

Databricks-Generative-AI-Engineer-Associate유효한최신 덤프공부 & Databricks-Generative-AI-Engineer- Associate최신덤프데모다운



2026 Itcertkr 최신 Databricks-Generative-AI-Engineer-Associate PDF 버전 시험 문제집과 Databricks-Generative-AI-Engineer-Associate 시험 문제 및 답변 무료 공유: <https://drive.google.com/open?id=1EuWXjcnBn63VIUNWFwo9BeJD7K5tJ-HL>

Itcertkr의Databricks Databricks-Generative-AI-Engineer-Associate 인증시험덤프는 자주 업데이트 되고, 오래 되고 더 이상 사용 하지 않는 문제들은 바로 삭제해버리며 새로운 최신 문제들을 추가 합니다. 이는 응시자가 확실하고도 빠르게Databricks Databricks-Generative-AI-Engineer-Associate덤프를 마스터하고Databricks Databricks-Generative-AI-Engineer-Associate시험을 패스할수 있도록 하는 또 하나의 보장입니다. 매력만점Databricks Databricks-Generative-AI-Engineer-Associate덤프 강력 추천합니다.

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

Section	Weight	Objectives
Topic 1: Data Preparation	14%	- Manage data quality and format - Implement data pipelines and transformations - Prepare and process data for LLM use
Topic 2: Evaluation and Monitoring	12%	- Iterate and improve solutions - Evaluate model and application performance - Monitor application behavior and outputs
Topic 3: Application Development	30%	- Integrate with MLflow - Build RAG applications - Develop LLM chains and workflows - Use Databricks Vector Search
Topic 4: Assembling and Deploying Apps	22%	- Manage application lifecycle - Deploy models via Model Serving - Package and deploy applications
Topic 5: Governance	8%	- Apply Unity Catalog for data governance - Manage access and permissions - Ensure compliance and security
Topic 6: Design Applications	14%	- Select appropriate models, tools and approaches - Decompose complex requirements into tasks - Design solution architecture

>> Databricks-Generative-AI-Engineer-Associate유효한 최신덤프공부 <<

최신버전 Databricks-Generative-AI-Engineer-Associate유효한 최신덤프공

부 덤프는 Databricks Certified Generative AI Engineer Associate 시험패스의 지름길

Demo를 다운받아 Databricks Databricks-Generative-AI-Engineer-Associate 덤프의 일부분 문제를 체험해보시고 구매하셔도 됩니다. 저희 Itcertkr에서는 Databricks Databricks-Generative-AI-Engineer-Associate 덤프의 일부분 문제를 샘플로 제공해드립니다. 덤프만 열공하시면 Databricks Databricks-Generative-AI-Engineer-Associate 시험패스가 가능하기에 저희 자료를 선택한걸 후회하지 않게 할 자신이 있습니다.

최신 Generative AI Engineer Databricks-Generative-AI-Engineer-Associate 무료샘플문제 (Q33-Q38):

질문 # 33

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. Throttle the incoming batch of requests manually to avoid rate limiting issues
- D. Switch to using External Models instead

정답: A

설명:

* 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.

Option B is ideal, offering flexibility and cost control, aligning expenses directly with the application's usage patterns.

질문 # 34

A Generative AI Engineer is developing a RAG application and would like to experiment with different embedding models to improve the application performance.

Which strategy for picking an embedding model should they choose?

- A. Pick the most recent and most performant open LLM released at the time
- B. Pick an embedding model trained on related domain knowledge
- C. Pick an embedding model with multilingual support to support potential multilingual user questions
- D. Pick the embedding model ranked highest on the Massive Text Embedding Benchmark (MTEB) leaderboard hosted by HuggingFace

정답: B

설명:

The task involves improving a Retrieval-Augmented Generation (RAG) application's performance by experimenting with embedding models. The choice of embedding model impacts retrieval accuracy, which is critical for RAG systems. Let's evaluate the options based on Databricks Generative AI Engineer best practices.

Option A: Pick an embedding model trained on related domain knowledge

Embedding models trained on domain-specific data (e.g., industry-specific corpora) produce vectors that better capture the semantics of the application's context, improving retrieval relevance. For RAG, this is a key strategy to enhance performance.

Databricks Reference: "For optimal retrieval in RAG systems, select embedding models aligned with the domain of your data" ("Building LLM Applications with Databricks," 2023).

Option B: Pick the most recent and most performant open LLM released at the time LLMs are not embedding models; they generate text, not embeddings for retrieval. While recent LLMs may be performant for generation, this doesn't address the embedding step in RAG. This option misunderstands the component being selected.

Databricks Reference: Embedding models and LLMs are distinct in RAG workflows: "Embedding models convert text to vectors, while LLMs generate responses" ("Generative AI Cookbook").

Option C: Pick the embedding model ranked highest on the Massive Text Embedding Benchmark (MTEB) leaderboard hosted by HuggingFace. The MTEB leaderboard ranks models across general tasks, but high overall performance doesn't guarantee suitability for a specific domain. A top-ranked model might excel in generic contexts but underperform on the engineer's unique data.

Databricks Reference: General performance is less critical than domain fit: "Benchmark rankings provide a starting point, but domain-specific evaluation is recommended" ("Databricks Generative AI Engineer Guide").

Option D: Pick an embedding model with multilingual support to support potential multilingual user questions. Multilingual support is useful only if the application explicitly requires it. Without evidence of multilingual needs, this adds complexity without guaranteed performance gains for the current use case.

Databricks Reference: "Choose features like multilingual support based on application requirements" ("Building LLM-Powered Applications").

Conclusion: Option A is the best strategy because it prioritizes domain relevance, directly improving retrieval accuracy in a RAG system-aligning with Databricks' emphasis on tailoring models to specific use cases.

질문 # 35

A Generative AI Engineer has a provisioned throughput model serving endpoint as part of a RAG application and would like to monitor the serving endpoint's incoming requests and outgoing responses. The current approach is to include a micro-service in between the endpoint and the user interface to write logs to a remote server.

Which Databricks feature should they use instead which will perform the same task?

- A. DBSQL
- B. Vector Search
- C. Lakeview
- **D. Inference Tables**

정답: D

설명:

Problem Context: The goal is to monitor the serving endpoint for incoming requests and outgoing responses in a provisioned throughput model serving endpoint within a Retrieval-Augmented Generation (RAG) application. The current approach involves using a microservice to log requests and responses to a remote server, but the Generative AI Engineer is looking for a more streamlined solution within Databricks.

Explanation of Options:

* Option A: Vector Search: This feature is used to perform similarity searches within vector databases.

It doesn't provide functionality for logging or monitoring requests and responses in a serving endpoint, so it's not applicable here.

* Option B: Lakeview: Lakeview is not a feature relevant to monitoring or logging request-response cycles for serving endpoints. It might be more related to viewing data in Databricks Lakehouse but doesn't fulfill the specific monitoring requirement.

* Option C: DBSQL: Databricks SQL (DBSQL) is used for running SQL queries on data stored in Databricks, primarily for analytics purposes. It doesn't provide the direct functionality needed to monitor requests and responses in real-time for an inference endpoint.

* Option D: Inference Tables: This is the correct answer. Inference Tables in Databricks are designed to store the results and metadata of inference runs. This allows the system to log incoming requests and outgoing responses directly within Databricks, making it an ideal choice for monitoring the behavior of a provisioned serving endpoint. Inference Tables can be queried and analyzed, enabling easier monitoring and debugging compared to a custom microservice.

Thus, Inference Tables are the optimal feature for monitoring request and response logs within the Databricks infrastructure for a model serving endpoint.

질문 # 36

A Generative AI 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. 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.
- C. Train the LLM to generate a single, comprehensive response without interacting with any external tools, relying solely on its pre-trained knowledge.
- 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.

정답: A

설명:

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.

질문 # 37

A Generative AI Engineer would like an LLM to generate formatted JSON from emails. This will require parsing and extracting the following information: order ID, date, and sender email. Here's a sample email:

Date: April 23, 2024
Time: 4:22 PM
From: anjali.thayer@computex.org
To: cust_service@realtek.com
Subject: Shipment details

Hey there,

I have a shipment (order ID is CD34RFT) can you please send me an update?

Thank you,
Anjali

They will need to write a prompt that will extract the relevant information in JSON format with the highest level of output accuracy. Which prompt will do that?

- A. You will receive customer emails and need to extract date, sender email, and order ID. Return the extracted information in JSON format.
- B. You will receive customer emails and need to extract date, sender email, and order ID. Return the extracted information in a human-readable format.

- C. You will receive customer emails and need to extract date, sender email, and order ID. Return the extracted information in JSON format.
Here's an example: {"date": "April 16, 2024", "sender_email": "sarah.lee925@gmail.com", "order_id": "RE987D"}
- D. You will receive customer emails and need to extract date, sender email, and order ID. You should return the date, sender email, and order ID information in JSON format.

정답: C

설명:

* Problem Context: The goal is to parse emails to extract certain pieces of information and output this in a structured JSON format. Clarity and specificity in the prompt design will ensure higher accuracy in the LLM's responses.

* Explanation of Options:

Option A: Provides a general guideline but lacks an example, which helps an LLM understand the exact format expected.

Option B: Includes a clear instruction and a specific example of the output format. Providing an example is crucial as it helps set the pattern and format in which the information should be structured, leading to more accurate results.

Option C: Does not specify that the output should be in JSON format, thus not meeting the requirement.

Option D: While it correctly asks for JSON format, it lacks an example that would guide the LLM on how to structure the JSON correctly.

Therefore, Option B is optimal as it not only specifies the required format but also illustrates it with an example, enhancing the likelihood of accurate extraction and formatting by the LLM.

질문 # 38

.....

Itcertkr는 여러분이 빠른 시일 내에 Databricks Databricks-Generative-AI-Engineer-Associate 인증시험을 효과적으로 터득할 수 있는 사이트입니다. Databricks Databricks-Generative-AI-Engineer-Associate 덤프는 보장하는 덤프입니다. 만약 시험에서 떨어지셨다고 하면 우리는 무조건 덤프전액 환불을 약속 드립니다. 우리 Itcertkr 사이트에서 Databricks Databricks-Generative-AI-Engineer-Associate 관련 자료의 일부분 문제와 답 등 샘플을 제공함으로 여러분은 무료로 다 운 받아 체험해보실 수 있습니다. 체험 후 우리의 Itcertkr에 신뢰감을 느끼게 됩니다. Itcertkr의 Databricks Databricks-Generative-AI-Engineer-Associate 덤프로 자신 있는 시험준비를 하세요.

Databricks-Generative-AI-Engineer-Associate 최신 덤프 데모 다운 : https://www.itcertkr.com/Databricks-Generative-AI-Engineer-Associate_exam.html

- Databricks-Generative-AI-Engineer-Associate 최고 품질 인증시험 대비자료 □ Databricks-Generative-AI-Engineer-Associate 최신 핫덤프 □ Databricks-Generative-AI-Engineer-Associate 인증시험 공부 □ 시험 자료를 무료로 다운로드하려면 ☀ www.dumpstap.com □ ☀ □ 을 통해 ➡ Databricks-Generative-AI-Engineer-Associate □ 를 검색하십시오 □ 오 Databricks-Generative-AI-Engineer-Associate 인증시험 공부
- Databricks-Generative-AI-Engineer-Associate 퍼펙트 최신 버전 자료 □ Databricks-Generative-AI-Engineer-Associate 퍼펙트 덤프 최신 데모 □ Databricks-Generative-AI-Engineer-Associate 응시자료 □ 무료로 다운로드하려면 > www.itdumps.com <로 이동하여 “Databricks-Generative-AI-Engineer-Associate”를 검색하십시오 □ 오 Databricks-Generative-AI-Engineer-Associate 덤프 문제
- Databricks-Generative-AI-Engineer-Associate 퍼펙트 덤프 최신 데모 □ Databricks-Generative-AI-Engineer-Associate 퍼펙트 덤프 공부자료 □ Databricks-Generative-AI-Engineer-Associate 최고 품질 인증시험 대비자료 ↔ (www.koreadumps.com) 에서 ➡ Databricks-Generative-AI-Engineer-Associate □ □ □ 를 검색하고 무료로 다운로드 하세요 □ 오 Databricks-Generative-AI-Engineer-Associate 시험대비 공부 문제
- 시험대비 Databricks-Generative-AI-Engineer-Associate 유효한 최신덤프 공부 최신 버전 덤프 샘플 문제 다운 받기 □ □ ▶ www.itdumps.com <웹사이트를 열고 ➡ Databricks-Generative-AI-Engineer-Associate □ 를 검색하여 무료 다운로드 Databricks-Generative-AI-Engineer-Associate PDF
- 퍼펙트한 Databricks-Generative-AI-Engineer-Associate 유효한 최신덤프 공부 최신 덤프 공부 □ □ kr.fast2test.com □ 웹사이트에서 > Databricks-Generative-AI-Engineer-Associate □ 를 열고 검색하여 무료 다운로드 Databricks-Generative-AI-Engineer-Associate 인증시험대비 공부자료
- Databricks-Generative-AI-Engineer-Associate 최고 품질 인증시험 대비자료 □ Databricks-Generative-AI-Engineer-Associate 덤프 최신자료 □ Databricks-Generative-AI-Engineer-Associate 높은 통과율 공부자료 □ ➡ www.itdumps.com □ 은 ➡ Databricks-Generative-AI-Engineer-Associate □ □ □ 무료 다운로드를 받을 수 있는 최고의 사이트입니다 □ 오 Databricks-Generative-AI-Engineer-Associate 최신 핫덤프
- Databricks-Generative-AI-Engineer-Associate 유효한 최신덤프 공부 시험대비 덤프 공부 □ 무료 다운로드를 위해 지금 ➡ www.pass4test.net □ 에서 「 Databricks-Generative-AI-Engineer-Associate 」 검색 □ 오 Databricks-Generative-AI-Engineer-Associate 최고 품질 인증시험 대비자료
- Databricks-Generative-AI-Engineer-Associate 퍼펙트 최신 버전 자료 □ Databricks-Generative-AI-Engineer-Associate

