

GES-C01 Pdf Dumps | GES-C01 Reliable Braindumps Sheet



The Snowflake PDF Questions format designed by the ExamDiscuss will facilitate its consumers. Its portability helps you carry on with the study anywhere because it functions on all smart devices. You can also make notes or print out the SnowPro® Specialty: Gen AI Certification Exam (GES-C01) pdf questions. The simple, systematic, and user-friendly Interface of the SnowPro® Specialty: Gen AI Certification Exam (GES-C01) PDF dumps format will make your preparation convenient.

In fact, a number of qualifying exams and qualifications will improve your confidence and sense of accomplishment to some extent, so our GES-C01 learning materials can be your new target. When we get into the job, our GES-C01 learning materials may bring you a bright career prospect. Companies need employees who can create more value for the company, but your ability to work directly proves your value. Our GES-C01 Learning Materials can help you improve your ability to work in the shortest amount of time, thereby surpassing other colleagues in your company, for more promotion opportunities and space for development. Believe it or not that up to you, our GES-C01 learning material is powerful and useful, it can solve all your stress and difficulties in reviewing the GES-C01 exams.

>> [GES-C01 Pdf Dumps](#) <<

Snowflake GES-C01 Pdf Dumps Exam | GES-C01 Reliable Braindumps Sheet – 100% free

In order to better meet users' needs, our GES-C01 study materials have set up a complete set of service system, so that users can enjoy our professional one-stop service. We not only in the pre-sale for users provide free demo, when buy the user can choose in we provide in the three versions, at the same time, our GES-C01 Study Materials also provides 24-hour after-sales service, even if you are failing the exam, don't pass the exam, the user may also demand a full refund with purchase vouchers, make the best use of the test data, not for the user to increase the economic burden.

Snowflake SnowPro® Specialty: Gen AI Certification Exam Sample Questions (Q245-Q250):

NEW QUESTION # 245

A new Gen AI team member attempts to use Document AI to process a batch of 1,500 scanned image files (JPG) that are 70 MB each, stored in an internal stage that was created without specifying an encryption type. Their '!PREDICT' queries consistently fail with various errors. Which of the following are valid reasons for the '!PREDICT' queries to fail in this scenario?

- A. Processing 1,500 documents in one query exceeds the maximum limit for Document AI.
- B. JPG is an unsupported file format for Document AI.
- C. The individual JPG files exceed the maximum supported file size for Document AI.
- D. The team member's role lacks the database role, which is essential for using Document AI functions.
- E. The internal stage was not created with 'ENCRYPTION = (TYPE = 'SNOWFLAKE SSE'), which is a requirement for Document AI.

Answer: A,C,D,E

Explanation:

Option A is correct because internal stages used with Document AI must specify 'ENCRYPTION = (TYPE = 'SNOWFLAKE SSE')' when created. Option B is correct as the database role is required for the account role to use Document AI functions to extract information. Option C is correct because Document AI supports processing a maximum of 1,000 documents in one query, so 1,500 documents would exceed this limit. Option D is correct because documents processed by Document AI must be 50 MB or less in size, and the 70 MB files exceed this limit. Option E is incorrect because JPG is listed as a supported file format for Document AI.

NEW QUESTION # 246

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 configurations correctly specify the necessary Python dependencies and GPU utilization for this inference service, adhering to Snowflake's recommendations?

```
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  
)
```

- A.
- B.



```
tp_model = reg.log_model(\n    my_hf_model,\n    model_name="image_classifier",\n    version_name="v1",\n    pip_requirements=["torch", "transformers", "opencv-python"]\n    sample_input_data=sample_data
```

```
ip_model.create_service(\n    service_name="image_inference_service",\n    service_compute_pool="my_gpu_pool",\n    gpu_requests="1",\n    ingress_enabled=True
```

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

- C.  examdiscuss.com
- D.  examdiscuss.com

```
conda_model = reg.log_model(\n    my_hf_model,\n    model_name="image_classifier",\n    version_name="v2",\n    conda_dependencies=["conda-forge::pytorch", "conda-forge::transformers", "conda-forge::opencv-python"]\n    sample_input_data=sample_data\n)\n\nconda_model.create_service(\n    service_name="image_inference_service",\n    service_compute_pool="my_gpu_pool",\n    gpu_requests="1",\n    ingress_enabled=True\n)
```

```

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
)

```

- E.

Answer: B

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. 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'. Option C is incorrect because is chosen, which will not provide GPU acceleration required by the model. 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, Snowflake explicitly recommends 'using only 'conda_dependencies' or only 'pip_requirements', not both' for managing dependencies to avoid potential conflicts.

NEW QUESTION # 247

A data application developer is building a Streamlit chat application within Snowflake. This application uses a RAG pattern to answer user questions about a knowledge base, leveraging a Cortex Search Service for retrieval and an LLM for generating responses. The developer wants to ensure responses are relevant, concise, and structured. Which of the following practices are crucial when integrating Cortex Search with Snowflake Cortex LLM functions like AI_COMPLETE for this RAG chatbot?

- A. For performance and cost optimization, it is always recommended to query Cortex Search and the LLM function within a single
- B. Using the
- C. The
- D. The retrieved context from Cortex Search should be directly concatenated with the user's prompt as input to the
- E. To maintain conversational context in a multi-turn chat, the developer should pass all previous user prompts and model responses in the

Answer: B,E

Explanation:

Option A is incorrect. The user's query is typically embedded (e.g., using `EMBED_TEXT`) to perform a similarity search against the Cortex Search Service. The "retrieved documents" (context) are then passed to the 'AI_COMPLETE' function, not the embedding function itself. Option B is correct because to provide a stateful, conversational experience, all previous user prompts and model responses should be passed in the 'prompt_or_history' array to the 'COMPLETE' or 'AI_COMPLETE' function. Option C is incorrect. While concatenation is a method, for better accuracy and control, the retrieved context should be integrated into a well-engineered prompt, often using tags or specific instructions, rather than just raw concatenation, to guide the LLM's response. Option D is correct because 'AI_COMPLETE Structured Outputs' allows you to supply a JSON schema that completion responses must follow, reducing the need for post-processing and enabling seamless integration with systems requiring deterministic responses. Option E is incorrect. While keeping processing within Snowflake is good for data governance, complex RAG pipelines often involve multiple distinct steps (query embedding, search, retrieval, LLM completion) that may benefit from a staged approach rather than a single monolithic SQL statement. The optimal approach depends on the specific complexity and performance requirements, and a single 'STRUCTURE' for the "entire" RAG flow might not always be the most efficient or practical solution.

NEW QUESTION # 248

A data engineer is configuring a Document AI pipeline to process scanned PDF invoices stored in an internal stage named 'invoice_docs_stage'. After uploading the PDF files, they execute an extracting query using '!PREDICT'. The query consistently returns the error:

```
{ "processingErrors": [ "File extension does not match actual mime type. Mime-Type: application/octet-stream" ] }
```

Which of the following is the most likely cause of this error?

- A. The 'GET_PRESIGNED_URL' function used in the '!PREDICT' query has an expired URL.
- B. The internal stage was not created with 'SNOWFLAKE_SSE' encryption enabled.**
- C. The Document AI model build is attempting to process more than 1000 documents in a single query.
- D. The PDF documents exceed the maximum allowed file size of 50 MB.
- E. The documents contain non-English text, which is not fully supported by Document AI for optimal results.

Answer: B

Explanation:

The error message 'File extension does not match actual mime type. Mime-Type: application/octet-stream' is a specific error documented for DocumentAI when internal stages are not created with 'SNOWFLAKE_SSE' encryption. For internal stages, Document AI requires server-side encryption to be enabled. Options A, C, and D would typically result in different error messages or behaviors. Option E refers to language support, which might impact accuracy but is not the cause of a file format identification error.

NEW QUESTION # 249

A security-conscious data scientist in an Azure East US 2 (Virginia) account wants to fine-tune a mistral-7b model for a specific text summarization task and then deploy it for real-time inference using the Cortex REST API. The mistral-7b base model is natively available for fine-tuning in Azure East US 2 (Virginia). For subsequent inference using the fine-tuned model, they need to understand the regional and cross-region inference considerations. Which of the following statements are correct?

- The fine-tuning job for `mistral-7b` can proceed in Azure East US 2 (Virginia) as it is a supported region for the `FINETUNE ('CREATE')` function.
- To use the fine-tuned `mistral-7b` model for inference via the Cortex REST API in a region where the base `mistral-7b` model is not natively supported, the `CORTEX_ENABLED_CROSS_REGION` account parameter must be configured.
- User inputs and outputs for inference requests to the fine-tuned model via cross-region inference will be stored or cached within the remote processing region to optimize performance.
- The cost for inference on the fine-tuned `mistral-7b` model will be based on the number of tokens processed for both prompt and completion, with an additional charge for Cortex Guard tokens if enabled.
- Snowflake Copilot, another Gen AI feature, also supports cross-region inference and its availability can be controlled by the `CORTEX_ENABLED_CROSS_REGION` parameter.

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

Answer: A,B,D,E

Explanation:

Option A is correct. Azure East US 2 (Virginia) is explicitly listed as a supported region for the `FINETUNE ('CREATE')` function for base models like `mistral-7b`. Option B is correct because support for inference of fine-tuned models is available in regions that support the `COMPLETE` function for the base model, and the `CORTEX_ENABLED_CROSS_REGION` parameter enables inference requests to be processed in a different region from the default. The Cortex REST API is an interface for the `COMPLETE` function. Option C is incorrect; user inputs, service generated prompts, and outputs are not stored or cached during cross-region inference. Option D is correct because for functions that generate new text, such as inference with a fine-tuned model (which uses `AI_COMPLETE`), both input and output tokens are billable. If Cortex Guard is enabled, its usage incurs an additional cost for tokens. Option E is correct as Snowflake Copilot's availability in unsupported regions can be enabled by setting the `CORTEX_ENABLED_CROSS_REGION` parameter.

NEW QUESTION # 250

.....

We also provide you with customizable desktop Central Finance in SnowPro® Specialty: Gen AI Certification Exam (GES-C01) practice test software and web-based Snowflake GES-C01 practice exam. You can adjust timings and SnowPro® Specialty: Gen

AI Certification Exam (GES-C01) questions number of our GES-C01 practice exams according to your training needs. These Snowflake GES-C01 Practice Tests simulate the real GES-C01 exam pattern, track your progress, and help you overcome mistakes. Our GES-C01 desktop software is compatible with Windows.

GES-C01 Reliable Braindumps Sheet: <https://www.examdiscuss.com/Snowflake/exam/GES-C01/>

We have made all efforts to update our products in order to help you deal with any change, making you confidently take part in the GES-C01 exam, Did you do it, If you have your own ambition of realizing personal dreams, our GES-C01 free questions will help you realize them successfully, Our GES-C01 study materials take the clients' needs to pass the test smoothly into full consideration, Preparing with real GES-C01 exam questions is one of the finest strategies for cracking the exam in one go.

Moving the Insertion Point with Click and Type, Then reconnect GES-C01 external power and any other cables you may have had to disconnect to get at the drive) and reboot.

We have made all efforts to update our products in order to help you deal with any change, making you confidently take part in the GES-C01 Exam, Did you do it?

New GES-C01 Pdf Dumps 100% Pass | Valid GES-C01 Reliable Braindumps Sheet: SnowPro® Specialty: Gen AI Certification Exam

If you have your own ambition of realizing personal dreams, our GES-C01 free questions will help you realize them successfully, Our GES-C01 study materials take the clients' needs to pass the test smoothly into full consideration.

Preparing with real GES-C01 exam questions is one of the finest strategies for cracking the exam in one go.

- New GES-C01 Cram Materials GES-C01 Braindumps Downloads GES-C01 Dumps Torrent Search for ► GES-C01 ↳ on ➡ www.dumps4pdf.com immediately to obtain a free download GES-C01 Exam Certification
- GES-C01 Pdf Dumps - SnowPro® Specialty: Gen AI Certification Exam Realistic Reliable Braindumps Sheet Free PDF Quiz  Search on  www.pdfvce.com  for  GES-C01  to obtain exam materials for free download New GES-C01 Dumps Pdf
- Your Investment with www.torrentvalid.com Snowflake GES-C01 Exam Questions is Secured Search for 《 GES-C01 》 and download exam materials for free through www.torrentvalid.com Reliable GES-C01 Dumps Files
- Accurate GES-C01 Pdf Dumps | Easy To Study and Pass Exam at first attempt - Authoritative GES-C01: SnowPro® Specialty: Gen AI Certification Exam Search on { www.pdfvce.com } for 「 GES-C01 」 to obtain exam materials for free download Test GES-C01 Engine Version
- Free PDF Quiz GES-C01 - SnowPro® Specialty: Gen AI Certification Exam – Efficient Pdf Dumps Copy URL ✓ www.torrentvalid.com ✓ open and search for GES-C01 to download for free GES-C01 Online Exam
- Snowflake GES-C01 Exam Questions – Secret To Pass On First Attempt Download GES-C01 for free by simply entering www.pdfvce.com website GES-C01 Dumps Torrent
- Snowflake GES-C01 Exam Questions – Secret To Pass On First Attempt Copy URL 「 www.vceengine.com 」 open and search for > GES-C01 to download for free GES-C01 Dumps Torrent
- Valid GES-C01 Exam Tutorial Reliable GES-C01 Exam Sample Trustworthy GES-C01 Dumps Search for ► GES-C01 and download it for free on ✓ www.pdfvce.com ✓ website GES-C01 Exam Format
- Reliable GES-C01 Dumps Files Valid GES-C01 Exam Tutorial GES-C01 Testking Learning Materials Enter  www.examdiscuss.com  and search for GES-C01 to download for free New GES-C01 Test Duration
- Buy Pdfvce Snowflake GES-C01 Questions Now And Get Free Updates Copy URL ► www.pdfvce.com ↳ open and search for ➡ GES-C01 to download for free New GES-C01 Test Duration
- Snowflake GES-C01 Exam Questions – Secret To Pass On First Attempt Easily obtain ⇒ GES-C01 ⇌ for free download through ➡ www.prep4pass.com GES-C01 Braindumps Downloads
- www.pcsq28.com, shortcourses.russellcollege.edu.au, elearning.eauqardho.edu.so, www.stes.tyc.edu.tw, study.stcs.edu.np, bbs.starcg.net, 6.k1668.cn, www.hgglz.com, xm.wztc58.cn, study.stcs.edu.np, Disposable vapes