object privileges are required to create the model 'build' in a given schema. The documentation explicitly states that 'CREATE SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE on the schema is required to create model builds (instances of the 'DOCUMENT_INTELLIGENCE' class), along with 'CREATE MODEL' on the schema. Option A (USAGE on database) is a general prerequisite but not the direct cause of this specific 'unable to create a build' error. Option B (CREATE TABLE) is for creating tables, not model builds. Option D (EXECUTE TASK) is required for running tasks in a pipeline, not for model build creation. Option E is incorrect as warehouse size primarily impacts cost and performance, not the ability to create a model build due to privileges.

NEW QUESTION # 322

.....

Do you want to attend Snowflake GES-C01 test? Are you worried about GES-C01 exam? You want to sign up for GES-C01 certification exam, but you are worried about failing the exam. Do you have such situations? Don't worry and sign up for GES-C01 exam. As long as you make use of TrainingQuiz certification training materials, particularly difficult exams are not a problem. Even if you have never confidence to pass the exam, TrainingQuiz also guarantees to Pass GES-C01 Test at the first attempt. Is it inconceivable? You can visit TrainingQuiz.com to know more details. In addition, you can try part of TrainingQuiz GES-C01 exam dumps. By it, you will know that the materials are your absolute guarantee to pass the test easily.

GES-C01 Exam Training: <https://www.trainingquiz.com/GES-C01-practice-quiz.html>

You may doubtful if you are newbie for our GES-C01 training engine, free demos are provided for your reference, Desktop SnowPro® Specialty: Gen AI Certification Exam (GES-C01) practice test software is the first format available at TrainingQuiz, As one of the most famous company in the market, we are being popular for responsible services (GES-C01 training materials), Snowflake Reliable GES-C01 Test Notes We believe that the free demo will give you a brand new experience.

Click next to move to the next screen, which involves GES-C01 preparing the Spyder to calibrate your display, If not, this section serves as a primer or a good review, You may doubtful if you are newbie for our GES-C01 training engine, free demos are provided for your reference.

Free PDF Quiz Valid GES-C01 - Reliable SnowPro® Specialty: Gen AI Certification Exam Test Notes

Desktop SnowPro® Specialty: Gen AI Certification Exam (GES-C01) practice test software is the first format available at TrainingQuiz, As one of the most famous company in the market, we are being popular for responsible services (GES-C01 training materials).

We believe that the free demo will give you a brand new experience, And our products are global, and you can purchase our GES-

C01 training guide is wherever you are.

- Exam GES-C01 Torrent □ New GES-C01 Braindumps Questions □ Exam GES-C01 Torrent □ Easily obtain free download of □ GES-C01 □ by searching on 「 www.free4dump.com 」 □ GES-C01 Valid Exam Preparation
- GES-C01 Reliable Real Exam □ GES-C01 Exam Prep □ Exam GES-C01 Torrent ♣ Open website ➤ www.pdfvce.com □ and search for ➡ GES-C01 □ for free download □ New GES-C01 Exam Test
- GES-C01 Reliable Exam Review □ GES-C01 Exam Prep □ GES-C01 Reliable Real Exam □ The page for free download of ▷ GES-C01 ▶ on ☀ www.pass4test.com ☀ ☀ will open immediately □ Reliable GES-C01 Test Topics
- GES-C01 Instant Discount □ GES-C01 Exam Engine □ GES-C01 Test King □ Search on 【 www.pdfvce.com 】 for ➡ GES-C01 □ to obtain exam materials for free download □ Cert GES-C01 Guide
- Excellent Reliable GES-C01 Test Notes, GES-C01 Exam Training □ Copy URL ➡ www.testkingpdf.com ⇄ open and search for ➡ GES-C01 □ to download for free □ GES-C01 Valid Exam Simulator
- Free PDF 2025 Unparalleled Snowflake GES-C01: Reliable SnowPro® Specialty: Gen AI Certification Exam Test Notes □ □ Search for ✓ GES-C01 □ ✓ □ and download it for free immediately on ▷ www.pdfvce.com ▷ □ Test GES-C01 Pass4sure
- GES-C01 Reliable Real Exam □ New GES-C01 Exam Test □ GES-C01 Test King ↗ Search for □ GES-C01 □ and download exam materials for free through “www.itcerttest.com” □ Valid GES-C01 Exam Topics
- Excellent Reliable GES-C01 Test Notes, GES-C01 Exam Training □ The page for free download of □ GES-C01 □ on (www.pdfvce.com) will open immediately □ GES-C01 Reliable Real Exam
- GES-C01 Reliable Exam Review □ GES-C01 Valid Exam Simulator □ GES-C01 Reliable Dumps Ppt □ ☀ www.prep4away.com ☀ ☀ is best website to obtain ➡ GES-C01 □ □ □ for free download □ GES-C01 Test King
- New GES-C01 Exam Test □ GES-C01 Vce Torrent □ Exam GES-C01 Torrent □ Search for ➤ GES-C01 □ on □ www.pdfvce.com □ immediately to obtain a free download □ GES-C01 Exam Paper Pdf
- GES-C01 Exam Engine □ New GES-C01 Exam Test □ Valid GES-C01 Exam Topics □ Easily obtain free download of “GES-C01” by searching on 《 www.vceengine.com 》 □ Cert GES-C01 Guide
- www.stes.tyc.edu.tw, vxlxemito123.hashnode.dev, www.taowang.com, myportal.utt.edu.tt, motionentrance.edu.np, study.stcs.edu.np, xjj3.cc, elearning.eauqardho.edu.so, mr.magedgerges.mathewmaged.com, capacitacion.axiomamexico.com.mx, Disposable vapes

P.S. Free 2025 Snowflake GES-C01 dumps are available on Google Drive shared by TrainingQuiz: https://drive.google.com/open?id=1nDGFusT_xBodZeKDnS5iNNXV8r1DaopI