

Databricks Certified Generative AI Engineer Associate exam vce torrent & Databricks-Generative-AI-Engineer-Associate pdf dumps & Databricks Certified Generative AI Engineer Associate valid study prep



BTW, DOWNLOAD part of ExamDiscuss Databricks-Generative-AI-Engineer-Associate dumps from Cloud Storage:
<https://drive.google.com/open?id=1vmkMIRhizPJIU8IS9BGQz3jicqUNYSr>

If you want to pass Databricks Databricks-Generative-AI-Engineer-Associate exam and get a high paying job in the industry; if you are searching for the perfect Databricks-Generative-AI-Engineer-Associate exam prep material to get your dream job, then you must consider using our Databricks Certified Generative AI Engineer Associate exam products to improve your skillset. We have curated new Databricks-Generative-AI-Engineer-Associate Questions Answers to help you prepare for the exam. It can be your golden ticket to pass the Databricks Databricks-Generative-AI-Engineer-Associate test on the first attempt. We are providing latest Databricks-Generative-AI-Engineer-Associate PDF question answers to help you prepare exam while working in the office to save your time.

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

Topic	Details
Topic 1	<ul style="list-style-type: none">• Governance: Generative AI Engineers who take the exam get knowledge about masking techniques, guardrail techniques, and legal licensing requirements in this topic.
Topic 2	<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 3	<ul style="list-style-type: none">• Application Development: In this topic, Generative AI Engineers learn about tools needed to extract data, Langchain• similar tools, and assessing responses to identify common issues. Moreover, the topic includes questions about adjusting an LLM's response, LLM guardrails, and the best LLM based on the attributes of the application.
Topic 4	<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.

Topic 5	<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 filtering extraneous content in source documents. Lastly, Generative AI Engineers also learn about extracting document content from provided source data and format.
---------	---

>> Databricks-Generative-AI-Engineer-Associate Study Tool <<

Databricks-Generative-AI-Engineer-Associate Actual Collection: Databricks Certified Generative AI Engineer Associate - Databricks-Generative-AI-Engineer-Associate Quiz Braindumps & Databricks-Generative-AI-Engineer-Associate Exam Guide

In order to let you have a deep understanding of our Databricks-Generative-AI-Engineer-Associate learning guide, our company designed the trial version for our customers. We will provide you with the trial version of our study materials before you buy our products. If you want to know our Databricks-Generative-AI-Engineer-Associate training materials, you can download the trial version from the web page of our company. If you use the trial version of our Databricks-Generative-AI-Engineer-Associate Study Materials, you will find that our products are very useful for you to pass your exam and get the certification. If you buy our Databricks-Generative-AI-Engineer-Associate exam questions, we can promise that you will enjoy a discount.

Databricks Certified Generative AI Engineer Associate Sample Questions (Q52-Q57):

NEW QUESTION # 52

A Generative AI Engineer is building a RAG application that will rely on context retrieved from source documents that are currently in PDF format. These PDFs can contain both text and images. They want to develop a solution using the least amount of lines of code.

Which Python package should be used to extract the text from the source documents?

- A. **unstructured**
- B. beautifulsoup
- C. flask
- D. numpy

Answer: A

Explanation:

* Problem Context: The engineer needs to extract text from PDF documents, which may contain both text and images. The goal is to find a Python package that simplifies this task using the least amount of code.

* Explanation of Options:

- * Option A: flask: Flask is a web framework for Python, not suitable for processing or extracting content from PDFs.
- * Option B: beautifulsoup: Beautiful Soup is designed for parsing HTML and XML documents, not PDFs.
- * Option C: unstructured: This Python package is specifically designed to work with unstructured data, including extracting text from PDFs. It provides functionalities to handle various types of content in documents with minimal coding, making it ideal for the task.
- * Option D: numpy: Numpy is a powerful library for numerical computing in Python and does not provide any tools for text extraction from PDFs.

Given the requirement, Option C (unstructured) is the most appropriate as it directly addresses the need to efficiently extract text from PDF documents with minimal code.

NEW QUESTION # 53

A Generative AI Engineer is building a production-ready LLM system which replies directly to customers.

The solution makes use of the Foundation Model API via provisioned throughput. They are concerned that the LLM could potentially respond in a toxic or otherwise unsafe way. They also wish to perform this with the least amount of effort.

Which approach will do this?

- A. Host Llama Guard on Foundation Model API and use it to detect unsafe responses

- B. Ask users to report unsafe responses
- C. Add some LLM calls to their chain to detect unsafe content before returning text
- D. Add a regex expression on inputs and outputs to detect unsafe responses.

Answer: A

Explanation:

The task is to prevent toxic or unsafe responses in an LLM system using the Foundation Model API with minimal effort. Let's assess the options.

- * Option A: Host Llama Guard on Foundation Model API and use it to detect unsafe responses
 - * Llama Guard is a safety-focused model designed to detect toxic or unsafe content. Hosting it via the Foundation Model API (a Databricks service) integrates seamlessly with the existing system, requiring minimal setup (just deployment and a check step), and leverages provisioned throughput for performance.
 - * Databricks Reference: "Foundation Model API supports hosting safety models like Llama Guard to filter outputs efficiently" ("Foundation Model API Documentation," 2023).
- * Option B: Add some LLM calls to their chain to detect unsafe content before returning text
 - * Using additional LLM calls (e.g., prompting an LLM to classify toxicity) increases latency, complexity, and effort (crafting prompts, chaining logic), and lacks the specificity of a dedicated safety model.
- * Option C: Add a regex expression on inputs and outputs to detect unsafe responses
 - * Regex can catch simple patterns (e.g., profanity) but fails for nuanced toxicity (e.g., sarcasm, context-dependent harm), requiring significant manual effort to maintain and update rules.
- * Option D: Ask users to report unsafe responses
 - * User reporting is reactive, not preventive, and places burden on users rather than the system. It doesn't limit unsafe outputs proactively and requires additional effort for feedback handling.

* Databricks Reference: "Proactive guardrails are preferred over user-driven monitoring" ("Databricks Generative AI Engineer Guide").

Conclusion: Option A (Llama Guard on Foundation Model API) is the least-effort, most effective approach, leveraging Databricks' infrastructure for seamless safety integration.

NEW QUESTION # 54

A Generative AI Engineer is tasked with deploying an application that takes advantage of a custom MLflow Pyfunc model to return some interim results.

How should they configure the endpoint to pass the secrets and credentials?

- A. Pass the secrets in plain text
- B. Use spark.conf.set()
- C. Pass variables using the Databricks Feature Store API
- D. Add credentials using environment variables

Answer: D

Explanation:

Context: Deploying an application that uses an MLflow Pyfunc model involves managing sensitive information such as secrets and credentials securely.

Explanation of Options:

- * Option A: Use spark.conf.set(): While this method can pass configurations within Spark jobs, using it for secrets is not recommended because it may expose them in logs or Spark UI.
- * Option B: Pass variables using the Databricks Feature Store API: The Feature Store API is designed for managing features for machine learning, not for handling secrets or credentials.
- * Option C: Add credentials using environment variables: This is a common practice for managing credentials in a secure manner, as environment variables can be accessed securely by applications without exposing them in the codebase.
- * Option D: Pass the secrets in plain text: This is highly insecure and not recommended, as it exposes sensitive information directly in the code.

Therefore, Option C is the best method for securely passing secrets and credentials to an application, protecting them from exposure.

NEW QUESTION # 55

A Generative AI 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, call_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. **transcript Volume**
- E. call_rep_history

Answer: B,D

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 # 56

A Generative AI Engineer has already trained an LLM on Databricks and it is now ready to be deployed.

Which of the following steps correctly outlines the easiest process for deploying a model on Databricks?

- A. Save the model along with its dependencies in a local directory, build the Docker image, and run the Docker container
- B. Log the model as a pickle object, upload the object to Unity Catalog Volume, register it to Unity Catalog using MLflow, and start a serving endpoint
- C. **Log the model using MLflow during training, directly register the model to Unity Catalog using the MLflow API, and start a serving endpoint**
- D. Wrap the LLM's prediction function into a Flask application and serve using Gunicorn

Answer: C

NEW QUESTION # 57

.....

Passing Databricks certification Databricks-Generative-AI-Engineer-Associate exam is not simple. Choose the right training is the first step to your success and choose a good resource of information is your guarantee of success. While the product of ExamDiscuss is a good guarantee of the resource of information. If you choose the ExamDiscuss product, it not only can 100% guarantee you to pass Databricks Certification Databricks-Generative-AI-Engineer-Associate Exam but also provide you with a year-long free update.

Exam Databricks-Generative-AI-Engineer-Associate Introduction:

<https://www.examdiscuss.com/Databricks/exam/Databricks-Generative-AI-Engineer-Associate/>

- Professional Databricks-Generative-AI-Engineer-Associate Study Tool - 100% Pass Databricks-Generative-AI-Engineer-Associate Exam Copy URL www.dumpsquestion.com open and search for [Databricks-Generative-AI-Engineer-Associate](#) to download for free Training Databricks-Generative-AI-Engineer-Associate Pdf
- Pass Guaranteed Quiz Accurate Databricks - Databricks-Generative-AI-Engineer-Associate - Databricks Certified Generative AI Engineer Associate Study Tool Search for [Databricks-Generative-AI-Engineer-Associate](#) on www.pdfvce.com immediately to obtain a free download Relevant Databricks-Generative-AI-Engineer-Associate Questions
- Relevant Databricks-Generative-AI-Engineer-Associate Questions Databricks-Generative-AI-Engineer-Associate Exam Questions Vce Databricks-Generative-AI-Engineer-Associate Latest Exam Cost www.prepawayte.com is best website to obtain "Databricks-Generative-AI-Engineer-Associate" for free download Relevant Databricks-Generative-AI-Engineer-Associate Questions
- Pass Guaranteed Updated Databricks - Databricks-Generative-AI-Engineer-Associate Study Tool www.pdfvce.com is best website to obtain [Databricks-Generative-AI-Engineer-Associate](#) for free download Latest Databricks-Generative-AI-Engineer-Associate Real Test
- 100% Pass Quiz 2026 Databricks - Databricks-Generative-AI-Engineer-Associate Study Tool Search for [Databricks-Generative-AI-Engineer-Associate](#) and download it for free immediately on { www.pdfdumps.com } Latest Databricks-Generative-AI-Engineer-Associate Real Test
- Latest Databricks-Generative-AI-Engineer-Associate Exam Labs Reliable Databricks-Generative-AI-Engineer-Associate Exam Blueprint Databricks-Generative-AI-Engineer-Associate Pdf Braindumps Copy URL www.pdfvce.com open and search for [Databricks-Generative-AI-Engineer-Associate](#) to download for free Databricks-Generative-AI-Engineer-Associate Test Simulator Fee
- 100% Pass Quiz 2026 Databricks - Databricks-Generative-AI-Engineer-Associate Study Tool Immediately open www.practicevce.com and search for [Databricks-Generative-AI-Engineer-Associate](#) to obtain a free download New Databricks-Generative-AI-Engineer-Associate Dumps
- Databricks-Generative-AI-Engineer-Associate Exam Questions Vce Databricks-Generative-AI-Engineer-Associate Test Simulator Fee Databricks-Generative-AI-Engineer-Associate Exam Questions Vce Copy URL "www.pdfvce.com" open and search for [Databricks-Generative-AI-Engineer-Associate](#) to download for free Relevant Databricks-Generative-AI-Engineer-Associate Questions
- Pass Guaranteed Updated Databricks - Databricks-Generative-AI-Engineer-Associate Study Tool Search for [Databricks-Generative-AI-Engineer-Associate](#) on www.vceengine.com immediately to obtain a free download New Databricks-Generative-AI-Engineer-Associate Dumps
- Databricks-Generative-AI-Engineer-Associate Study Tool Exam Instant Download | Updated Databricks Exam Databricks-Generative-AI-Engineer-Associate Introduction Immediately open "www.pdfvce.com" and search for ([Databricks-Generative-AI-Engineer-Associate](#)) to obtain a free download New Databricks-Generative-AI-Engineer-Associate Exam Pass4sure
- Pass Guaranteed Quiz Accurate Databricks - Databricks-Generative-AI-Engineer-Associate - Databricks Certified Generative AI Engineer Associate Study Tool Open website [www.practicevce.com] and search for { [Databricks-Generative-AI-Engineer-Associate](#) } for free download Latest Databricks-Generative-AI-Engineer-Associate Exam Labs
- www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, confengine.com, hhi.instructure.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, www.stes.tyc.edu.tw, myspace.com, www.stes.tyc.edu.tw, skyhighes.in, www.stes.tyc.edu.tw, Disposable vapes

P.S. Free 2026 Databricks Databricks-Generative-AI-Engineer-Associate dumps are available on Google Drive shared by ExamDiscuss: <https://drive.google.com/open?id=1vmkMIRhizPJIU8IS9BGQz3jicqUNYSr>

