


Efficient Databricks-Generative-AI-Engineer-Associate Reliable Test Online Supply you Fast-Download Certification Exam Cost for Databricks-Generative-AI- Engineer-Associate: Databricks Certified Generative AI Engineer Associate to Study casually



BTW, DOWNLOAD part of Real4dumps Databricks-Generative-AI-Engineer-Associate dumps from Cloud Storage:
https://drive.google.com/open?id=1Q9Xa5q0o-qApikQGJI-ldnVFwvXc_CY7

Our Databricks-Generative-AI-Engineer-Associate training materials are famous for instant access to download. You can receive your downloading link and password within ten minutes, so that you can start your learning as early as possible. In order to build up your confidence for Databricks-Generative-AI-Engineer-Associate exam materials, we are pass guarantee and money back guarantee, and if you fail to pass the exam, we will give you full refund. In addition, Databricks-Generative-AI-Engineer-Associate test materials cover most of knowledge points for the exam, therefore you can master the major points for the exam as well as improve your professional ability in the process of learning.

Whether for a student or an office worker, obtaining Databricks-Generative-AI-Engineer-Associate certificate can greatly enhance the individual's competitiveness in the future career. Try our Databricks-Generative-AI-Engineer-Associate study materials, which are revised by hundreds of experts according to the changes in the syllabus and the latest developments in theory and practice. Once you choose Databricks-Generative-AI-Engineer-Associate training dumps, passing the exam one time is no longer a dream.

>> **Databricks-Generative-AI-Engineer-Associate Reliable Test Online** <<

Databricks-Generative-AI-Engineer-Associate Certification Exam Cost | Databricks-Generative-AI-Engineer-Associate Official Cert Guide

If you buy the Databricks Databricks-Generative-AI-Engineer-Associate practice materials within one year you can enjoy free updates. Being the most competitive and advantageous company in the market, our Databricks Certified Generative AI Engineer Associate Databricks-Generative-AI-Engineer-Associate exam questions have help tens of millions of exam candidates, realized their dreams all these years. What you can harvest is not only certificate but of successful future from now on just like our former clients.

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

Topic	Details
Topic 1	<ul style="list-style-type: none">• 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 2	<ul style="list-style-type: none"> • 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 3	<ul style="list-style-type: none"> • 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.

Databricks Certified Generative AI Engineer Associate Sample Questions (Q34-Q39):

NEW QUESTION # 34

A Generative AI Engineer is designing a chatbot for a gaming company that aims to engage users on its platform while its users play online video games.

Which metric would help them increase user engagement and retention for their platform?

- A. Diversity of responses
- B. Lack of relevance
- C. Randomness
- D. Repetition of responses

Answer: A

Explanation:

In the context of designing a chatbot to engage users on a gaming platform, diversity of responses (option B) is a key metric to increase user engagement and retention. Here's why:

Diverse and Engaging Interactions:

A chatbot that provides varied and interesting responses will keep users engaged, especially in an interactive environment like a gaming platform. Gamers typically enjoy dynamic and evolving conversations, and diversity of responses helps prevent monotony, encouraging users to interact more frequently with the bot.

Increasing Retention:

By offering different types of responses to similar queries, the chatbot can create a sense of novelty and excitement, which enhances the user's experience and makes them more likely to return to the platform.

Why Other Options Are Less Effective:

A (Randomness): Random responses can be confusing or irrelevant, leading to frustration and reducing engagement.

C (Lack of Relevance): If responses are not relevant to the user's queries, this will degrade the user experience and lead to disengagement.

D (Repetition of Responses): Repetitive responses can quickly bore users, making the chatbot feel uninteresting and reducing the likelihood of continued interaction.

Thus, diversity of responses (option B) is the most effective way to keep users engaged and retain them on the platform.

NEW QUESTION # 35

A Generative AI Engineer is tasked with developing a RAG application that will help a small internal group of experts at their company answer specific questions, augmented by an internal knowledge base. They want the best possible quality in the answers, and neither latency nor throughput is a huge concern given that the user group is small and they're willing to wait for the best answer. The topics are sensitive in nature and the data is highly confidential and so, due to regulatory requirements, none of the information is allowed to be transmitted to third parties.

Which model meets all the Generative AI Engineer's needs in this situation?

- A. BGE-large
- B. Llama2-70B
- C. OpenAI GPT-4
- D. Dolly 1.5B

Answer: A

Explanation:

* Problem Context: The Generative AI Engineer needs a model for a Retrieval-Augmented Generation (RAG) application that provides high-quality answers, where latency and throughput are not major concerns. The key factors are confidentiality and sensitivity of the data, as well as the requirement for all processing to be confined to internal resources without external data transmission.

* Explanation of Options:

Option A: Dolly 1.5B: This model does not typically support RAG applications as it's more focused on image generation tasks.

Option B: OpenAI GPT-4: While GPT-4 is powerful for generating responses, its standard deployment involves cloud-based processing, which could violate the confidentiality requirements due to external data transmission.

Option C: BGE-large: The BGE (Big Green Engine) large model is a suitable choice if it is configured to operate on-premises or within a secure internal environment that meets regulatory requirements. Assuming this setup, BGE-large can provide high-quality answers while ensuring that data is not transmitted to third parties, thus aligning with the project's sensitivity and confidentiality needs.

Option D: Llama2-70B: Similar to GPT-4, unless specifically set up for on-premises use, it generally relies on cloud-based services, which might risk confidential data exposure.

Given the sensitivity and confidentiality concerns, BGE-large is assumed to be configurable for secure internal use, making it the optimal choice for this scenario.

NEW QUESTION # 36

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

Answer: B,C

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 D 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 # 37

All of the following are Python APIs used to query Databricks foundation models. When running in an interactive notebook, which of the following libraries does not automatically use the current session credentials?

- A. REST API via requests library
- B. OpenAI client

- C. Databricks Python SDK
- D. MLflow Deployments SDK

Answer: A

Explanation:

When working within a Databricks notebook, several high-level SDKs are "Databricks-aware." The MLflow Deployments SDK (C) and the Databricks Python SDK (D) are designed to automatically look for the DATABRICKS_HOST and DATABRICKS_TOKEN environment variables provided by the notebook context. The OpenAI client (A), when configured for Databricks via Mosaic AI Gateway, also typically handles authentication via workspace integration in recent versions. However, the REST API via the requests library (B) is a generic Python HTTP client. It has no intrinsic knowledge of the Databricks environment. To use it, an engineer must manually extract the token (e.g., via dbutils.notebook.entry_point...) and explicitly pass it in the Authorization: Bearer <token> header of the request. Without this manual step, the requests library will fail with a 401 Unauthorized error.

NEW QUESTION # 38

A Generative AI Engineer is deciding between using LSH (Locality Sensitive Hashing) and HNSW (Hierarchical Navigable Small World) for indexing their vector database. Their top priority is semantic accuracy. Which approach should the Generative AI Engineer use to evaluate these two techniques?

- A. Compare the Levenshtein distances of returned results against a representative sample of test inputs
- **B. Compare the cosine similarities of the embeddings of returned results against those of a representative sample of test inputs**
- C. Compare the Recall-Oriented-Understudy for Gisting Evaluation (ROUGE) scores of returned results for a representative sample of test inputs
- D. Compare the Bilingual Evaluation Understudy (BLEU) scores of returned results for a representative sample of test inputs

Answer: B

NEW QUESTION # 39

.....

Obtaining a certificate has many benefits, you can strengthen your competitive force in the job market, enter a better company, and double your wage etc. Databricks-Generative-AI-Engineer-Associate exam bootcamp of us will help you get the certificate successfully. With experienced experts to edit and verify, Databricks-Generative-AI-Engineer-Associate exam dumps are high quality and accuracy. You can pass the exam just one time. In addition, Databricks-Generative-AI-Engineer-Associate Exam Bootcamp contain both questions and answers, and you can check the answer easily. Free update for 365 days is available. Our system will send the latest version of Databricks-Generative-AI-Engineer-Associate exam dumps to you automatically.

Databricks-Generative-AI-Engineer-Associate Certification Exam Cost: https://www.real4dumps.com/Databricks-Generative-AI-Engineer-Associate_examcollection.html

- Databricks-Generative-AI-Engineer-Associate Valid Test Dumps Valid Databricks-Generative-AI-Engineer-Associate Exam Tips Databricks-Generative-AI-Engineer-Associate Mock Exam Go to website www.pdf4dumps.com open and search for \Rightarrow Databricks-Generative-AI-Engineer-Associate \Leftarrow to download for free Dump Databricks-Generative-AI-Engineer-Associate File
- Databricks-Generative-AI-Engineer-Associate Latest Dumps: Databricks Certified Generative AI Engineer Associate - Databricks Certified Generative AI Engineer Associate Exam Cram Search for Databricks-Generative-AI-Engineer-Associate and download it for free on " www.pdfvce.com " website Databricks-Generative-AI-Engineer-Associate Reliable Dumps Pdf
- Databricks-Generative-AI-Engineer-Associate Latest Dumps: Databricks Certified Generative AI Engineer Associate - Databricks Certified Generative AI Engineer Associate Exam Cram Open website \Rightarrow www.examcollectionpass.com and search for \triangleright Databricks-Generative-AI-Engineer-Associate \triangleleft for free download Test Databricks-Generative-AI-Engineer-Associate Pass4sure
- 100% Pass Quiz 2026 Databricks Databricks-Generative-AI-Engineer-Associate: High Hit-Rate Databricks Certified Generative AI Engineer Associate Reliable Test Online Open [www.pdfvce.com] enter Databricks-Generative-AI-Engineer-Associate and obtain a free download Valid Databricks-Generative-AI-Engineer-Associate Exam Tips
- Valid Databricks-Generative-AI-Engineer-Associate Test Syllabus Vce Databricks-Generative-AI-Engineer-Associate Format Dump Databricks-Generative-AI-Engineer-Associate File Simply search for Databricks-Generative-AI-Engineer-Associate for free download on www.testkingpass.com Pdf Databricks-Generative-AI-Engineer-

