Snowflake GES-C01 Reliable Braindumps - Test GES-C01 Dump


For 'RAG_Config_A', instrument the 'generate_answer' function with @instrument(span_type=SpanAttributes.SpanType.GENERATION). For 'RAG_Config_B', instrument its equivalent 'generate_answer' function similarly, and ensure both are registered as part of distinct 'TruApp' versions or runs for comparison.

Enable cross-region inference using the CORTEX_ENABLED_CROSS_REGION parameter to ensure both 'llama3.1-8b' and 'mistral-7b' models are available, as this directly enables the comparison feature within AI Observability.

Instrument the context retrieval component in both configurations with @instrument(span_type=SpanAttributes.SpanType.RETRIEVAL) to allow for the calculation of 'context_relevance' metrics for each, which can then be used in comparative evaluations.

Create separate runs (using add_run() with distinct run_name or Label in the RunConfig) for each RAG configuration, specifying the respective LLM as llm_judge_name, and explicitly list 'answer_relevance' and 'groundedness' in the metrics parameter when calling compute_metrics().

Focus solely on 'prompt_tokens' and 'completion_tokens' via the 'CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY' view, as these metrics provide the most direct comparison of LLM performance for RAG applications.

BTW, DOWNLOAD part of TroytecDumps GES-C01 dumps from Cloud Storage: https://drive.google.com/open?id=1_oeQSD6_uH9yISd_wSmv6jbo5tTIXpn_

To avail of all these benefits you need to pass the GES-C01 exam which is a difficult exam that demands firm commitment and complete GES-C01 exam questions preparation. For the well and quick GES-C01 exam dumps preparation, you can get help from TroytecDumps GES-C01 Questions which will provide you with everything that you need to learn, prepare and pass the SnowPro® Specialty: Gen AI Certification Exam certification exam.

We are aimed to develop a long-lasting and reliable relationship with our customers who are willing to purchase our GES-C01 study materials. To enhance the cooperation built on mutual-trust, we will renovate and update our system for free so that our customers can keep on practicing our GES-C01 study materials without any extra fee. Meanwhile, to ensure that our customers have greater chance to pass the exam, we will make our GES-C01 test training keeps pace with the digitized world that change with each passing day. In this way, our endeavor will facilitate your learning as you can gain the newest information on a daily basis and keep being informed of any changes in GES-C01 test. Therefore, our customers can save their limited time and energy to stay focused on their study as we are in charge of the updating of our GES-C01 test training. It is our privilege and responsibility to render a good service to our honorable customers.

>> Snowflake GES-C01 Reliable Braindumps <<

GES-C01 Actual Test Guide Boosts Most efficient Exam Questions for Your SnowPro® Specialty: Gen AI Certification Exam Exam

Comparing to other training institution, our valid GES-C01 vce dumps are affordable, latest and cost-effective, which can overcome the difficulty of valid GES-C01 Actual Test and ensure you pass the exam. It can not only save your time and money, but also help you clear Snowflake practice exam with high rate.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q45-Q50):

NEW QUESTION # 45
A data application developer is tasked with building a multi-turn conversational AI application using Streamlit in Snowflake (SiS) that leverages the COMPLETE (SNOWFLAKE. CORTEX) LLM function. To ensure the conversation flows naturally and the LLM maintains context from previous interactions, which of the following is the most appropriate method for handling and passing the conversation history?

The developer should store the entire conversation history in a temporary table in Snowflake and query it with each new turn, passing only the latest user message to the COMPLETE function.

Snowflake automatically manages conversational context for COMPLETE within the session, so the developer only needs to pass the current user prompt as a string.

The conversation history must be explicitly managed within the Streamlit application's state, typically by initializing st.session_state.messages = [] and appending each user and assistant message as an object with 'role' and 'content' keys, then passing the full list to the prompt_or_history argument of COMPLETE.

The developer should concatenate all previous user prompts and assistant responses into a single, long string, and pass this as the <prompt> argument to COMPLETE for each turn.

The COMPLETE function has an optional 'conversation_id' parameter that automatically retrieves and manages conversation history when provided.

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

- E. Option A

Answer: D

NEW QUESTION # 46
A development team plans to utilize Snowpark Container Services (SPCS) for deploying a variety of AI/ML workloads, including custom LLMs and GPU-accelerated model training jobs. They are in the process of creating a compute pool and need to select the appropriate instance families and configurations. Which of the following statements about 'CREATE COMPUTE POOL' in SPCS are accurate?

- A. To support GPU-accelerated LLM inference and training, the 'INSTANCE_FAMILY' must be selected from a type starting with 'GPU' (e.g.,
  `GPU_NV_S`
  `or`
  `GPU_GCP_NV_L4_1_24G`
  `).`
- B. Setting 'AUTO RESUME = TRUE ensures that the compute pool automatically starts when a service or job is submitted to it, rather than requiring manual resumption.
- C. For cost optimization, 'AUTO SUSPEND SECS = 0' should be used to prevent automatic suspension of the compute pool, as suspension and resumption incur minimum billing durations.
- D. The 'MIN NODES' and 'MAX NODES parameters define the scaling range for the compute pool, and Snowflake automatically scales the pool within this range based on workload demand.
- E. Snowpark-optimized warehouses are the recommended compute pool type for all large-scale ML training workloads within SPCS due to their enhanced memory limits and CPU architectures.

Answer: A,B

Explanation:
Option A is correct. GPU-accelerated workloads, such as LLM inference and model training, require instance families specifically designed with GPUs. The documentation lists instance family names starting with 'GPU' for this purpose, such as or 'GPU_GCP NV L4 Option B is incorrect. While 'MIN NODES and 'MAX NODES define the range, the size of compute clusters in Snowpark Container Services does "not" auto-scale dynamically based on workload demand. Users must manually alter the number of instances at runtime using commands like 'ALTER SERVICE MIN INSTANCES = s. Snowflake does handle load balancing across instances within the configured node counts. Option C is correct. The 'AUTO_RESUME = TRUE parameter, when specified during compute pool creation, enables the pool to automatically resume operation when a service or job is submitted, removing the need for explicit SALTER COMPUTE POOL RESUME commands. option D is incorrect. Setting = prevents the compute pool from automatically suspending, meaning it will continue to consume credits even when idle. This would generally lead to higher costs, not cost optimization, unless the pool is constantly active. The default is 3600 seconds (1 hour). SPCS Compute Nodes have a minimum charge of five minutes when started or resumed, making intelligent use of auto-suspend important for cost management. Option E is incorrect. Snowpark-optimized warehouses are a type of 'virtual warehouse' and are recommended for Snowpark workloads with large memory requirements or specific CPU architecture, typically for single-node ML training workloads 'within a warehouse'. SPCS compute pools, however, provide their own dedicated instance families (CPU, HighMemory, GPU) for containerized workloads, abstracting the underlying infrastructure and supporting distributed GPU clusters directly within SPCS, not Snowpark-optimized warehouses as a 'compute pool type' for SPCS.

NEW QUESTION # 47
A data science team operating in the AWS Europe Central 1 (Frankfurt) region needs to leverage the `snowflake-llama-3.1-405b` model for complex generative AI tasks using the `AI_COMPLETE` function. They've noted that `snowflake-llama-3.1-405b` is not natively available in Frankfurt but is supported via cross-region inference from AWS US regions. The `ACCOUNTADMIN` has already configured the account to include `snowflake-llama-3.1-405b` in the `CORTEX_MODELS_ALLOWLIST`. Despite this, users are consistently encountering errors indicating 'model not found' or 'region not supported'. Which specific Snowflake account parameter, and what value, must be additionally configured by the `ACCOUNTADMIN` to enable the use of `snowflake-llama-3.1-405b` in this scenario?

- A. Cross-region inference is automatically managed by Snowflake for allowed models, implying that a new, larger virtual

warehouse is required to handle the cross- region data transfer overhead.

- B.
  The `CORTEX_MODELS_ALLOWLIST` must explicitly list both the source region and the target region for `snowflake-llama-3.1-405b`.
- C.
  A `GRANT USAGE ON LLM snowflake-llama-3.1-405b` statement must be executed for each user role, explicitly enabling cross-region functionality.
- D.
  The `CORTEX_ENABLED_CROSS_REGION` account parameter must be set to `ANY_REGION` or specifically include `AWS_US_EAST_1` to enable inference requests to be processed in a supported region where the model is available.
- E.
  The `ENABLE_CORTEX_ANALYST_MODEL_AZURE_OPENAI` parameter needs to be set to `TRUE` to allow cross-cloud access, even for Snowflake-hosted models.

Answer: D

Explanation:
Option A is correct. The `CORTEX_ENABLED_CROSS_REGION` account parameter controls whether inference requests can be processed in a different region from the default. To use models available via cross-region inference, this parameter must be set to `ANY_REGION` or a list explicitly including the supported cross-region (e.g., an AWS US region for `snowflake-llama-3.1-405b`). Without this setting, even if the model is in the allowlist, the system won't look for it in other regions. Option B is incorrect as `ENABLE_CORTEX_ANALYST_MODEL_AZURE_OPENAI` specifically pertains to legacy Azure OpenAI models for Cortex Analyst and is not relevant for Snowflake-hosted models or general cross-region inference for LLM functions. Option C is incorrect as access to LLMs is controlled by the `CORTEX_MODELS_ALLOWLIST` parameter and the `SNOWFLAKE.CORTEX_USER` role, not individual `GRANT USAGE ON LLM` statements. Option D is incorrect; the `CORTEX_MODELS_ALLOWLIST` specifies model names, not regions or combinations of regions. Option E is incorrect as cross-region inference does not automatically enable larger warehouses; while latency might be a consideration, it's explicitly controlled by the `CORTEX_ENABLED_CROSS_REGION` parameter, and it does not bypass the need for proper configuration.

NEW QUESTION # 48
A data engineering team is building a Retrieval Augmented Generation (RAG) pipeline that heavily relies on 'SNOWFLAKE.CORTEX.EMBED_TEXT 768' to process millions of documents daily. They need to optimize for both cost and retrieval quality. Which of the following statements are true regarding the cost and performance of 'EMBED_TEXT 768' in Snowflake? (Select all that apply)

- A. The 'EMBED TEXT 768' function, regardless of the 768-dimension model used, has a fixed cost of 1.50 Credits per one million Tokens processed.
- B. The 'EMBED_TEXT 768' function is billed based on the number of 'output tokens' generated by the embedding model, as this represents the computational complexity of the vector.
- C. For optimal retrieval quality in RAG scenarios, text should be split into chunks of no more than 512 tokens before being passed to 'EMBED TEXT 768', even if the model supports a larger context window.
- D. The 'snowflake-arctic-embed-m-vl .5 model, used by 'EMBED TEXT 768', has a context window of 512 tokens, and texts exceeding this length are truncated before embedding.
- E. To minimize costs for ' EMBED_TEXT 768 operations, it is recommended to execute queries using a smaller virtual warehouse (no larger than MEDIUM), as larger warehouses do not improve performance for these functions.

Answer: C,D,E

Explanation:
Option B is correct because Snowflake recommends executing queries that call Cortex AISQL functions, including , with a smaller warehouse (no larger than MEDIUM). Larger warehouses do not increase performance for these functions. Option C is correct because the 'snowflake-arctic-embed-m-vl .5 model, which can be used with 'EMBED TEXT 768', has a context window of 512 tokens. Input text exceeding this limit is truncated before embedding. Option D is correct because, for best search results with Cortex Search and RAG, Snowflake recommends splitting the text into chunks of no more than 512 tokens. This smaller chunk size typically results in higher retrieval and downstream LLM response quality. Option A is incorrect because for functions, only 'input tokens' are counted towards the billable total, not output tokens. Option E is incorrect because the cost for models (such as 'e5-base-v2 , snowflake-arctic-embed-m' , 'snowflake-arctic-embed-m-v1.5') is 0.03 Credits per one million Tokens processed. The 1.50 Credits per one million Tokens applies to the ' TRANSLATE' function, not 'EMBED_TEXT_768'.

NEW QUESTION # 49
A development team is preparing to deploy a new Retrieval-Augmented Generation (RAG) application written in Python. They intend to use Snowflake AI Observability with the TruLens SDK to capture detailed logs and traces for debugging and performance analysis. Which of the following configurations are essential prerequisites for enabling this logging capability effectively?

] The Python environment must have the TRULENS_OTEL_TRACING environment variable set to 1 before establishing a connection to Snowflake.

] The Python project requires the installation of `trulens-core`, `trulens-connectors-snowflake`, and `trulens-providers-cortex` packages, with versions 2.1.2 or ater.

] The account role used for the application must be granted both the `SNOWFLAKE.CORTEX_USER` database role and the `AI_OBSERVABILITY_EVENTS_LOOKUP` applicatio ole.

] The application must be designed to run exclusively within a Snowflake Notebook, as this environment automatically streams logs to AI Observability.

] The role performing the operations must have the `CREATE EXTERNAL AGENT` and `CREATE TASK` privileges on the schema where the application objects reside, in ddition to the global `EXECUTE TASK` privilege.

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

Answer: A,C,D,E

Explanation:
To enable AI Observability for a Python application in Snowflake, several prerequisites must be met. The Python environment needs he `TRULENS_OTEL_TRACING` environment variable set to `1`. Essential Python packages, specifically `trulens-core`, `trulens-connectors-snowflake`, and `trulens-providers-cortex`, must be installed with version 2.1.2 or later. The Snowflake account role used requires the SNOWFLAKE.CORTEX_USER` database role and the `AI_OBSERVABILITY_EVENTS_LOOKUP` application role. Additionally, the role needs specific privileges: `CREATE EXTERNAL AGENT` and `CREATE TASK` on the schema, and `EXECUTE TASK` globally. It is explicitly stated that he project cannot be run in a Snowflake Notebook for these features to work, making option D incorrect.

NEW QUESTION # 50
.....

We know making progress and getting the certificate of GES-C01 study materials will be a matter of course with the most professional experts in command of the newest and the most accurate knowledge in it. Our SnowPro® Specialty: Gen AI Certification Exam exam prep has taken up a large part of market. with decided quality to judge from customers' perspective, If you choose the right GES-C01 Practice Braindumps, it will be a wise decision. Our behavior has been strictly ethical and responsible to you, which is trust worthy.

Test GES-C01 Dump: https://www.troytecdumps.com/GES-C01-troytec-exam-dumps.html

A good choice can make one work twice the result with half the effort, and our GES-C01 study materials will be your right choice, TroytecDumps Test GES-C01 Dump exam preparation material is available in two different types, Snowflake GES-C01 Reliable Braindumps Those updates will be sent to you accordingly for one year freely, Snowflake GES-C01 Reliable Braindumps Answer: We provide 90 DAYS free updates.

See the full list and read about Pierre's process, The Next Big Thing—The Fund of Funds, A good choice can make one work twice the result with half the effort, and our GES-C01 Study Materials will be your right choice.

The Best GES-C01 Reliable Braindumps & Leading Offer in Qualification Exams & Correct Snowflake SnowPro® Specialty: Gen AI Certification Exam

TroytecDumps exam preparation material is available in two different GES-C01 Actual Test types, Those updates will be sent to you accordingly for one year freely, Answer: We provide 90 DAYS free updates.

Like a mini Snowflake Certification boot camp, you'll be prepared for what ever comes your way GES-C01 with the world's best Snowflake Certification practice test guaranteed to deliver you the Snowflake Certification certificate you have been struggling to obtain with Snowflake Certification dumps.

- Quiz 2025 Snowflake Updated GES-C01 Reliable Braindumps □ Open website " www.dumpsquestion.com " and search for ⇒ GES-C01 ⇐ for free download □GES-C01 Test Engine Version
- Reliable GES-C01 Reliable Braindumps – Marvelous Test Dump Provider for GES-C01: SnowPro® Specialty: Gen AI Certification Exam □ Go to website ✔ www.pdfvce.com □✔□ open and search for □ GES-C01 □ to download for free □GES-C01 Reliable Braindumps Files
- Snowflake - GES-C01 Perfect Reliable Braindumps □ Open website ➡ www.examcollectionpass.com □□□ and search

for □ GES-C01 □ for free download □Relevant GES-C01 Exam Dumps

- Free PDF Quiz Authoritative Snowflake - GES-C01 Reliable Braindumps □ Search for （ GES-C01 ） on { www.pdfvce.com } immediately to obtain a free download □GES-C01 Latest Test Bootcamp
- Actual Snowflake GES-C01 Exam Dumps – Pass Exam With Good Scores □ Download ☀ GES-C01 □☀□ for free by simply searching on { www.lead1pass.com } □Exam GES-C01 Fees
- GES-C01 Labs □ Exam GES-C01 Fees □ Valid GES-C01 Test Dumps ➽ Search for □ GES-C01 □ and easily obtain a free download on ▹ www.pdfvce.com ◃ □GES-C01 Labs
- Valid GES-C01 Test Dumps □ GES-C01 Test Engine Version □ Training GES-C01 Tools □ Immediately open 【 www.prep4away.com 】 and search for { GES-C01 } to obtain a free download □GES-C01 Exam Materials
- 2025 GES-C01 Reliable Braindumps 100% Pass | Professional GES-C01: SnowPro® Specialty: Gen AI Certification Exam 100% Pass □ ➡ www.pdfvce.com □□□ is best website to obtain " GES-C01 " for free download □Reliable GES-C01 Test Sample
- Snowflake - GES-C01 Perfect Reliable Braindumps □ Search for ➡ GES-C01 □□□ and obtain a free download on □ www.examcollectionpass.com □ □Valid GES-C01 Test Topics
- Snowflake - GES-C01 Perfect Reliable Braindumps □ Immediately open ➠ www.pdfvce.com □ and search for ➥ GES-C01 □ to obtain a free download □Exam GES-C01 Braindumps
- TOP GES-C01 Reliable Braindumps: SnowPro® Specialty: Gen AI Certification Exam - Trustable Snowflake Test GES-C01 Dump □ Search for ➥ GES-C01 □ on ➡ www.pass4leader.com □□□ immediately to obtain a free download □ □Valid Braindumps GES-C01 Ppt
- dougwar742.fireblogz.com, study.stcs.edu.np, www.stes.tyc.edu.tw, catarijohanna643.blogspot.com, training.yoodrive.com, mrhamed.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, 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, onlyfans-asia.com, iangree641.full-design.com, Disposable vapes

DOWNLOAD the newest TroytecDumps GES-C01 PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1_oeQSD6_uH9yISd_wSmv6jbo5tTIXpn_