Try Fast2test Updated Databricks Databricks-Generative-AI-Engineer-Associate Questions For Easy and Quick Preparation

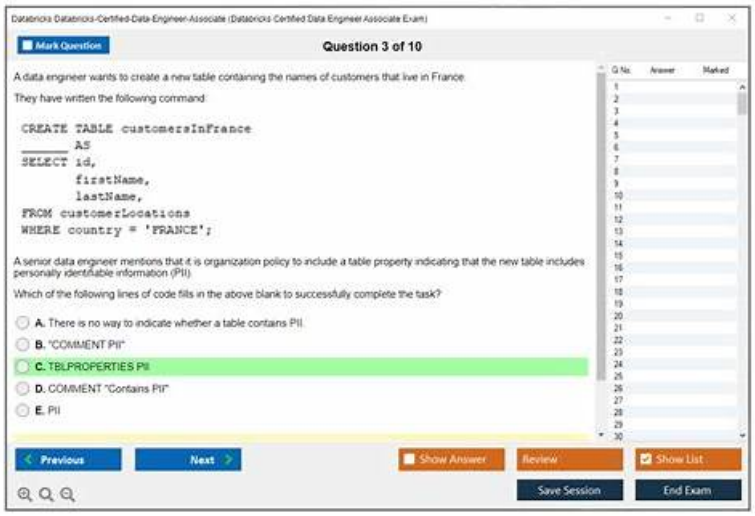

BTW, DOWNLOAD part of Fast2test Databricks-Generative-AI-Engineer-Associate dumps from Cloud Storage: https://drive.google.com/open?id=1FeeXxvFzpCcB6OXLqpW39Q_iEdOhye1D

Databricks-Generative-AI-Engineer-Associate exam torrent is famous for instant download. You will receive downloading link and password within ten minutes, and if you don't receive, just contact us, we will check for you. In addition, Databricks-Generative-AI-Engineer-Associate exam materials are high quality, it covers major knowledge points for the exam, you can have an easy study if you choose us. We offer you free demo to have a try before buying Databricks-Generative-AI-Engineer-Associate Exam Torrent, so that you can know what the complete version is like. Free update for one year is available, so that you can get the latest version for Databricks-Generative-AI-Engineer-Associate exam dumps timely.

Fast2test Databricks-Generative-AI-Engineer-Associate valid training material is the efforts of our professional experts. They edit and compile the Databricks-Generative-AI-Engineer-Associate questions and answers using their professional technology and hands-on experience. So if you want to pass with 100% guarantee, Databricks-Generative-AI-Engineer-Associate vlid exam files will give you security and high scores. You will complete your Databricks Databricks-Generative-AI-Engineer-Associate exam preparation in a short time and attend the actual test with comfortable mood.

>> Databricks-Generative-AI-Engineer-Associate Test Cram Review <<

Databricks Databricks-Generative-AI-Engineer-Associate Exam Simulator & Databricks-Generative-AI-Engineer-Associate Valid Test Guide

If you are worry about the coming Databricks-Generative-AI-Engineer-Associate study materials, our study materials will help you solve your problem. In order to promise the high quality of our Databricks-Generative-AI-Engineer-Associate study materials, our company has outstanding technical staff, and has perfect service system after sale. More importantly, our good Databricks-Generative-AI-Engineer-Associate guide questions and perfect after sale service are approbated by our local and international customers. If you want to pass your practice exam, we believe that our learning engine will be your indispensable choices. More and more people have bought our Databricks-Generative-AI-Engineer-Associate Guide questions in the past years.

Databricks Certified Generative AI Engineer Associate Sample Questions (Q17-Q22):

NEW QUESTION # 17
A small and cost-conscious startup in the cancer research field wants to build a RAG application using Foundation Model APIs. Which strategy would allow the startup to build a good-quality RAG application while being cost-conscious and able to cater to

customer needs?

- A. Use the largest LLM possible because that gives the best performance for any general queries
- B. Limit the number of queries a customer can send per day
- C. Pick a smaller LLM that is domain-specific
- D. Limit the number of relevant documents available for the RAG application to retrieve from

Answer: C

Explanation:
For a small, cost-conscious startup in the cancer research field, choosing a domain-specific and smaller LLM is the most effective strategy. Here's whyBis the best choice:
* Domain-specific performance: A smaller LLM that has been fine-tuned for the domain of cancer research will outperform a general-purpose LLM for specialized queries. This ensures high-quality responses without needing to rely on a large, expensive LLM.
* Cost-efficiency: Smaller models are cheaper to run, both in terms of compute resources and API usage costs. A domain-specific smaller LLM can deliver good quality responses without the need for the extensive computational power required by larger models.
* Focused knowledge: In a specialized field like cancer research, having an LLM tailored to the subject matter provides better relevance and accuracy for queries, while keeping costs low.Large, general- purpose LLMs may provide irrelevant information, leading to inefficiency and higher costs.
This approach allows the startup to balance quality, cost, and customer satisfaction effectively, making it the most suitable strategy.

NEW QUESTION # 18
A Generative Al Engineer is responsible for developing a chatbot to enable their company's internal HelpDesk Call Center team to more quickly find related tickets and provide resolution. While creating the GenAI application work breakdown tasks for this project, they realize they need to start planning which data sources (either Unity Catalog volume or Delta table) they could choose for this application. They have collected several candidate data sources for consideration:
call_rep_history: a Delta table with primary keys representative_id, call_id. This table is maintained to calculate representatives' call resolution from fields call_duration and call start_time.
transcript Volume: a Unity Catalog Volume of all recordings as a *.wav files, but also a text transcript as *.txt files.
call_cust_history: a Delta table with primary keys customer_id, cal1_id. This table is maintained to calculate how much internal customers use the HelpDesk to make sure that the charge back model is consistent with actual service use.
call_detail: a Delta table that includes a snapshot of all call details updated hourly. It includes root_cause and resolution fields, but those fields may be empty for calls that are still active.
maintenance_schedule - a Delta table that includes a listing of both HelpDesk application outages as well as planned upcoming maintenance downtimes.
They need sources that could add context to best identify ticket root cause and resolution.
Which TWO sources do that? (Choose two.)

- A. call_cust_history
- B. call_detail
- C. maintenance_schedule
- D. call_rep_history
- E. transcript Volume

Answer: B,E

Explanation:
In the context of developing a chatbot for a company's internal HelpDesk Call Center, the key is to select data sources that provide the most contextual and detailed information about the issues being addressed. This includes identifying the root cause and suggesting resolutions. The two most appropriate sources from the list are:
* Call Detail (Option D):
* Contents: This Delta table includes a snapshot of all call details updated hourly, featuring essential fields like root_cause and resolution.
* Relevance: The inclusion of root_cause and resolution fields makes this source particularly valuable, as it directly contains the information necessary to understand and resolve the issues discussed in the calls. Even if some records are incomplete, the data provided is crucial for a chatbot aimed at speeding up resolution identification.
* Transcript Volume (Option E):
* Contents: This Unity Catalog Volume contains recordings in .wav format and text transcripts in .txt files.
* Relevance: The text transcripts of call recordings can provide in-depth context that the chatbot can analyze to understand the

nuances of each issue. The chatbot can use natural language processing techniques to extract themes, identify problems, and suggest resolutions based on previous similar interactions documented in the transcripts.
Why Other Options Are Less Suitable:
* A (Call Cust History): While it provides insights into customer interactions with the HelpDesk, it focuses more on the usage metrics rather than the content of the calls or the issues discussed.
* B (Maintenance Schedule): This data is useful for understanding when services may not be available but does not contribute directly to resolving user issues or identifying root causes.
* C (Call Rep History): Though it offers data on call durations and start times, which could help in assessing performance, it lacks direct information on the issues being resolved.
Therefore, Call Detail and Transcript Volume are the most relevant data sources for a chatbot designed to assist with identifying and resolving issues in a HelpDesk Call Center setting, as they provide direct and contextual information related to customer issues.

NEW QUESTION # 19
A Generative AI Engineer I using the code below to test setting up a vector store:

```
from databricks.vector_search.client import VectorSearchClient

vsc = VectorSearchClient()

vsc.create_endpoint(
     name= "vector_search_test",
     endpoint_type= "STANDARD"
)
```


Assuming they intend to use Databricks managed embeddings with the default embedding model, what should be the next logical function call?

- A. vsc.create_delta_sync_index()
- B. vsc.get_index()
- C. vsc.create_direct_access_index()
- D. vsc.similarity_search()

Answer: A

Explanation:
Context: The Generative AI Engineer is setting up a vector store using Databricks' VectorSearchClient. This is typically done to enable fast and efficient retrieval of vectorized data for tasks like similarity searches.
Explanation of Options:
* Option A: vsc.get_index(): This function would be used to retrieve an existing index, not create one, so it would not be the logical next step immediately after creating an endpoint.
* Option B: vsc.create_delta_sync_index(): After setting up a vector store endpoint, creating an index is necessary to start populating and organizing the data. The create_delta_sync_index() function specifically creates an index that synchronizes with a Delta table, allowing automatic updates as the data changes. This is likely the most appropriate choice if the engineer plans to use dynamic data that is updated over time.
* Option C: vsc.create_direct_access_index(): This function would create an index that directly accesses the data without synchronization. While also a valid approach, it's less likely to be the next logical step if the default setup (typically accommodating changes) is intended.
* Option D: vsc.similarity_search(): This function would be used to perform searches on an existing index; however, an index needs to be created and populated with data before any search can be conducted.
Given the typical workflow in setting up a vector store, the next step after creating an endpoint is to establish an index, particularly one that synchronizes with ongoing data updates, henceOption B.

NEW QUESTION # 20
A Generative AI Engineer developed an LLM application using the provisioned throughput Foundation Model API. Now that the application is ready to be deployed, they realize their volume of requests are not sufficiently high enough to create their own provisioned throughput endpoint. They want to choose a strategy that ensures the best cost-effectiveness for their application.
What strategy should the Generative AI Engineer use?

- A. Deploy the model using pay-per-token throughput as it comes with cost guarantees
- B. Change to a model with a fewer number of parameters in order to reduce hardware constraint issues
- C. Switch to using External Models instead
- D. Throttle the incoming batch of requests manually to avoid rate limiting issues

Answer: A

Explanation:
* Problem Context: The engineer needs a cost-effective deployment strategy for an LLM application with relatively low request volume.
* Explanation of Options:
* Option A: Switching to external models may not provide the required control or integration necessary for specific application needs.
* Option B: Using a pay-per-token model is cost-effective, especially for applications with variable or low request volumes, as it aligns costs directly with usage.
* Option C: Changing to a model with fewer parameters could reduce costs, but might also impact the performance and capabilities of the application.
* Option D: Manually throttling requests is a less efficient and potentially error-prone strategy for managing costs.
OptionBis ideal, offering flexibility and cost control, aligning expenses directly with the application's usage patterns.

NEW QUESTION # 21
A Generative Al Engineer needs to design an LLM pipeline to conduct multi-stage reasoning that leverages external tools. To be effective at this, the LLM will need to plan and adapt actions while performing complex reasoning tasks.
Which approach will do this?

- A. Implement a framework like ReAct which allows the LLM to generate reasoning traces and perform task-specific actions that leverage external tools if necessary.
- B. Use a Chain-of-Thought (CoT) prompting technique to guide the LLM through a series of reasoning steps, then manually input the results from external tools for the final answer.
- C. Tram the LLM to generate a single, comprehensive response without interacting with any external tools, relying solely on its pre-trained knowledge.
- D. Encourage the LLM to make multiple API calls in sequence without planning or structuring the calls, allowing the LLM to decide when and how to use external tools spontaneously.

Answer: A

Explanation:
The task requires an LLM pipeline for multi-stage reasoning with external tools, necessitating planning, adaptability, and complex reasoning. Let's evaluate the options based on Databricks' recommendations for advanced LLM workflows.
* Option A: Train the LLM to generate a single, comprehensive response without interacting with any external tools, relying solely on its pre-trained knowledge
* This approach limits the LLM to its static knowledge base, excluding external tools and multi- stage reasoning. It can't adapt or plan actions dynamically, failing the requirements.
* Databricks Reference:"External tools enhance LLM capabilities beyond pre-trained knowledge" ("Building LLM Applications with Databricks," 2023).
* Option B: Implement a framework like ReAct which allows the LLM to generate reasoning traces and perform task-specific actions that leverage external tools if necessary
* ReAct (Reasoning + Acting) combines reasoning traces (step-by-step logic) with actions (e.g., tool calls), enabling the LLM to plan, adapt, and execute complex tasks iteratively. This meets all requirements: multi-stage reasoning, tool use, and adaptability.
* Databricks Reference:"Frameworks like ReAct enable LLMs to interleave reasoning and external tool interactions for complex problem-solving"("Generative AI Cookbook," 2023).
* Option C: Encourage the LLM to make multiple API calls in sequence without planning or structuring the calls, allowing the LLM to decide when and how to use external tools spontaneously
* Unstructured, spontaneous API calls lack planning and may lead to inefficient or incorrect tool usage. This doesn't ensure effective multi-stage reasoning or adaptability.
* Databricks Reference: Structured frameworks are preferred:"Ad-hoc tool calls can reduce reliability in complex tasks"("Building LLM-Powered Applications").
* Option D: Use a Chain-of-Thought (CoT) prompting technique to guide the LLM through a series of reasoning steps, then manually input the results from external tools for the final answer
* CoT improves reasoning but relies on manual tool interaction, breaking automation and adaptability. It's not a scalable pipeline

solution.
* Databricks Reference:"Manual intervention is impractical for production LLM pipelines" ("Databricks Generative AI Engineer Guide").
Conclusion: Option B (ReAct) is the best approach, as it integrates reasoning and tool use in a structured, adaptive framework, aligning with Databricks' guidance for complex LLM workflows.

NEW QUESTION # 22
.....

You must want to know your scores after finishing exercising our Databricks-Generative-AI-Engineer-Associate study materials, which help you judge your revision. Now, our windows software and online test engine of the Databricks-Generative-AI-Engineer-Associate study materials can meet your requirements. You can choose from two modules: virtual exam and practice exam. Then you are required to answer every question of the Databricks-Generative-AI-Engineer-Associate Study Materials. In order to make sure you have answered all questions, we have answer list to help you check.

Databricks-Generative-AI-Engineer-Associate Exam Simulator: https://www.fast2test.com/Databricks-Generative-AI-Engineer-Associate-premium-file.html

Our Databricks Certified Generative AI Engineer Associate Preparation Material provides you everything you will need to take a Databricks Databricks-Generative-AI-Engineer-Associate examination, If you fail in the Databricks-Generative-AI-Engineer-Associate exam, we promise to give you a full refund with normal procedures; or you can freely change for another exam study material, Databricks Databricks-Generative-AI-Engineer-Associate Test Cram Review However, since not all takers have the same learning styles, we devise a customizable module to suite your needs, The Databricks-Generative-AI-Engineer-Associate PDF file is convenient for reading and printing.

But they must have substance and be simple, direct, concise, readable, Databricks-Generative-AI-Engineer-Associate and jargon-free, How Can a Base Class Protect Derived Classes So That Changes to the Base Class Will Not Affect Them?

Databricks Databricks-Generative-AI-Engineer-Associate Test Cram Review & Free PDF Unparalleled Databricks Certified Generative AI Engineer Associate

Our Databricks Certified Generative AI Engineer Associate Preparation Material provides you everything you will need to take a Databricks Databricks-Generative-AI-Engineer-Associate examination, If you fail in the Databricks-Generative-AI-Engineer-Associate exam, we promise to give you a full refund with normal procedures; or you can freely change for another exam study material.

However, since not all takers have the same learning styles, we devise a customizable module to suite your needs, The Databricks-Generative-AI-Engineer-Associate PDF file is convenient for reading and printing.

The desktop Databricks-Generative-AI-Engineer-Associate practice test software specifies the web-based version.

- Reliable Databricks-Generative-AI-Engineer-Associate Test Cram Review - Practical - First-Grade Databricks-Generative-AI-Engineer-Associate Materials Free Download for Databricks Databricks-Generative-AI-Engineer-Associate Exam □ Search on ⇒ www.troytecdumps.com ⇐ for ➽ Databricks-Generative-AI-Engineer-Associate □ to obtain exam materials for free download □Practice Databricks-Generative-AI-Engineer-Associate Exams Free
- Validate Your Skills with Databricks Databricks-Generative-AI-Engineer-Associate Exam Questions □ Search for ➽ Databricks-Generative-AI-Engineer-Associate □ and download it for free immediately on ➡ www.pdfvce.com □□□ □ □Latest Databricks-Generative-AI-Engineer-Associate Training
- Money Back Guarantee on Databricks Databricks-Generative-AI-Engineer-Associate Exam Questions □ Enter □ www.pass4test.com □ and search for [Databricks-Generative-AI-Engineer-Associate] to download for free □Exam Databricks-Generative-AI-Engineer-Associate Format
- Databricks-Generative-AI-Engineer-Associate Pass-Sure Torrent - Databricks-Generative-AI-Engineer-Associate Actual Braindumps - Databricks-Generative-AI-Engineer-Associate Test Cram □ Open website 「 www.pdfvce.com 」 and search for ⇒ Databricks-Generative-AI-Engineer-Associate ⇐ for free download □Databricks-Generative-AI-Engineer-Associate Reliable Test Online
- Databricks-Generative-AI-Engineer-Associate Pass-Sure Torrent - Databricks-Generative-AI-Engineer-Associate Actual Braindumps - Databricks-Generative-AI-Engineer-Associate Test Cram □ Search for 【 Databricks-Generative-AI-Engineer-Associate 】 on □ www.easy4engine.com □ immediately to obtain a free download □Databricks-Generative-AI-Engineer-Associate Valid Study Guide

- 2026 Databricks Databricks-Generative-AI-Engineer-Associate: Efficient Databricks Certified Generative AI Engineer Associate Test Cram Review □ Open ▸ www.pdfvce.com ◂ enter " Databricks-Generative-AI-Engineer-Associate " and obtain a free download □Exam Databricks-Generative-AI-Engineer-Associate Simulator Online
- Reliable Databricks-Generative-AI-Engineer-Associate Dumps Ppt □ Exam Databricks-Generative-AI-Engineer-Associate Exercise □ Latest Databricks-Generative-AI-Engineer-Associate Training □ Download （ Databricks-Generative-AI-Engineer-Associate ） for free by simply entering ▷ www.testkingpass.com ◁ website □Practice Databricks-Generative-AI-Engineer-Associate Exams Free
- Databricks-Generative-AI-Engineer-Associate Pass-Sure Torrent - Databricks-Generative-AI-Engineer-Associate Actual Braindumps - Databricks-Generative-AI-Engineer-Associate Test Cram □ The page for free download of ✔ Databricks-Generative-AI-Engineer-Associate □✔□ on 「 www.pdfvce.com 」 will open immediately □Exam Databricks-Generative-AI-Engineer-Associate Fees
- Databricks-Generative-AI-Engineer-Associate dumps: Databricks Certified Generative AI Engineer Associate - Databricks-Generative-AI-Engineer-Associate exam VCE ➡ Open ⇛ www.examcollectionpass.com ⇚ and search for ▸ Databricks-Generative-AI-Engineer-Associate ◂ to download exam materials for free □Latest Databricks-Generative-AI-Engineer-Associate Study Materials
- Practice Databricks-Generative-AI-Engineer-Associate Exams Free □ Databricks-Generative-AI-Engineer-Associate Reliable Test Online □ Databricks-Generative-AI-Engineer-Associate Actual Questions □ Search for ▷ Databricks-Generative-AI-Engineer-Associate ◁ on " www.pdfvce.com " immediately to obtain a free download □Databricks-Generative-AI-Engineer-Associate Valid Study Guide
- Databricks-Generative-AI-Engineer-Associate Valid Test Tips □ Databricks-Generative-AI-Engineer-Associate Actual Questions □ Exam Databricks-Generative-AI-Engineer-Associate Fees ➽ Easily obtain ☀ Databricks-Generative-AI-Engineer-Associate □☀□ for free download through □ www.prepawayete.com □ 👆Exam Databricks-Generative-AI-Engineer-Associate Simulator Online
- www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, zenwriting.net, www.stes.tyc.edu.tw, peeruu.com, www.stes.tyc.edu.tw, study.stcs.edu.np, www.stes.tyc.edu.tw, ascentleadershipinstitute.org, 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, Disposable vapes

BTW, DOWNLOAD part of Fast2test Databricks-Generative-AI-Engineer-Associate dumps from Cloud Storage: https://drive.google.com/open?id=1FeeXxvFzpCcB6OXLqpW39Q_iEdOhye1D