GES-C01 Latest Study Questions & GES-C01 Original Questions

Option C is correct. In a production data pipeline, 'TRY_COMPLETE is preferred over 'AI_COMPLETE for robustness. If the model fails to generate a valid response (e.g., cannot adhere to the schema or encounters an internal error), 'TRY COMPLETE' returns 'NULL' instead of raising an error, allowing the pipeline to continue processing other records without interruption.
Option D is incorrect. The number of tokens processed (and thus billed) for 'AI_COMPLETE Structured Outputs does increase with schema complexity. Highly-structured responses, especially those with deep nesting, consume a larger number of input and output tokens.
Option E is incorrect. The specific requirements for 'additionalProperties' being 'false' and the 'required' field containing all property names only apply to OpenAI (GPT) models when using Structured Outputs. Other models do not strictly enforce these requirements, although including them may simplify schema management across different model types.

22.A Gen AI Specialist is setting up their Snowflake environment to deploy a high-performance open-source LLM for real-time inference using Snowpark Container Services (SPCS). They need to create a compute pool that can leverage NVIDIAAIOG GPUs to optimize model performance.
Which of the following SQL statements correctly creates a compute pool capable of supporting an intensive GPU usage scenario, such as serving LLMs, while adhering to common configuration best practices for a new, small-scale deployment in Snowpark Container Services?
A)

```
CREATE COMPUTE POOL llm_gpu_pool
   MIN_NODES = 1
   MAX_NODES = 2
   INSTANCE_FAMILY = CPU_X64_M
   AUTO_SUSPEND_SECS = 3600;
```

B)

```
CREATE COMPUTE POOL llm_gpu_pool
   MIN_NODES = 1
   MAX_NODES = 1
   INSTANCE_FAMILY = GPU_NV_S
   AUTO_RESUME = TRUE;
```

C)

```
CREATE COMPUTE POOL llm_gpu_pool
   MIN_NODES = 2
   MAX_NODES = 4
   INSTANCE_FAMILY = HIGHMEM_X64_L
   AUTO_SUSPEND_SECS = 0;
```

D)

19 / 23

What's more, part of that BraindumpsPrep GES-C01 dumps now are free: https://drive.google.com/open?id=1-GJ8YoidcdqSNDu9xzgnEB86VA04HqUN

When you first contact our software, different people will have different problems. Maybe you are not comfortable with our GES-C01 exam question and want to know more about our products and operations. As long as you have questions, you can send e-mail to us, we have online staff responsible for ensuring 24-hour service to help you solve all the problems about our GES-C01 Test Prep. After you purchase our GES-C01 quiz guide, we will still provide you with considerate services. Maybe you will ask whether we will charge additional service fees.

BraindumpsPrep provides numerous extra features to help you succeed on the GES-C01 exam, in addition to the Snowflake GES-C01 exam questions in PDF format and online practice test engine. These include 100% real questions and accurate answers, 1 year of free updates, a free demo of the Snowflake GES-C01 Exam Questions, a money-back guarantee in the event of failure, and a 20% discount. BraindumpsPrep is the ideal alternative for your GES-C01 test preparation because it combines all of these elements.

>> GES-C01 Latest Study Questions <<

Pass Guaranteed Quiz High Pass-Rate Snowflake - GES-C01 Latest Study Questions

With GES-C01 study engine, you will get rid of the dilemma that you work hard but cannot improve. With our GES-C01 learning materials, you can spend less time but learn more knowledge than others. GES-C01 exam questions will help you reach the peak of your career. Just think of that after you get the GES-C01 Certification, you will have a lot of opportunities of going to biger and better company and getting higher incomes! what a brighter future!

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q18-Q23):

NEW QUESTION # 18
A retail company wants to implement an automated data pipeline in Snowflake to analyze daily customer reviews. The goal is to enrich a 'product_reviews_sentiment' table with sentiment categories (e.g., 'positive', 'neutral', 'negative') for each new review. They require the sentiment to be returned as a JSON object for downstream processing and need the pipeline to handle potential LLM errors gracefully without stopping. Assuming a stream 'new reviews_stream' monitors a 'customer _ reviews' table, which approach effectively uses a Snowflake Cortex function for this scenario?

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

Answer: A

Explanation:
Option C is the most effective approach for this scenario. It correctly uses "SNOWFLAKE.CORTEX.TRY COMPLETE", which performs the same operation as 'COMPLETE but returns NULL instead of raising an error when the operation cannot be performed, making the pipeline more robust to LLM issues. The 'response_format option ensures the output adheres to a specified JSON schema for structured sentiment categories, meeting the requirement for structured output. This is integrated within a 'MERGE statement in a task for incremental processing of new data from Option A suggests a Python UDF with 'COMPLETE. While feasible, 'TRY_COMPLETE is explicitly designed for graceful error handling in pipelines, which 'COMPLETE lacks by default. Option B uses 'SNOWFLAKE.CORTEX.SENTIMENT , which returns a numeric score (e.g., 0.5424458), not a categorical JSON object, requiring additional post-processing logic for categorization. Option D uses for summarization and 'AI CLASSIFY for classification. While 'AI_CLASSIFY can categorize, the request is for sentiment of 'each' review, and 'AI_AGG' would aggregate before classifying, not fulfilling the individual review sentiment requirement. Option E suggests a dynamic table, but dynamic tables currently do not support incremental refresh with 'COMPLETE (or 'AI_COMPLETE) functions, making them unsuitable for continuous LLM-based processing in this manner. Furthermore, 'COMPLETE does not offer the graceful error handling of 'TRY COMPLETE'.

NEW QUESTION # 19
A data processing team is using Snowflake Document AI to extract data from incoming supplier invoices. They observe that many documents are failing to process, and successful extractions are taking longer than expected, leading to increased costs. Upon investigation, they find error messages such as
. Additionally, their 'X-LARGE virtual warehouse is constantly active, contributing to higher-than-anticipated bills. Which two of the following actions are essential steps to troubleshoot and address the root causes of these processing errors and optimize their Document AI pipeline?

- A. Increase the 'max_tokens' parameter within the '!PREDICT' function options to accommodate longer document responses from the model.
- B. Redefine extraction questions to be more generic and encompassing, reducing the number of distinct questions needed per document.
- C. Implement a pre-processing step to split documents exceeding 125 pages or 50 MB into smaller, compliant files before loading to the stage.
- D. Configure the internal stage used for storing invoices with 'ENCRYPTION = (TYPE = 'SNOWFLAKE SSE')'.
- E. Scale down the virtual warehouse to 'X-SMALL' or 'SMALL' size, as larger warehouses do not increase Document AI query processing speed and incur unnecessary costs.

Answer: C,D

Explanation:

The error messages 'Document has too many pages. Actual: 130. Maximum: 125.' and File exceeds maximum size. Actual: 54096026 bytes. Maximum: 50000000 bytes: directly indicate that the documents do not meet Document AI's input requirements, which specify a maximum of 125 pages and 50 MB file size. Therefore, implementing a pre-processing step to split or resize these documents is an essential solution (Option B). The error 'cannot identify image file < io.BytesIO object at Ox...>' is a known issue that occurs when an internal stage used for Document AI is not configured with encryption. Correctly configuring the stage with this encryption type is crucial for resolving this processing error (Option D). Option A, while addressing cost optimization, is not a root cause of the 'processing errors' themselves, although it is a best practice for cost governance as larger warehouses do not increase Document AI query processing speed. Option C is incorrect; best practices for question optimization suggest being specific, not generic. Option E is incorrect as 'max_tokenS relates to the length of the model's output, not the input document's size or page limits.

NEW QUESTION # 20
An enterprise is deploying a new RAG application using Snowflake Cortex Search on a large dataset of customer support tickets. The operations team is concerned about managing compute costs and ensuring efficient index refreshes for the Cortex Search Service, which needs to be updated hourly. Which of the following considerations and configurations are relevant for optimizing cost and performance of the Cortex Search Service in this scenario?

- A. CHANGE_TRACKING
- B. The primary cost driver for Cortex Search is the number of search queries executed against the service, with the volume of indexed data (GB/month) having a minimal impact on overall billing.
- C. For embedding text, selecting a model like
- D. The
- E. For optimal performance and cost efficiency, Snowflake recommends using a dedicated warehouse of size no larger than MEDIUM for each Cortex Search Service.

Answer: A,C,D,E

Explanation:
Option A is correct because a Cortex Search Service requires a virtual warehouse to refresh the service, which runs queries against base objects when they are initialized and refreshed, incurring compute costs. Option B is correct because the cost of embedding models varies. For example, 'snowflake-arctic-embed-m-v1.5 costs 0.03 credits per million tokens, while 'voyage-multilingual-2 costs 0.07 credits per million tokens. Choosing a more cost-effective model like 'snowflake-arctic-embed-m-v1.5' for English-only data can reduce token costs. Option C is correct because Snowflake recommends using a dedicated warehouse of size no larger than MEDIUM for each Cortex Search Service to achieve optimal performance. Option D is correct because change tracking is required for the Cortex Search Service to be able to detect and process updates to the base table, enabling incremental refreshes that are more efficient than full re-indexing. Option E is incorrect because Cortex Search Services incur costs based on virtual warehouse compute for refreshes, 'EMBED_TEXT_TOKENS' cost per input token, and a charge of 6.3 Credits per GB/mo of indexed data. The volume of indexed data has a significant impact, not minimal.

NEW QUESTION # 21
A data engineering team is building a pipeline to process legal documents using Snowflake Cortex functions. They aim to extract specific entities and summarize key clauses while being highly cost-conscious. To optimize token-based costs, which of the following practices should they implement when using Cortex LLM functions?

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

Answer: A,D

Explanation:

NEW QUESTION # 22
An ML engineer is deploying a custom PyTorch-based image classification model, obtained from Hugging Face, to Snowpark Container Services (SPCS). The deployment requires GPU acceleration on a compute pool named 'my_gpu_pool' and specific

Python packages ('torch', 'transformerS, 'opencv-python'). The scenario dictates that 'opencv-python' is only available via PyPI, while 'torch' and 'transformers' can be sourced from either conda-forge or PyPI. The engineer uses the Snowflake Model Registry to log the model. Which of the following 'log model' and 'create_service' configurations correctly specify the necessary Python dependencies and GPU utilization for this inference service, adhering to Snowflake's recommendations?

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

Answer: D

Explanation:
Option A is correct. The 'pip_requirements' argument can be used to specify all necessary Python packages, including 'torch 'transformers', and 'opencv-python', which are commonly available on PyPI. The 'create_service' call correctly specifies and to leverage GPU acceleration, as SPCS supports GPU instances for ML workloads. This approach aligns with the Snowflake recommendation to use either 'conda_dependencies' or 'pip_requirements', but not both, for dependency management. Option B is incorrect because 'opencv-python' is specified as only available via PyPI in the scenario, meaning it cannot be installed via 'conda-forge' . While 'conda-forge' is the default for SPCS dependencies, not all packages are available there. Option C is incorrect because is chosen, which will not provide GPU acceleration required by the PyTorch model for GPU-powered inference. GPU acceleration requires a GPU compute pool. Option D is incorrect because 'opencv-python' is not available through Anaconda channels (as per the scenario that it is PyPI only), and for other conda packages, explicitly specifying the 'conda- forge' channel (e.g., is the recommended practice for SPCS dependencies if they are not in the Snowflake Anaconda channel. Option E is incorrect because, while it correctly separates conda and pip dependencies for specific packages, Snowflake explicitly recommends 'using only 'conda_dependencies' or only 'pip_requirements', not both' for managing dependencies to avoid potential conflicts.

NEW QUESTION # 23
.....

You may want to know our different versions of GES-C01 exam questions. Firstly, PDF version is easy to read and print. Secondly software version simulates the real GES-C01 actual test guide, but it can only run on Windows operating system. Thirdly, online version supports for any electronic equipment and also supports offline use. For the first time, you need to open GES-C01 Exam Questions in online environment, and then you can use it offline. All in all, helping our candidates to pass the exam successfully is what we always looking for. Our GES-C01 actual test guide is your best choice.

GES-C01 Original Questions: https://www.briandumpsprep.com/GES-C01-prep-exam-braindumps.html

If you have got GES-C01 latest dumps, your IT professional ability will be approved by a lot of IT companies, Our GES-C01 training materials are designed carefully, Though the content is the same with all versions of the GES-C01 study materials, the displays are totally different, GES-C01 exam cram is high-quality, and you can pass your exam by using them, Snowflake GES-C01 Latest Study Questions The first and the most important aspect is the pass rate which is concerned by the most customers, we have a high pas rate as 98% to 100%, which is unique in the market!

But according to Hartman it is often described as magical" by diehard fans, Free BraindumpsPrep GES-C01 Exam BraindumpsPrep Practice Test Demo is Worth a Try, If you have got GES-C01 latest dumps, your IT professional ability will be approved by a lot of IT companies.

Want to Know Your Readiness for Snowflake GES-C01 Exam? Take Our Online Practice Test

Our GES-C01 training materials are designed carefully, Though the content is the same with all versions of the GES-C01 study materials, the displays are totally different.

GES-C01 exam cram is high-quality, and you can pass your exam by using them, The first and the most important aspect is the pass rate which is concerned by the most GES-C01 customers, we have a high pas rate as 98% to 100%, which is unique in the market!

- Simulations GES-C01 Pdf □ GES-C01 Valid Dump □ GES-C01 Test Engine □ Enter ➡ www.troytecdumps.com □□□ and search for ➤ GES-C01 □ to download for free □GES-C01 Test Braindumps
- GES-C01 Latest Exam Answers □ GES-C01 Exam Fee □ Test GES-C01 Valid □ Simply search for 【 GES-C01 】

for free download on ✔ www.pdfvce.com □✔□ □Test GES-C01 Valid

- GES-C01 Exam Engine !! GES-C01 Reliable Test Dumps □ Simulations GES-C01 Pdf □ Enter （ www.easy4engine.com ） and search for 【 GES-C01 】 to download for free □GES-C01 Reliable Exam Simulations
- GES-C01 Exams Training □ GES-C01 Valid Dumps Book □ GES-C01 Reliable Exam Simulations □ Enter ➡ www.pdfvce.com □□□ and search for 【 GES-C01 】 to download for free □Test GES-C01 Valid
- Valid GES-C01 Exam Experience □ Valid GES-C01 Exam Experience □ New GES-C01 Test Experience □ Simply search for □ GES-C01 □ for free download on □ www.troytecdumps.com □ □GES-C01 Test Questions Vce
- Quiz Snowflake - GES-C01 - Pass-Sure SnowPro® Specialty: Gen AI Certification Exam Latest Study Questions □ Easily obtain free download of ➡ GES-C01 □□□ by searching on 【 www.pdfvce.com 】 □GES-C01 Reliable Test Dumps
- GES-C01 Latest Exam Answers □ GES-C01 New Braindumps □ GES-C01 Test Braindumps □ Immediately open ➡ www.pdfdumps.com □□□ and search for 「 GES-C01 」 to obtain a free download □GES-C01 Test Questions Vce
- GES-C01 Valid Mock Test □ GES-C01 Test Questions Vce □ GES-C01 Valid Dumps Book □ Search on { www.pdfvce.com } for " GES-C01 " to obtain exam materials for free download □GES-C01 Reliable Exam Simulations
- GES-C01 Test Questions Vce □ GES-C01 Reliable Exam Simulations □ GES-C01 Reliable Exam Simulations □ Enter 【 www.prep4away.com 】 and search for ⇛ GES-C01 ⇚ to download for free □GES-C01 Test Price
- GES-C01 Valid Mock Test □ GES-C01 Valid Dump □ Simulations GES-C01 Pdf □ Search for ➽ GES-C01 □ and obtain a free download on { www.pdfvce.com } □GES-C01 New Braindumps
- GES-C01 Dumps Guide: SnowPro® Specialty: Gen AI Certification Exam - GES-C01 Actual Test - GES-C01 Exam Torrent □ Go to website ✔ www.prepawaypdf.com □✔□ open and search for □ GES-C01 □ to download for free □ □GES-C01 Valid Dump
- www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, bbs.t-firefly.com, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, Disposable vapes

P.S. Free & New GES-C01 dumps are available on Google Drive shared by BraindumpsPrep: https://drive.google.com/open?id=1-GJ8YoidcdqSNDu9xzgnEB86VA04HqUN