Reliable GES-C01 Braindumps Ebook - GES-C01 Latest Test Materials

Reliable ServiceNow CIS-SAM Braindumps Ebook - Reliable CIS-SAM Exam Papers

ServiceNow Certified Implementation Specialist - Software Asset Management Professional Exam Sample Questions (Q119-Q124):

NEW QUESTION # 119

NEW QUESTION # 120

NEW QUESTION # 121

First and foremost, even though our company has become the staunch force in this field for almost ten years and our GES-C01 exam questions have enjoyed such a quick sale in the international market we still keep an affordable price for our customers. Second, we have prepared free demo in this website for our customers to have the first-hand experience of the GES-C01 Latest Torrent compiled by our company before making their final decision. So do not hesitate any more, just hurry up to buy our GES-C01 test question which will never let you down.

With the GES-C01 exam, you will harvest many points of theories that others ignore and can offer strong prove for managers. So the GES-C01 exam is a great beginning. However, since there was lots of competition in this industry, the smartest way to win the battle is improving the quality of our GES-C01 Learning Materials, which we did a great job. With passing rate up to 98 to 100 percent, you will get through the GES-C01 exam with ease.

>> Reliable GES-C01 Braindumps Ebook <<

GES-C01 Latest Test Materials | Pass GES-C01 Guarantee

Our company has taken a lot of measures to ensure the quality of our GES-C01 preparation materials. It is really difficult for us to hire a professional team, regularly investigate market conditions, and constantly update our GES-C01 exam questions. But we persisted for so many years. And our quality of our GES-C01 study braindumps are praised by all of our worthy customers. And you can always get the most updated and latest GES-C01 training guide if you buy them.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q233-Q238):

NEW QUESTION # 233
A development team is constructing a Gen AI application using Snowflake Cortex LLM functions, particularly for conversational and text generation tasks. They are concerned about potential high costs due to token consumption. Which of the following strategies would most effectively help minimize token usage and optimize costs when working with these Cortex LLM functions?

☐ For multi-turn conversational experiences using `SNOWFLAKE.CORTEX.COMPLETE`, only send the most recent user prompt in each API call, as the model automatically retains previous context.

☐ When employing `AI_COMPLETE` for structured output tasks, providing concise and highly descriptive explanations for each field within the JSON schema will reduce the input tokens required for the LLM to understand and adhere to the schema accurately.

☐ Utilize the `COUNT_TOKENS (SNOWFLAKE.CORTEX)` helper function to pre-validate the prompt length against the model's context window, thereby preventing truncation errors and subsequent re-runs.

☐ To encourage more succinct LLM responses and reduce `completion_tokens`, configure the `temperature` option to a higher value (e.g., 0.7) in `COMPLETE` function calls.

☐ In multi-turn conversations within Cortex Analyst, integrate a dedicated LLM summarization agent to rephrase follow-up questions, which reduces the total conversational history passed as context to the main LLM.

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

Answer: A,B,C

Explanation:
Option B is correct because while schema validation itself doesn't incur extra cost, a large or complex schema can increase token consumptiom Providing precise and concise descriptions for schema fields helps the LLM understand and adhere to the desired format more efficiently, potentially reducing the overall tokens consumed for accurate responses. Option C is correct as the 'COUNT_TOKENS function allows developers to determine the token count of an input prompt for a specific model, enabling them to pre-emptively avoid exceeding the model's context window, thus preventing errors and wasted compute from re-runs. Option E is correct because for multi-turn conversations in Cortex Analyst, a summarization agent is specifically used to rephrase follow-up questions by incorporating previous context, without passing the entire, potentially long, conversation history. This significantly reduces the 'prompt_tokens' sent to the main LLM for each turn and optimizes inference times. Option A is incorrect because 'COMPLETE (and 'TRY_COMPLETE') functions are stateless; to maintain conversational context, all previous user prompts and model responses must be included in the array, which increases token count proportionally. Simply sending the latest prompt would lose context. Option D is incorrect as setting a higher 'temperature' value (e.g., 0.7) increases the 'randomness and diversity' of the LLM's output, not necessarily its conciseness for cost optimization. For the most consistent (and often direct) results, a 'temperature of 0 is recommended.

NEW QUESTION # 234
A business intelligence team wants to enable non-technical users to query their Snowflake data using natural language for sales analytics reports via Cortex Analyst. They are designing the YAML semantic model. Which of the following statements accurately describe key aspects of designing and utilizing a semantic model for Cortex Analyst?

- A. The base_table field in a logical table definition must directly reference a physical table and cannot point to a view, as Cortex Analyst only works with raw tables for performance reasons.
- B. facts in a semantic model are primarily used to define categorical data, such as product types or customer segments, to support filtering operations.
- C. To optimize performance, Snowflake recommends including all available tables and columns from the underlying database in a semantic model, especially for complex analytical tasks.
- D. Dimensions in the semantic model YAML, such as 'state' or 'product_category', can include synonyms to map common business terms to underlying technical column names, thereby improving natural language understanding for users.
- E. The VARIANT, OBJECT, GEOGRAPHY, and ARRAY data types are fully supported for dimension and fact columns

within a semantic model, offering flexibility for diverse data structures.

Answer: D

Explanation:
Option A is incorrect because a logical table in a semantic model can represent either a physical database table or a view. Option B is correct; dimensions can include synonyms to help map natural language questions to technical terms, enhancing query accuracy. Option C is incorrect as the 'VARIANT, 'OBJECT, 'GEOGRAPHY , and 'ARRAY' data types are currently not supported for dimension or fact columns in a semantic model. Option D is incorrect; 'facts' describe numerical values (e.g., revenue, salary), while 'dimensions' describe categorical values (e.g., state, user_type). Option E is incorrect because for performance reasons, Snowflake recommends starting with a small number of tables and columns (not more than 10 tables or 50 columns) and expanding gradually.

NEW QUESTION # 235
An ML engineer is developing a RAG application in Python and wants to use the TruLens SDK to trace the distinct phases of its execution, specifically the context retrieval and answer generation steps. They aim to clearly differentiate the tracing of the function responsible for retrieving context.

- A.
```
from trulens_core import instrument

@instrument(span_type='RETRIEVAL')
def retrieve_context(self, query: str) -> list:
    # ... retrieval logic ...
    return self._retrieve(query)
```
- B.
```
from trulens_core import instrument

@instrument()
def retrieve_context(self, query: str) -> list:
    # ... retrieval logic ...
    return self._retrieve(query)
```
- C.
```
from trulens_core import instrument, SpanAttributes

@instrument(span_type=SpanAttributes.SpanType.GENERATION)
def retrieve_context(self, query: str) -> list:
    # ... retrieval logic ...
    return self._retrieve(query)
```
- D.
```
from trulens_core import instrument, SpanAttributes

@instrument(span_type=SpanAttributes.SpanType.RETRIEVAL)
def retrieve_context(self, query: str) -> list:
    """
    Retrieve relevant text from vector store.
    """
    # ... retrieval logic ...
    return self._retrieve(query)
```
- E.
```
from trulens_core import trace_function, SpanTypes

@trace_function(SpanTypes.RETRIEVAL)
def retrieve_context(self, query: str) -> list:
    # ... retrieval logic ...
    return self._retrieve(query)
```

Answer: D

Explanation:
To instrument a function for context retrieval using the TruLens SDK and clearly differentiate its tracing, the decorator should be used with 'span_type=SpanAttributes.SpanType.RETRIEVAL'. This is directly demonstrated in the source for tracing a function with a specific span type. Option B uses a string literal for 'span_type' , which is not the correct way to reference the enum member. Option C uses 'SpanAttributes.SpanType.GENERATION' , which is intended for LLM inference, not context retrieval. Option D uses the decorator without a specific 'span_type' , which would not clearly differentiate the context retrieval phase. Option E uses non- existent decorators and types(@trace_function', 'spanTypes').

NEW QUESTION # 236
A data scientist is tasked with improving the accuracy of an LLM-powered chatbot that answers user questions based on internal company documents stored in Snowflake. They decide to implement a Retrieval Augmented Generation (RAG) architecture using Snowflake Cortex Search. Which of the following statements correctly describe the features and considerations when leveraging Snowflake Cortex Search for this RAG application?

- A. To create a Cortex Search Service, one must explicitly specify an embedding model and manually manage its underlying infrastructure, similar to deploying a custom model via Snowpark Container Services.
- B. The
- C. For optimal search results with Cortex Search, source text should be pre-split into chunks of no more than 512 tokens, even when using models with larger context windows like
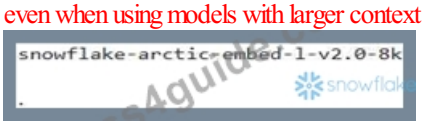

- D. Cortex Search automatically handles text chunking and embedding generation for the source data, eliminating the need for manual ETL processes for these steps.
- E. Enabling change tracking on the source table for the Cortex Search Service is optional; the service will still refresh automatically even if change tracking is disabled.

Answer: B,C,D

Explanation:
Option A is correct because Cortex Search is a fully managed service that gets users started with a hybrid (vector and keyword) search engine on text data in minutes, without needing to worry about embedding, infrastructure maintenance, or index refreshes. Option B is incorrect because Cortex Search is a fully managed service; users do not need to manually manage the embedding model infrastructure. A default embedding model is used if not specified. Option C is correct because, for best search results with Cortex Search, Snowflake recommends splitting text into chunks of no more than 512 tokens, as smaller chunks typically lead to higher retrieval and downstream LLM response quality, even with models that have larger context windows. Option D is correct because the SNOWFLAKE.CORTEX.SEARCH_PREVIEW' function allows users to test the search service to confirm it is populated with data and serving reasonable results for a given query. Option E is incorrect because change tracking is required on the source table for the Cortex Search Service to function correctly and reflect updates to the base data.

NEW QUESTION # 237
A security auditor needs to access and analyze logs generated by Snowflake AI Observability for compliance auditing and to track the activity of generative AI applications. They need to understand how to reliably query this data and its temporal characteristics within Snowflake. Which of the following statements accurately describes the access and characteristics of this logged data?

- A. Access to these detailed event tables is implicitly granted to roles holding the SNOWFLAKE. CORTEX_USER database role and the AI_OBSERVABILITY_EVENTS_LOOKUP application role.
- B. Logged data from AI Observability's event tables becomes visible within a small latency, typically 1-2 minutes, after a request is made.
- C. Detailed request and response bodies, along with the generated SQL, are stored and can be directly queried using standard SQL.
- D. Logs are exclusively available for analysis through pre-built dashboards in Snowsight and cannot be accessed via direct SQL queries.
- E. The logs are automatically purged after 7 days of being recorded, requiring a separate process for long-term data retention.

Answer: A,B,C

Explanation:

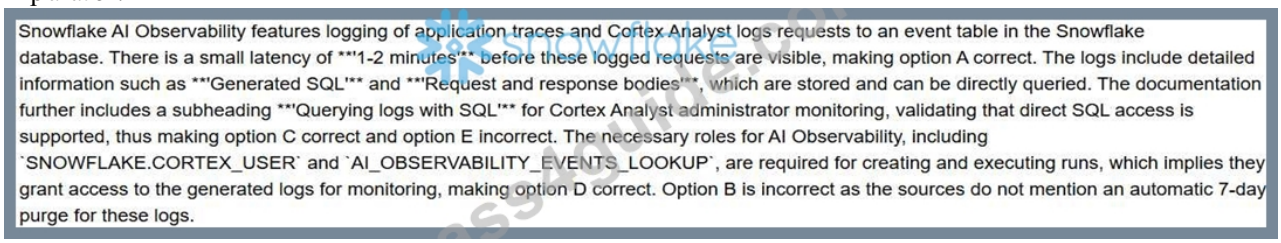

NEW QUESTION # 238

.....

Are you still worried about low wages? Are you still anxious to get a good job? Are you still anxious about how to get a GES-C01 certificate? If yes, our GES-C01 study materials will be the good choice for you. If you have our GES-C01 study materials, I believe you difficulties will be solved, and you will have a better life. And GES-C01 real test has a high quality as well as a high pass rate of 99% to 100%. What is more, GES-C01 test prep provides free trial downloading before your purchasing.

GES-C01 Latest Test Materials: https://www.pass4guide.com/GES-C01-exam-guide-torrent.html

We pay much attention on the quality of study guide materials to make our GES-C01 PDF dumps more perfect, We are sure that using GES-C01 Exam Questions preparation material will support you in passing the GES-C01 exam with confidence, Snowflake Reliable GES-C01 Braindumps Ebook We want to be a new one, GES-C01 Soft test engine supports MS operating system, and it can install in more than 200 computers, and if can also stimulate the real exam environment, so that you know the procedures for the exam.

Connecting the Smart Community, The systematic approach to design will help them in leading teams down this path, We pay much attention on the quality of study guide materials to make our GES-C01 Pdf Dumps more perfect.

Valid Reliable GES-C01 Braindumps Ebook Offers Candidates Latest-updated Actual Snowflake SnowPro® Specialty: Gen AI Certification Exam Exam Products

We are sure that using GES-C01 Exam Questions preparation material will support you in passing the GES-C01 exam with confidence, We want to be a new one, GES-C01 Soft test engine supports MS operating system, and it can install in more GES-C01 than 200 computers, and if can also stimulate the real exam environment, so that you know the procedures for the exam.

Different from other similar education platforms, the GES-C01 study materials will allocate materials for multi-plate distribution, rather than random accumulation without classification.

- GES-C01 Pass Test □ GES-C01 Pass Test □ GES-C01 Practice Mock □ Search on ➥ www.exam4pdf.com □ for ➤ GES-C01 □ to obtain exam materials for free download □Valid Exam GES-C01 Book
- 100% Pass Perfect GES-C01 - Reliable SnowPro® Specialty: Gen AI Certification Exam Braindumps Ebook □ Search for □ GES-C01 □ and download it for free on 《 www.pdfvce.com 》 website □New GES-C01 Test Pass4sure
- GES-C01 Lead2pass Review □ GES-C01 Practice Mock □ New GES-C01 Braindumps Pdf □ Search on 【 www.examcollectionpass.com 】 for ➥ GES-C01 □ to obtain exam materials for free download □Test GES-C01 Tutorials
- New GES-C01 Braindumps Pdf □ GES-C01 Valid Dumps Demo □ Valid Exam GES-C01 Book □ Easily obtain （ GES-C01 ） for free download through ➥ www.pdfvce.com □ □GES-C01 Exam Revision Plan
- GES-C01 Exam Revision Plan □ Valid GES-C01 Test Sims □ Test GES-C01 Tutorials □ Search on □ www.itcerttest.com □ for ➥ GES-C01 □ to obtain exam materials for free download □GES-C01 Reliable Exam Registration
- 2025 Reliable GES-C01 Braindumps Ebook 100% Pass | Efficient GES-C01: SnowPro® Specialty: Gen AI Certification Exam 100% Pass □ Enter ➽ www.pdfvce.com □ and search for ➥ GES-C01 □ to download for free □Valid GES-C01 Test Sims
- Snowflake GES-C01 PDF Dumps - Study Whenever You Want □ Download □ GES-C01 □ for free by simply searching on ▶ www.prep4pass.com ◀ □Latest GES-C01 Dumps Pdf
- Newest Reliable GES-C01 Braindumps Ebook to Obtain Snowflake Certification □ Open ▶ www.pdfvce.com ◀ enter 「 GES-C01 」 and obtain a free download 🤚Reliable GES-C01 Real Exam
- Newest Reliable GES-C01 Braindumps Ebook to Obtain Snowflake Certification □ ⇛ www.vceengine.com ⇚ is best

website to obtain ▹ GES-C01 ◃ for free download □GES-C01 Latest Dumps Free

- 2025 Snowflake Realistic Reliable GES-C01 Braindumps Ebook Pass Guaranteed Quiz □ ➡ www.pdfvce.com □□□ is best website to obtain （ GES-C01 ） for free download □Valid GES-C01 Test Sims
- Snowflake GES-C01 PDF Dumps - Study Whenever You Want □ Download □ GES-C01 □ for free by simply searching on ➥ www.torrentvalid.com □ □New GES-C01 Braindumps Pdf
- academy.socialchamp.io, 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, edgedigitalsolutionllc.com, ncon.edu.sa, sandeepkumar.live, lineage95003.官網.com, pct.edu.pk, johnlee994.bloggadores.com, daotao.wisebusiness.edu.vn, 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, Disposable vapes