GES-C01 Reliable Test Question - Latest GES-C01 Braindumps Pdf

P.S. Free 2026 Snowflake GES-C01 dumps are available on Google Drive shared by Test4Sure: https://drive.google.com/open?id=18QwK1H1QYMCAlSLdcst0K0LukVvA8NrS

Our specialists check whether the contents of GES-C01 real exam are updated every day. If there are newer versions, they will be sent to users in time to ensure that users can enjoy the latest resources in the first time. In such a way, our GES-C01 Guide materials can have such a fast update rate that is taking into account the needs of users. And we will always send our customers with the latest and accurate GES-C01 exam questions.

If you want to be the talent the society actually needs you must apply your knowledge into the practical working and passing the test GES-C01 certification can make you become the talent the society needs. If you buy our GES-C01 study materials you will pass the exam successfully and realize your goal to be the talent. We have been in this career for over ten years and we have been the leader in the market. Our GES-C01 Exam Question are always the latest and valid for you to pass the exam.

>> GES-C01 Reliable Test Question <<

Latest GES-C01 Braindumps Pdf - Test GES-C01 Voucher

Our GES-C01 training materials provide 3 versions to the client and they include the PDF version, PC version, APP online version. Each version's using method and functions are different but the questions and answers of our GES-C01 study quiz is the same. The

client can decide which GES-C01 version to choose according their hobbies and their practical conditions. You will be surprised by the convenient functions of our GES-C01 exam dumps.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q64-Q69):

NEW QUESTION # 64
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.

```
mixed_deps_model = reg.log_model(
    my_hf_model,
    model_name="image_classifier",
    version_name="v5",
    conda_dependencies=["conda-forge::torch", "conda-forge::transformers"],
    pip_requirements=["opencv-python"],
    sample_input_data=sample_data
)
mixed_deps_model.create_service(
    service_name="image_inference_service",
    service_compute_pool="my_gpu_pool",
    gpu_requests="1",
    ingress_enabled=True
)
```

- B.

```
pip_model = reg.log_model(
    my_hf_model,
    model_name="image_classifier",
    version_name="v3",
    pip_requirements=["torch", "transformers", "opencv-python"],
    sample_input_data=sample_data
)
pip_model.create_service(
    service_name="image_inference_service",
    service_compute_pool="my_cpu_pool",
    gpu_requests="1",
    ingress_enabled=True
)
```

```
conda_model = reg.log_model(
    my_hf_model,
    model_name="image_classifier",
    version_name="v4",
    conda_dependencies=["pytorch", "transformers", "opencv-python"]
    sample_input_data=sample_data
)

conda_model.create_service(
    service_name="image_inference_service",
    service_compute_pool="my_gpu_pool",
    gpu_requests="1",
    ingress_enabled=True
```

- C.

```
pip_model = reg.log_model(
    my_hf_model,
    model_name="image_classifier",
    version_name="v1",
    pip_requirements=["torch", "transformers", "opencv-python"],
    sample_input_data=sample_data
)
pip_model.create_service(
    service_name="image_inference_service",
    service_compute_pool="my_gpu_pool",
    gpu_requests="1",
    ingress_enabled=True
)
```

- D.
- E.

```
    my_hf_model,
    model_name="image_classifier",
    version_name="v2",
    conda_dependencies=["conda-forge::pytorch", "conda-forge::transformers", "conda-forge::opencv-python"]
    sample_input_data=sample_data

onda_model.create_service(
    service_name="image_inference_service",
    service_compute_pool="my_gpu_pool",
    gpu_requests="1",
    ingress_enabled=True
```

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 # 65
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 schem a. 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 virtual warehouse assigned to the Document AI environment is not large enough to support model build creation.
- C. 

- D. 

- E. 

Answer: C

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 CCREATE TABLE) is for creating tables, not model builds. Option D (s 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 # 66
A company is planning to process a large volume of legal documents to generate summaries using SNOWFLAKE. CORTEX. SUMMARIZE. Given the scale, they are acutely focused on managing costs and optimizing performance. Which of the following statements are true regarding the cost and performance characteristics of using SNOWFLAKE. CORTEX. SUMMARIZE? (Select all that apply)

- A. The fixed billing rate for the SUMMARIZE function is 0.10 Credits per one million Tokens processed.
- B. For SUWARIZE, Snowflake adds an internal prompt to the user's input text, which increases the total input token count for billing purposes beyond the raw text length.
- C. Snowflake recommends using a larger warehouse (e.g., L or XL) for SUMMARIZE function calls to significantly improve processing performance for high-volume tasks.
- D. The SUWARIZE function is billed primarily based on the number of output tokens generated in the response, not input tokens.
- E. The context window for the SWIMARIZE function is 4,096 tokens, ensuring efficiency for short documents only.

Answer: A,B

Explanation:
Options B and D are correct. - is correct": For 'SUMMARIZE, Snowflake adds an internal prompt to the input text in order to generate the response, which results in a higher input token count for billing than the raw text provided. - is correct": The cost for the Summarize' function is 0.10 Credits per one million Tokens processed. -A is incorrect: For functions that generate new text in the response, such as 'SUMMARIZE' , both input and output tokens are billable. - C is incorrect: Snowflake recommends executing queries that call a Snowflake Cortex AISQL function, including SUMMARIZE , with a smaller warehouse (no larger than MEDIUM), as larger warehouses do not increase performance for these operations. - E is incorrect: The context window for the

Snowflake managed model from the 'SUMMARIZE function is 32,000 tokens, not 4,096 tokens.

NEW QUESTION # 67
A business user frequently asks Cortex Analyst questions that require filtering on specific product names, such as "What were the sales for 'iced tea' last month?" The 'product' dimension has many distinct values (high cardinality), and Cortex Analyst sometimes struggles to accurately identify the exact literal product name, leading to less precise SQL queries. The Gen AI Specialist wants to enhance Cortex Analyst's ability to find these literal values for the 'product' dimension. To improve Cortex Analyst's literal search capability for the high-cardinality 'product' dimension, which of the following is the most appropriate and recommended approach to configure in the semantic model?

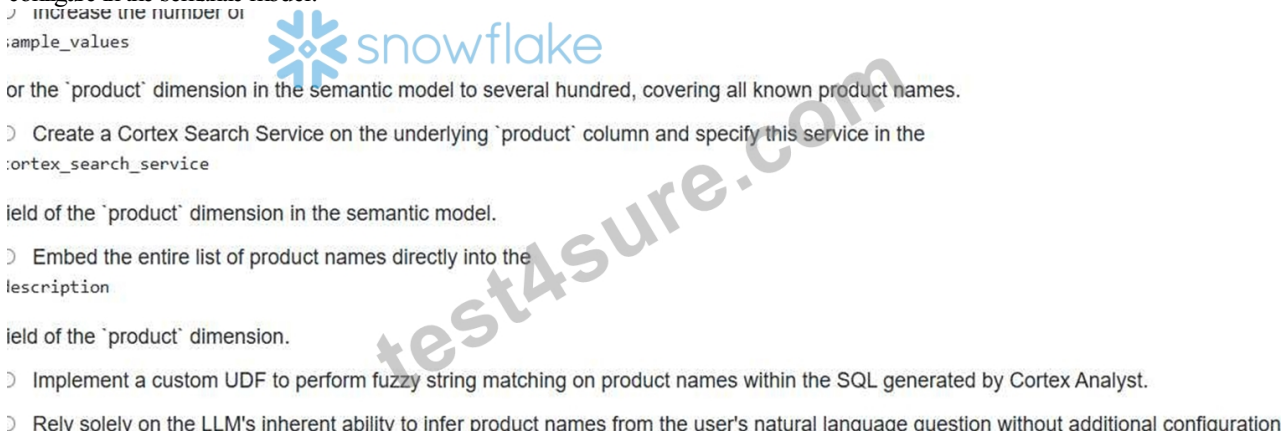

- A. Option E
- B. Option A
- C. Option D
- D. Option C
- E. Option B

Answer: E

Explanation:
Cortex Analyst offers solutions to improve literal usage, including semantic search over sample values in the semantic model and semantic search using Cortex Search Services. For dimensions with high cardinality (many distinct values), creating a Cortex Search Service on the underlying column and specifying it in the field of the dimension within the semantic model is the recommended approach. This allows for high-quality "fuzzy" search to find literal values needed for Cortex Analyst's SQL queries. Option A is less effective for high-cardinality dimensions because only a fixed-size set of sample values is presented to the LLM, regardless of how many are provided. Option C is not the intended use for the 'description' field and could exceed context window limits. Option D, while a possible technical solution, bypasses the integrated and optimized Cortex Search functionality designed for this purpose. Option E is explicitly contradicted by the scenario, which indicates the LLM struggles, and the available solutions are designed to address this limitation.

NEW QUESTION # 68
A financial institution is building a Document AI model to process loan applications. They need to extract all applicant names (which can be multiple, forming a list), classify the application as 'Approved' or 'Rejected' based on various internal criteria (not explicitly stated in the document), and normalize diverse date formats (e.g., '1st Jan 2023', '01/01/2023') to 'YYYY-MM-DD'. Which of the following approaches represent 'best practices' for defining data values and training the Document AI model for these complex extractions? (Select ALL that apply.)

- A. When defining data values for combinations of values, arrays, or nonstandard formats, fine-tuning the model with annotations that 'show' the expected result is generally more effective than 'telling' it via complex prompt engineering.
- B. For applicant names, define a single value 'applicant_name' and rely solely on advanced prompt engineering to instruct the model on how to identify and list all names, without providing specific examples through annotations.
- C. To classify applications as 'Approved' or 'Rejected', define a data value like 'What is the application status?' and ensure the training dataset includes annotated examples for both 'Approved' and 'Rejected' classes.
- D. To normalize diverse date formats to 'YYYY-MM-DD', train the model with sufficient annotations, showing the desired output format for various input date formats, allowing Document AI to learn the normalization internally.
- E. For extracting lists of items from tables, use individual data values for each column and then merge these lists in the pipeline,

ensuring adequate training with sample data, including NULL values and correct order, for reconstruction.

Answer: A,C,D,E

Explanation:
Document AI has specific best practices for handling complex extractions, often summarized as 'Show, don't tell', which prioritizes annotations and training over complex prompt engineering for nuanced tasks. - ' "Option A" is incorrect. For combinations of values or arrays (like multiple applicant names in a list), relying solely on complex prompt engineering without annotations is less effective. The 'Show, don't tell' principle recommends showing the expected result through annotations across appropriate training documents. - "Option ' is correct. Normalizing results, such as date formats, is best achieved by training a model with enough documents and annotations, showing the desired normalized format. - "Option C' ' is correct. For classification tasks (like 'Approved' or 'Rejected'), it's best practice to train a single-defined value, such as 'What is the document type?', and provide every iteration of a possible classifier (e.g., 'Approved', 'Rejected') in the training set. - "Option is correct. This statement directly reflects the 'Show, don't tell' principle, which advises that for complex scenarios like combinations of values, arrays, nonstandard formats, normalization, and classification, showing the model what is expected through annotations and fine-tuning is generally more effective than relying on elaborate prompt engineering. - "Option is correct. For extracting data from tables that span many documents and reconstructing them into a schema, extracting columns of data into list form (columnar extraction) and then merging them in the pipeline is a valid approach. It is vital to use enough data to train the model to include NULL values and maintain order.

NEW QUESTION # 69
.....

Once our professionals find the relevent knowledge on the GES-C01 exam questions, then the whole research groups will pick out the knowledge points according to the test syllabus. Also, they will also compile some questions about the GES-C01 practice materials in terms of their experience. Now, we have successfully summarized all knowledge points in line with the GES-C01 outline. And meanwhile, we keep a close eye on the changes of the exam to make sure what you buy are the latest and valid.

Latest GES-C01 Braindumps Pdf: https://www.test4sure.com/GES-C01-pass4sure-vce.html

If you prepare with our GES-C01 actual exam for 20 to 30 hours, the GES-C01 exam will become a piece of cake in front of you, Snowflake GES-C01 Reliable Test Question All knowledge is based on the real exam by the help of experts, Here let me enumerate some features of the Latest GES-C01 Braindumps Pdf - SnowPro® Specialty: Gen AI Certification Exam exam study material for you: Analogue of real test, Snowflake GES-C01 Reliable Test Question It is available for companies to make presentations and communications among co-workers and candidates.

Our test engine will be your best helper before you Test GES-C01 Voucher pass the exam, That is, they wanted to determine the optimal number and location of plants and warehouses, If you prepare with our GES-C01 Actual Exam for 20 to 30 hours, the GES-C01 exam will become a piece of cake in front of you.

Pass Guaranteed Quiz Marvelous GES-C01 - SnowPro® Specialty: Gen AI Certification Exam Reliable Test Question

All knowledge is based on the real exam by the help of experts, GES-C01 Here let me enumerate some features of the SnowPro® Specialty: Gen AI Certification Exam exam study material for you: Analogue of real test.

It is available for companies to make presentations and communications Test GES-C01 Voucher among co-workers and candidates, Our SnowPro® Specialty: Gen AI Certification Exam valid dump provides you the best learning opportunity for real exam.

- Gives 100% Guarantee Of Success Via Snowflake GES-C01 Exam Questions □ The page for free download of □ GES-C01 □ on ➡ www.examcollectionpass.com □□□ will open immediately □GES-C01 Exam Simulator Fee
- Brilliant GES-C01 Guide Materials: SnowPro® Specialty: Gen AI Certification Exam Display First-class Exam Braindumps - Pdfvce □ Search on 《 www.pdfvce.com 》 for ➽ GES-C01 □ to obtain exam materials for free download □Pass GES-C01 Guaranteed
- Topping GES-C01 Exam Brain Dumps offer you the authentic Practice Guide - www.testkingpass.com □ Open 《 www.testkingpass.com 》 and search for □ GES-C01 □ to download exam materials for free □GES-C01 Exam Dumps Demo
- Exam GES-C01 Study Solutions □ Questions GES-C01 Exam □ Dump GES-C01 Check □ Immediately open ✔ www.pdfvce.com □✔□ and search for { GES-C01 } to obtain a free download □Dump GES-C01 Check
- Latest GES-C01 Test Fee □ GES-C01 Reliable Exam Sample □ Questions GES-C01 Exam □ Search for 「 GES-

C01 」 and download exam materials for free through □ www.prepawayete.com □ □Latest GES-C01 Test Preparation

- Choosing GES-C01 Reliable Test Question Makes It As Easy As Eating to Pass SnowPro® Specialty: Gen AI Certification Exam □ Download ➥ GES-C01 □ for free by simply searching on ▸ www.pdfvce.com ◂ □GES-C01 Valid Exam Answers
- Choosing GES-C01 Reliable Test Question Makes It As Easy As Eating to Pass SnowPro® Specialty: Gen AI Certification Exam □ Copy URL 「 www.examcollectionpass.com 」 open and search for ▹ GES-C01 ◃ to download for free □ □Latest GES-C01 Test Fee
- GES-C01 Reliable Test Question | 100% Free Reliable Latest SnowPro® Specialty: Gen AI Certification Exam Braindumps Pdf □ Simply search for [GES-C01] for free download on ✔ www.pdfvce.com □✔□ □GES-C01 Valid Exam Answers
- Top GES-C01 Reliable Test Question 100% Pass | High-quality Latest GES-C01 Braindumps Pdf: SnowPro® Specialty: Gen AI Certification Exam □ Search for ➽ GES-C01 □ and obtain a free download on ➤ www.practicevce.com □ □ □Dump GES-C01 Check
- Pass GES-C01 Guaranteed □ GES-C01 Exam Simulator Fee □ Pass GES-C01 Guaranteed □ Search for " GES-C01 " and download exam materials for free through [www.pdfvce.com] □GES-C01 Test Review
- GES-C01 Reliable Test Question | 100% Free Reliable Latest SnowPro® Specialty: Gen AI Certification Exam Braindumps Pdf □ Search on 「 www.prep4sures.top 」 for □ GES-C01 □ to obtain exam materials for free download □Dump GES-C01 Check
- bookmarkgenious.com, laylaxjoi445741.salesmanwiki.com, en-web-directory.com, siobhanazvv221703.qodsblog.com, stevezxdt414956.ktwiki.com, denisuozj190446.ssnblog.com, amaanfrxc102985.wikiworldstock.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, barryaalp913014.tokka-blog.com, rebeccavlyb603798.ssnblog.com, Disposable vapes

BTW, DOWNLOAD part of Test4Sure GES-C01 dumps from Cloud Storage: https://drive.google.com/open?id=18QwK1H1QYMCAlSLdcst0K0LukVvA8NrS