

Databricks-Generative-AI-Engineer-Associate Unlimited Exam Practice & Databricks-Generative-AI-Engineer-Associate Reliable Exam Online



BTW, DOWNLOAD part of GetValidTest Databricks-Generative-AI-Engineer-Associate dumps from Cloud Storage:
https://drive.google.com/open?id=1v5fzAxy6SIW9ka_b_2iAgyJSa9i99TFM

Here in this Desktop practice test software, the Databricks Certified Generative AI Engineer Associate (Databricks-Generative-AI-Engineer-Associate) practice questions given are very relevant to the actual Databricks Certified Generative AI Engineer Associate (Databricks-Generative-AI-Engineer-Associate) exam. It is compatible with Windows computers. GetValidTest provides its valued customers with customizable Databricks Certified Generative AI Engineer Associate (Databricks-Generative-AI-Engineer-Associate) practice exam sessions. The Databricks Certified Generative AI Engineer Associate (Databricks-Generative-AI-Engineer-Associate) practice test software also keeps track of the previous Databricks Databricks-Generative-AI-Engineer-Associate practice exam attempts.

Databricks Databricks-Generative-AI-Engineer-Associate Exam Syllabus Topics:

Topic	Details
Topic 1	• Data Preparation: Generative AI Engineers covers a chunking strategy for a given document structure and model constraints. The topic also focuses on filter extraneous content in source documents. Lastly, Generative AI Engineers also learn about extracting document content from provided source data and format.
Topic 2	• Assembling and Deploying Applications: In this topic, Generative AI Engineers get knowledge about coding a chain using a pyfunc mode, coding a simple chain using langchain, and coding a simple chain according to requirements. Additionally, the topic focuses on basic elements needed to create a RAG application. Lastly, the topic addresses sub-topics about registering the model to Unity Catalog using MLflow.
Topic 3	• Evaluation and Monitoring: This topic is all about selecting an LLM choice and key metrics. Moreover, Generative AI Engineers learn about evaluating model performance. Lastly, the topic includes sub-topics about inference logging and usage of Databricks features.
Topic 4	• Governance: Generative AI Engineers who take the exam get knowledge about masking techniques, guardrail techniques, and legal • licensing requirements in this topic.

Databricks Databricks-Generative-AI-Engineer-Associate Reliable Exam Online | Databricks-Generative-AI-Engineer-Associate Latest Test Preparation

Our Databricks-Generative-AI-Engineer-Associate training dumps are deemed as a highly genius invention so all exam candidates who choose our Databricks-Generative-AI-Engineer-Associate exam questions have analogous feeling that high quality our practice materials is different from other practice materials in the market. So our Databricks-Generative-AI-Engineer-Associate study braindumps are a valuable invest which cost only tens of dollars but will bring you permanent reward. So many our customers have benefited form our Databricks-Generative-AI-Engineer-Associate preparation quiz, so will you!

Databricks Certified Generative AI Engineer Associate Sample Questions (Q53-Q58):

NEW QUESTION # 53

A Generative AI Engineer is creating an LLM-powered application that will need access to up-to-date news articles and stock prices.

The design requires the use of stock prices which are stored in Delta tables and finding the latest relevant news articles by searching the internet.

How should the Generative AI Engineer architect their LLM system?

- A. Query the Delta table for volatile stock prices and use an LLM to generate a search query to investigate potential causes of the stock volatility.
- B. Use an LLM to summarize the latest news articles and lookup stock tickers from the summaries to find stock prices.
- C. Download and store news articles and stock price information in a vector store. Use a RAG architecture to retrieve and generate at runtime.
- **D. Create an agent with tools for SQL querying of Delta tables and web searching, provide retrieved values to an LLM for generation of response.**

Answer: D

Explanation:

To build an LLM-powered system that accesses up-to-date news articles and stock prices, the best approach is to create an agent that has access to specific tools (option D).

* **Agent with SQL and Web Search Capabilities:** By using an agent-based architecture, the LLM can interact with external tools. The agent can query Delta tables (for up-to-date stock prices) via SQL and perform web searches to retrieve the latest news articles. This modular approach ensures the system can access both structured (stock prices) and unstructured (news) data sources dynamically.

* **Why This Approach Works:**

* **SQL Queries for Stock Prices:** Delta tables store stock prices, which the agent can query directly for the latest data.

* **Web Search for News:** For news articles, the agent can generate search queries and retrieve the most relevant and recent articles, then pass them to the LLM for processing.

* **Why Other Options Are Less Suitable:**

* **A (Summarizing News for Stock Prices):** This convoluted approach would not ensure accuracy when retrieving stock prices, which are already structured and stored in Delta tables.

* **B (Stock Price Volatility Queries):** While this could retrieve relevant information, it doesn't address how to obtain the most up-to-date news articles.

* **C (Vector Store):** Storing news articles and stock prices in a vector store might not capture the real-time nature of stock data and news updates, as it relies on pre-existing data rather than dynamic querying.

Thus, using an agent with access to both SQL for querying stock prices and web search for retrieving news articles is the best approach for ensuring up-to-date and accurate responses.

NEW QUESTION # 54

A Generative AI Engineer is creating an LLM system that will retrieve news articles from the year 1918 and related to a user's query

and summarize them. The engineer has noticed that the summaries are generated well but often also include an explanation of how the summary was generated, which is undesirable.

Which change could the Generative AI Engineer perform to mitigate this issue?

- A. Tune the chunk size of news articles or experiment with different embedding models.
- B. Split the LLM output by newline characters to truncate away the summarization explanation.
- C. Revisit their document ingestion logic, ensuring that the news articles are being ingested properly.
- **D. Provide few shot examples of desired output format to the system and/or user prompt.**

Answer: D

Explanation:

To mitigate the issue of the LLM including explanations of how summaries are generated in its output, the best approach is to adjust the training or prompt structure. Here's why Option D is effective:

* Few-shot Learning: By providing specific examples of how the desired output should look (i.e., just the summary without explanation), the model learns the preferred format. This few-shot learning approach helps the model understand not only what content to generate but also how to format its responses.

* Prompt Engineering: Adjusting the user prompt to specify the desired output format clearly can guide the LLM to produce summaries without additional explanatory text. Effective prompt design is crucial in controlling the behavior of generative models.

Why Other Options Are Less Suitable:

* A: While technically feasible, splitting the output by newline and truncating could lead to loss of important content or create awkward breaks in the summary.

* B: Tuning chunk sizes or changing embedding models does not directly address the issue of the model's tendency to generate explanations along with summaries.

* C: Revisiting document ingestion logic ensures accurate source data but does not influence how the model formats its output.

By using few-shot examples and refining the prompt, the engineer directly influences the output format, making this approach the most targeted and effective solution.

NEW QUESTION # 55

After changing the response generating LLM in a RAG pipeline from GPT-4 to a model with a shorter context length that the company self-hosts, the Generative AI Engineer is getting the following error:

```
{ "error_code": "BAD_REQUEST", "message": "Bad request: rpc error: code = InvalidArgument desc = prompt token count (4595) cannot exceed 4096..." }
```

What TWO solutions should the Generative AI Engineer implement without changing the response generating model? (Choose two.)

- **A. Reduce the number of records retrieved from the vector database**
- B. Reduce the maximum output tokens of the new model
- C. Use a smaller embedding model to generate
- D. Retrain the response generating model using ALiBi
- **E. Decrease the chunk size of embedded documents**

Answer: A,E

Explanation:

* Problem Context: After switching to a model with a shorter context length, the error message indicating that the prompt token count has exceeded the limit suggests that the input to the model is too large.

* Explanation of Options:

* Option A: Use a smaller embedding model to generate- This wouldn't necessarily address the issue of prompt size exceeding the model's token limit.

* Option B: Reduce the maximum output tokens of the new model- This option affects the output length, not the size of the input being too large.

* Option C: Decrease the chunk size of embedded documents- This would help reduce the size of each document chunk fed into the model, ensuring that the input remains within the model's context length limitations.

* Option D: Reduce the number of records retrieved from the vector database- By retrieving fewer records, the total input size to the model can be managed more effectively, keeping it within the allowable token limits.

* Option E: Retrain the response generating model using ALiBi- Retraining the model is contrary to the stipulation not to change the response generating model.

Options C and E are the most effective solutions to manage the model's shorter context length without changing the model itself, by

adjusting the input size both in terms of individual document size and total documents retrieved.

NEW QUESTION # 56

Generative AI Engineer at an electronics company just deployed a RAG application for customers to ask questions about products that the company carries. However, they received feedback that the RAG response often returns information about an irrelevant product.

What can the engineer do to improve the relevance of the RAG's response?

- A. Use a different LLM to improve the generated response
- **B. Assess the quality of the retrieved context**
- C. Implement caching for frequently asked questions
- D. Use a different semantic similarity search algorithm

Answer: B

Explanation:

In a Retrieval-Augmented Generation (RAG) system, the key to providing relevant responses lies in the quality of the retrieved context. Here's why option A is the most appropriate solution:

* Context Relevance: The RAG model generates answers based on retrieved documents or context. If the retrieved information is about an irrelevant product, it suggests that the retrieval step is failing to select the right context. The Generative AI Engineer must first assess the quality of what is being retrieved and ensure it is pertinent to the query.

* Vector Search and Embedding Similarity: RAG typically uses vector search for retrieval, where embeddings of the query are matched against embeddings of product descriptions. Assessing the semantic similarity search process ensures that the closest matches are actually relevant to the query.

* Fine-tuning the Retrieval Process: By improving the retrieval quality, such as tuning the embeddings or adjusting the retrieval strategy, the system can return more accurate and relevant product information.

* Why Other Options Are Less Suitable:

* B (Caching FAQs): Caching can speed up responses for frequently asked questions but won't improve the relevance of the retrieved content for less frequent or new queries.

* C (Use a Different LLM): Changing the LLM only affects the generation step, not the retrieval process, which is the core issue here.

* D (Different Semantic Search Algorithm): This could help, but the first step is to evaluate the current retrieval context before replacing the search algorithm.

Therefore, improving and assessing the quality of the retrieved context (option A) is the first step to fixing the issue of irrelevant product information.

NEW QUESTION # 57

A Generative AI Engineer is developing a patient-facing healthcare-focused chatbot. If the patient's question is not a medical emergency, the chatbot should solicit more information from the patient to pass to the doctor's office and suggest a few relevant pre-approved medical articles for reading. If the patient's question is urgent, direct the patient to calling their local emergency services.

Given the following user input:

"I have been experiencing severe headaches and dizziness for the past two days." Which response is most appropriate for the chatbot to generate?

- **A. Please call your local emergency services.**
- B. Please provide your age, recent activities, and any other symptoms you have noticed along with your headaches and dizziness.
- C. Here are a few relevant articles for your browsing. Let me know if you have questions after reading them.
- D. Headaches can be tough. Hope you feel better soon!

Answer: A

Explanation:

* Problem Context: The task is to design responses for a healthcare-focused chatbot that appropriately addresses the urgency of a patient's symptoms.

* Explanation of Options:

* Option A: Suggesting articles might be suitable for less urgent inquiries but is inappropriate for symptoms that could indicate a serious condition.

- * Option B: Given the description of severe symptoms like headaches and dizziness, directing the patient to emergency services is prudent. This aligns with medical guidelines that recommend immediate professional attention for such severe symptoms.
- * Option C: Offering well-wishes does not address the potential seriousness of the symptoms and lacks appropriate action.
- * Option D: While gathering more information is part of a detailed assessment, the immediate need here suggests a more urgent response.

Given the potential severity of the described symptoms, Option B is the most appropriate, ensuring the chatbot directs patients to seek urgent care when needed, potentially saving lives.

NEW QUESTION # 58

.....

GetValidTest is intent on keeping up with the latest technologies and applying them to the exam questions and answers not only on the content but also on the displays. That is why our pass rate is high as 98% to 100%. The data are unique-particular in this career. With our Databricks-Generative-AI-Engineer-Associate study torrent, you can enjoy the leisure study experience as well as pass the Databricks-Generative-AI-Engineer-Associate Exam with success ensured. For the content of our Databricks-Generative-AI-Engineer-Associate preparation materials is simplified by our professional experts and the displays are designed effectually. Just try and enjoy it!

Databricks-Generative-AI-Engineer-Associate Reliable Exam Online: <https://www.getvalidtest.com/Databricks-Generative-AI-Engineer-Associate-exam.html>

- Databricks-Generative-AI-Engineer-Associate dumps torrent - Databricks-Generative-AI-Engineer-Associate exam VCE - Databricks-Generative-AI-Engineer-Associate VCE PDF ☐ Enter ➡ www.prep4pass.com ☐ and search for ➡ Databricks-Generative-AI-Engineer-Associate ☐ to download for free ☐ Databricks-Generative-AI-Engineer-Associate Original Questions
- Unparalleled Databricks Databricks-Generative-AI-Engineer-Associate Unlimited Exam Practice With Interactive Test Engine - The Best Databricks-Generative-AI-Engineer-Associate Reliable Exam Online ☐ Enter ➡ www.pdfvce.com ☐ and search for ➡ Databricks-Generative-AI-Engineer-Associate ☐ to download for free * Databricks-Generative-AI-Engineer-Associate Exam Objectives Pdf
- Databricks-Generative-AI-Engineer-Associate Exam Objectives Pdf ☐ Databricks-Generative-AI-Engineer-Associate Latest Real Exam ☐ Databricks-Generative-AI-Engineer-Associate Valid Exam Duration ☐ Open website ☐ www.testsimulate.com ☐ and search for ⇒ Databricks-Generative-AI-Engineer-Associate ⇐ for free download ☐ Reliable Databricks-Generative-AI-Engineer-Associate Test Tutorial
- Databricks-Generative-AI-Engineer-Associate dumps torrent - Databricks-Generative-AI-Engineer-Associate exam VCE - Databricks-Generative-AI-Engineer-Associate VCE PDF ☐ Search for (Databricks-Generative-AI-Engineer-Associate) and download exam materials for free through ➡ www.pdfvce.com ☐ ☐ Databricks-Generative-AI-Engineer-Associate Valid Exam Duration
- New Databricks-Generative-AI-Engineer-Associate Braindumps ☐ Instant Databricks-Generative-AI-Engineer-Associate Discount ☐ Practice Databricks-Generative-AI-Engineer-Associate Test Online ☐ The page for free download of [Databricks-Generative-AI-Engineer-Associate] on ☐ www.pass4test.com ☐ will open immediately ☐ Databricks-Generative-AI-Engineer-Associate Exam Questions Pdf
- Updated Databricks Databricks-Generative-AI-Engineer-Associate Exam Questions BUNDLE PACK ☐ Open website ▶ www.pdfvce.com ◀ and search for ➡ Databricks-Generative-AI-Engineer-Associate ☐☐☐ for free download ☐ ☐ Databricks-Generative-AI-Engineer-Associate Original Questions
- Databricks-Generative-AI-Engineer-Associate PDF VCE ☐ PDF Databricks-Generative-AI-Engineer-Associate Download ☐ Databricks-Generative-AI-Engineer-Associate Reliable Exam Questions ☐ Search for ➤ Databricks-Generative-AI-Engineer-Associate ☐ and easily obtain a free download on ▶ www.testsdumps.com ◀ ☐ Latest Databricks-Generative-AI-Engineer-Associate Real Test
- Databricks-Generative-AI-Engineer-Associate Valid Exam Duration ☐ Instant Databricks-Generative-AI-Engineer-Associate Discount ☐ Databricks-Generative-AI-Engineer-Associate Mock Test ☐ Search for ➤ Databricks-Generative-AI-Engineer-Associate ☐ and obtain a free download on ▶ www.pdfvce.com ◀ ☐ Well Databricks-Generative-AI-Engineer-Associate Prep
- Unparalleled Databricks Databricks-Generative-AI-Engineer-Associate Unlimited Exam Practice With Interactive Test Engine - The Best Databricks-Generative-AI-Engineer-Associate Reliable Exam Online ☐ Enter ▶ www.real4dumps.com ◀ and search for ➤ Databricks-Generative-AI-Engineer-Associate ☐ to download for free ☐ Databricks-Generative-AI-Engineer-Associate Valid Exam Duration
- New Databricks-Generative-AI-Engineer-Associate Braindumps ☐ Databricks-Generative-AI-Engineer-Associate Latest Real Exam ☐ Databricks-Generative-AI-Engineer-Associate Reliable Exam Questions ☐ [www.pdfvce.com] is best website to obtain ☼ Databricks-Generative-AI-Engineer-Associate ☐☼☐ for free download ☐ Databricks-Generative-AI-Engineer-Associate Exam Objectives Pdf

- Databricks-Generative-AI-Engineer-Associate Unlimited Exam Practice Exam Latest Release | Updated Databricks-Generative-AI-Engineer-Associate Reliable Exam Online ☐ Easily obtain free download of 《 Databricks-Generative-AI-Engineer-Associate 》 by searching on ☒ www.prep4away.com ☐ ☒ ☐ Instant Databricks-Generative-AI-Engineer-Associate Discount
- www.wcs.edu.eu, learn.educatingeverywhere.com, mainpostvfx.com, www.goodgua.com, mikemil988.activosblog.com, 123.59.83.120:8080, yxy99.top, mikemil988.blogdomago.com, adamree449.aboutyoublog.com, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, Disposable vapes

DOWNLOAD the newest GetValidTest Databricks-Generative-AI-Engineer-Associate PDF dumps from Cloud Storage for free:
https://drive.google.com/open?id=1v5fzAxy6SIW9ka_b_2iAgyJSa9i99TFM