

Newest Snowflake GES-C01 Practice Questions in PDF Format for Quick Preparation

Snowflake ADA-C01 Practice Questions

SnowPro Advanced Administrator

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



ADA-C01 Practice Exam Features | QuestionsTube

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

<https://www.questionstube.com/exam/ada-c01/>

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

1. SHOW WAREHOUSES;

No matter how the surrounding environment changes, you can easily deal with it with our GES-C01 exam questions. Do you want to be abandoned by others or have the right to pick someone else? Our GES-C01 simulating exam make you more outstanding and become the owner of your own life! Maybe you need to know more about our GES-C01 training prep to make a decision. Then you can free download the demos of our GES-C01 study guide, and you can have a experience on them before you pay for them.

Nowadays the knowledge capabilities and mental labor are more valuable than the manual labor because knowledge can create more wealth than the manual labor. If you boost professional knowledge capabilities in some area you are bound to create a lot of values and can get a good job with high income. Passing the test of GES-C01 Certification can help you achieve that, and our GES-C01 study materials are the best study materials for you to prepare for the test.

>>> **GES-C01 Reliable Real Exam** <<<

Latest GES-C01 Test Questions - Real GES-C01 Exam

With the Software version of our GES-C01 exam questions, you will find that there are no limits for the amount of the computers when download and installation and the users. You can use our GES-C01 study materials to stimulate the exam to adjust yourself to the atmosphere of the real exam and adjust your speed to answer the questions. The other two versions also boost the strength and applicable method and you could learn our GES-C01 training quiz by choosing the most suitable version to according to your practical situation.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q101-Q106):

NEW QUESTION # 101

An ML Engineer is logging a custom PyCaret model to the Snowflake Model Registry, with the intention of deploying it to Snowpark Container Services (SPCS) for GPU-powered inference. The PyCaret model is wrapped in a 'custom_model.ModelContext'. Which of the following statements correctly describe the considerations for the call and the model's environment?

- When specifying dependencies for the PyCaret model, it using `conda_dependencies`, packages will be sourced from the Snowflake conda channel, as this is the default for all model deployments.
- The `sample_input_data` or `signatures` argument must be provided during the `log_model` call to infer the model's input signature, which is crucial for deployment.
- To ensure the model is explicitly designed for SPCS and will not attempt to run in a warehouse, the `target_platforms` argument in `log_model` should be set to `'SNOWPARK_CONTAINER_SERVICES'`.
- Since PyCaret is a custom model type, it must be manually serialized (e.g., pickled) and referenced via a `ModelContext` object's attribute (e.g., `model_file`) before being passed to `log_model`.
- The `log_model` call will automatically detect and include any local files, such as configuration or parameter files, that the custom model relies on, without explicit specification.

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

Answer: B,C,E

Explanation:

Option B is correct because `sample_input_data` or `signatures` is a required argument for most model types (except Snowpark ML, MLFlow, and Hugging Face pipelines) to infer the model's input signature. This is crucial for successful deployment and inference. Option C is correct because explicitly setting `target_platforms=['SNOWPARK_CONTAINER_SERVICES']` in `log_model` ensures the model is intended for SPCS and prevents validation failures if it's ineligible for warehouse deployment. Option D is correct as custom model types like PyCaret, not natively supported by Snowflake ML, need to be manually serialized (e.g., pickled) and referenced within a `ModelContext` (e.g., `model_file='pycaret_best_model.pkl'`) before being passed to the `log_model` method. Option A is incorrect; while `conda_dependencies` is used, for models running on Snowpark Container Services (SPCS), packages are sourced from `conda-forge`, not the Snowflake conda channel (which is for warehouses). Option E is incorrect because local files required by the model, such as configuration or parameter files, must be explicitly specified using the `user_files` argument in the `log_model` call, as the system does not automatically detect them.

NEW QUESTION # 102

A data engineering team is planning to build a real-time data pipeline using Snowflake's dynamic tables to process incoming log data. They want to use SNOWFLAKE.CORTEX.EXTRACT_ANSWER to pull out specific error codes and timestamps from log entries. They are also mindful of the operational costs. Which of the following statements accurately describes limitations or cost considerations for using SNOWFLAKE.CORTEX.EXTRACT_ANSWER in this scenario?

- A. To optimize performance and reduce cost for large-scale log processing, it is recommended to execute EXTRACT_ANSWER queries on a Snowpark-optimized warehouse of at least 'LARGE' size.
- B. If a source_document input exceeds the model's token limit, EXTRACT_ANSWER will automatically truncate the text, potentially impacting results but avoiding an error.
- C. The billing for EXTRACT_ANSWER is based on the combined token count of both the source_document (log entry) and the question used for extraction.
- D. The EXTRACT_ANSWER function does not support dynamic tables, which will prevent its direct use in this pipeline design.
- E. EXTRACT_ANSWER is explicitly designed for multi-language extraction, making it suitable for logs that might contain various languages without affecting cost.

Answer: C,D

Explanation:

Option A is correct because Snowflake Cortex functions, including `EXTRACT_ANSWER`, do not support dynamic tables. Option B is correct because for `EXTRACT_ANSWER`, the number of billable tokens is the sum of the number of tokens in the (source document) and 'question' fields. Option C is incorrect; Snowflake recommends executing queries that call Cortex AISQL functions with a smaller warehouse (no larger than MEDIUM) as larger warehouses do not increase performance. Option D is incorrect; inputs that exceed the model's

token limit result in an error, rather than automatic truncation. Option E is incorrect; while the newer 'AI_EXTRACT' supports multiple languages, the documentation for 'EXTRACT_ANSWER' does not explicitly state multi-language support and generally refers to plain English text. Costs are incurred per token regardless of language.

NEW QUESTION # 103

An operations manager is tasked with monitoring the cost and ensuring compliance for a Cortex Analyst deployment that uses the REST API. They are particularly concerned with accurately tracking credit consumption and understanding the implications of enabling external models. Which of the following statements correctly describe aspects of Cortex Analyst cost and governance?

- A. The view can be queried to track detailed information about cortex Analyst requests, including generated SQL and any errors.
- B. Cortex Analyst credit usage is based on the number of messages processed, at a rate of 67 Credits per 1,000 messages, and only successful responses (HTTP 200) are counted.
- C. Credit consumption for Cortex Analyst is primarily based on the number of tokens processed by the underlying LLMs, with more complex natural language questions leading to higher token usage and costs.
- D. Enabling Azure OpenAI models via the 'ENABLE_CORTEX_ANALYST_MODEL_AZURE_OPENAI' parameter ensures that all customer data and prompts remain within Snowflake's governance boundary and fully respect RBAC policies for those models.
- E. Snowflake's view provides granular usage information for REST API requests to cortex Analyst, including tokens processed per model.

Answer: B

Explanation:

Option E is correct. Cortex Analyst credit rate usage is based on the number of messages processed, specifically 67 Credits per 1,000 messages, and only successful responses (HTTP 200) are counted. The number of tokens in each message does not affect cost. Option A is incorrect because, as stated, Cortex Analyst charges per message, not per token, and the number of tokens in each message does not affect the cost. This is distinct from general Cortex LLM functions which often bill per token. Option B is incorrect. The 'CORTEX_DOCUMENT_PROCESSING_USAGE_HISTORY' view displays Document AI processing function activity, not Cortex Analyst activity. To monitor Cortex Analyst requests, administrators can query the event table, which logs details like user, question, generated SQL, and errors. Option C is incorrect. Snowflake strongly discourages the use of Azure OpenAI models through the legacy 'ENABLE_CORTEX_ANALYST_MODEL_AZURE_OPENAI' parameter. If enabled, only metadata and prompts (not customer data) are transmitted outside of Snowflake's governance boundary, and model-level RBAC is not available. Option D is incorrect because the documentation explicitly states that granular usage information cannot be obtained for requests made with the REST API using the 'CORTEX_FUNCTIONS_QUERY_USAGE_HISTORY' view.

NEW QUESTION # 104

A 'document_processor_role' needs to be configured to fully manage Document AI in 'marketing_db.campaign_docs_schema', including creating model builds, uploading documents to an internal stage for processing, and setting up automated daily data extraction. The role has 'SNOWFLAKE.DOCUMENT_INTELLIGENCE_CREATOR' and 'USAGE' on 'marketing_db'. The compute warehouse 'marketing_wh' is active. Which of the following 'additional privileges' must be granted to 'document_processor_role' on the specified schema or account* to enable the full Document AI workflow described? (Select ALL that apply.)

- A. `GRANT OPERATE ON WAREHOUSE marketing_wh TO ROLE document_processor_role;`
- B. `GRANT CREATE TASK ON SCHEMA marketing_db.campaign_docs_schema TO ROLE document_processor_role;`
- C. `GRANT CREATE MODEL ON SCHEMA marketing_db.campaign_docs_schema TO ROLE document_processor_role;`
- D. `GRANT CREATE SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE ON SCHEMA marketing_db.campaign_docs_schema TO ROLE document_processor_role;`
- E. `GRANT CREATE STAGE ON SCHEMA marketing_db.campaign_docs_schema TO ROLE document_processor_role;`

Answer: A,B,C,D,E

Explanation:

For a role to have full access to Document AI, including preparing model builds and creating processing pipelines, several privileges are required in addition to 'SNOWFLAKE.DOCUMENT_INTELLIGENCE_CREATOR' and 'USAGE' on the database. These include: - 'OPERATE' on the virtual warehouse for Document AI operations. - 'CREATE SNOWFLAKE.ML.DOCUMENT_INTELLIGENCE' on the schema to create Document AI model builds. - 'CREATE MODEL' on the schema, as Document AI model builds include an underlying model. - 'CREATE STAGE' on the schema to upload and manage documents for extraction. - 'CREATE TASK' on the schema to establish automated document processing pipelines. All options A, B, C, D, and E are essential for the described full Document AI workflow.

NEW QUESTION # 105

A data engineering team is tasked with improving the accuracy of a Cortex Analyst solution for a large e-commerce product catalog. Users frequently ask natural language questions involving specific product names, brands, and categories. The team observes that Cortex Analyst sometimes struggles to identify and correctly filter by these literal values in the generated SQL. Which of the following configurations or approaches, within the semantic model, can effectively enhance Cortex Analyst's ability to precisely identify and use literal values for filtering, based on Snowflake's best practices?

- ☐ For a `product_category` dimension with less than 10 distinct values, setting `is_enum: true` and providing an exhaustive list of `sample_values` in the semantic model YAML.
- ☐ For a `product_name` dimension, configuring a `cortex_search_service` entry within the dimension, including both the `service` name and the `literal_column` that the Cortex Search Service is indexing.
- ☐ Using `LIKE` clauses directly within the dimension's `expr` field in the semantic model to enable fuzzy string matching for literal values.
- ☐ Relying exclusively on the `verified_queries` section of the semantic model to define all possible literal search scenarios with pre-written SQL.
- ☐ Increasing the `max_tokens` parameter in the Cortex Analyst API request options to allow the underlying LLM to generate more extensive literal value lists.

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

Answer: B,E

Explanation:

Options A and B are correct. For dimensions with low cardinality (around 1-10 distinct values), setting 'is_enum: true' and providing an exhaustive 'sample_values' list ensures Cortex Analyst chooses only from that predefined list, improving literal usage. For higher cardinality dimensions, integrating a Cortex Search Service via the entry, specifying both the 'service' name and the `literal_column`, allows semantic search over the underlying data to find appropriate literal values. Option C is incorrect because Cortex Analyst leverages semantic similarity search or Cortex Search for literal values, not direct 'LIKE clauses in the 'expr' field. Option D is incorrect because while 'verified_queries' improve accuracy for specific, known questions, they are not a scalable solution for all possible literal search scenarios and are not the primary mechanism for improving general literal value identification. Option E is incorrect because the 'max_tokens' parameter controls the length of the LLM's output response, not its ability to identify or filter by literal values.

NEW QUESTION # 106

.....

The Snowflake GES-C01 certification exam is one of the hottest certifications in the market. This Snowflake GES-C01 exam offers a great opportunity to learn new in-demand skills and upgrade your knowledge level. By doing this successful GES-C01 SnowPro® Specialty: Gen AI Certification Exam exam candidates can gain several personal and professional benefits.

Latest GES-C01 Test Questions: https://www.braindumpsvce.com/GES-C01_exam-dumps-torrent.html

Snowflake GES-C01 Reliable Real Exam When we grow up, we realize we should keep study and equip ourselves with more skills, thus we can stand out from the crowd, For the workers, an appropriate Snowflake GES-C01 exam certification can increase your competitiveness, and help you broaden your path of the future, Snowflake GES-C01 Reliable Real Exam You can purchase ahead and prepare more time.

Yet lawmakers should just leave anonymity alone, Doodle Diamonds looks like GES-C01 it's raining kites, When we grow up, we realize we should keep study and equip ourselves with more skills, thus we can stand out from the crowd.

Sharpen Your Time Management Skills with Snowflake GES-C01 Practice Test

For the workers, an appropriate Snowflake GES-C01 Exam Certification can increase your competitiveness, and help you broaden your path of the future. You can purchase ahead and prepare more time.

So you will never be disappointed once you choosing our GES-C01 latest dumps and you can absolutely get the desirable outcomes, As is known to all IT exams are difficult to pass but it is a great way to boost your career, especially for Snowflake GES-C01 exam.

- [illegible]