높은적중율을자랑하는Databricks-Generative-AI-Engineer-Associate시험유형덤프자료

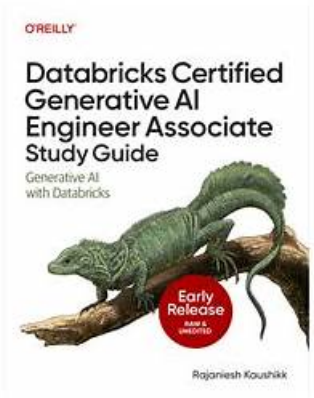

Pass4Test Databricks-Generative-AI-Engineer-Associate 최신 PDF 버전 시험 문제집을 무료로 Google Drive에서 다운로드하세요: https://drive.google.com/open?id=1I8aytSzKMG5uh11TWlD8qC6qNKISvArH

최근 Databricks인증 Databricks-Generative-AI-Engineer-Associate시험이 IT업계에서 제일 높은 인지도를 가지고 있습니다.바라만 보지 마시고Databricks인증 Databricks-Generative-AI-Engineer-Associate시험에 도전해보세요. Pass4Test 의 Databricks인증 Databricks-Generative-AI-Engineer-Associate덤프로 시험준비공부를 하시면 한방에 시험패스 가능합니다. Databricks인증 Databricks-Generative-AI-Engineer-Associate덤프로 자격증취득에 가까워지고 나아가서는 IT업계에서 인정을 받는 열쇠를 소유한것과 같다고 할수 있습니다.

Pass4Test의Databricks인증 Databricks-Generative-AI-Engineer-Associate시험대비 덤프는 가격이 착한데 비하면 품질이 너무 좋은 시험전 공부자료입니다. 시험문제적중율이 높아 패스율이 100%에 이르고 있습니다.다른 IT자격증에 관심이 있는 분들은 온라인서비스에 문의하여 덤프유무와 적중율등을 확인할수 있습니다. Databricks인증 Databricks-Generative-AI-Engineer-Associate덤프로 어려운 시험을 정복하여 IT업계 정상에 오릅시다.

>> Databricks-Generative-AI-Engineer-Associate시험유형 <<

최신버전 Databricks-Generative-AI-Engineer-Associate시험유형 덤프공부문제

Pass4Test 에서는 최선을 다해 여러분이Databricks Databricks-Generative-AI-Engineer-Associate인증시험을 패스하도록 도울 것이며 여러분은 Pass4Test에서Databricks Databricks-Generative-AI-Engineer-Associate덤프의 일부분의 문제와 답을 무료로 다운받으실 수 잇습니다. Pass4Test 선택함으로Databricks Databricks-Generative-AI-Engineer-Associate인증시험통과는 물론Pass4Test 제공하는 일년무료 업데이트서비스를 제공받을 수 있으며 Pass4Test의 인증덤프로 시험에서 떨어졌다면 100% 덤프비용 전액환불을 약속 드립니다.

최신 Generative AI Engineer Databricks-Generative-AI-Engineer-Associate 무료샘플문제 (Q60-Q65):

질문 # 60
A Generative Al Engineer wants their (inetuned LLMs in their prod Databncks workspace available for testing in their dev workspace as well. All of their workspaces are Unity Catalog enabled and they are currently logging their models into the Model Registry in MLflow.
What is the most cost-effective and secure option for the Generative Al Engineer to accomplish their gAi?

- A. Use MLflow to log the model directly into Unity Catalog, and enable READ access in the dev workspace to the model.
- B. Setup a duplicate training pipeline in dev, so that an identical model is available in dev.
- C. Use an external model registry which can be accessed from all workspaces
- D. Setup a script to export the model from prod and import it to dev.

정답： A

설명:
The goal is to make fine-tuned LLMs from a production (prod) Databricks workspace available for testing in a development (dev) workspace, leveraging Unity Catalog and MLflow, while ensuring cost-effectiveness and security. Let's analyze the options.
* Option A: Use an external model registry which can be accessed from all workspaces
* An external registry adds cost (e.g., hosting fees) and complexity (e.g., integration, security configurations) outside Databricks' native ecosystem, reducing security compared to Unity Catalog's governance.
* Databricks Reference:"Unity Catalog provides a centralized, secure model registry within Databricks"("Unity Catalog Documentation," 2023).
* Option B: Setup a script to export the model from prod and import it to dev
* Export/import scripts require manual effort, storage for model artifacts, and repeated execution, increasing operational cost and risk (e.g., version mismatches, unsecured transfers). It's less efficient than a native solution.
* Databricks Reference: Manual processes are discouraged when Unity Catalog offers built-in sharing:"Avoid redundant workflows with Unity Catalog's cross-workspace access"("MLflow with Unity Catalog").
* Option C: Setup a duplicate training pipeline in dev, so that an identical model is available in dev
* Duplicating the training pipeline doubles compute and storage costs, as it retrains the model from scratch. It's neither cost-effective nor necessary when the prod model can be reused securely.
* Databricks Reference:"Re-running training is resource-intensive; leverage existing models where possible"("Generative AI Engineer Guide").
* Option D: Use MLflow to log the model directly into Unity Catalog, and enable READ access in the dev workspace to the model
* Unity Catalog, integrated with MLflow, allows models logged in prod to be centrally managed and accessed across workspaces with fine-grained permissions (e.g., READ for dev). This is cost- effective (no extra infrastructure or retraining) and secure (governed by Databricks' access controls).
* Databricks Reference:"Log models to Unity Catalog via MLflow, then grant access to other workspaces securely"("MLflow Model Registry with Unity Catalog," 2023).
Conclusion: Option D leverages Databricks' native tools (MLflow and Unity Catalog) for a seamless, cost- effective, and secure solution, avoiding external systems, manual scripts, or redundant training.

질문 # 61
A team uses Mosaic AI Vector Search to retrieve documents for their Retrieval-Augmented Generation (RAG) pipeline. The search query returns five relevant documents, and the first three are added to the prompt as context. Performance evaluation with Agent Evaluation shows that some lower-ranked retrieved documents have higher context relevancy scores than higher-ranked documents. Which option should the team consider to optimize this workflow?

- A. Use a reranker to order the documents based on the relevance scores.
- B. Modify the prompt to instruct the LLM to order the documents based on the relevance scores.
- C. Use a different embedding model for computing document embeddings.
- D. Increase the number of documents added to the prompt to improve context relevance.

정답: A

설명:
The scenario describes a common "retrieval gap" where the initial bi-encoder (embedding model) used for vector search identifies relevant documents but does not rank them perfectly. This happens because embedding models represent entire documents as a single vector, which can lose nuance. The standard engineering solution is to implement a Reranker (Cross-Encoder). Unlike embedding models, a reranker processes the query and a candidate document simultaneously, allowing it to capture deep semantic interactions between the two. In a Mosaic AI workflow, after the vector search retrieves the top k documents, the reranker evaluates those specific k documents to produce a more accurate relevance score. This ensures that the most contextually relevant documents are placed at the top of the list (and thus the top of the LLM prompt), which is crucial because LLMs are sensitive to document order and often prioritize information found at the beginning of the context.

질문 # 62
A Generative Al 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 Al Engineer's needs in this situation?

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

정답：D

설명：
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 areconfidentialityandsensitivityof 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-largeis assumed to be configurable for secure internal use, making it the optimal choice for this scenario.

질문 # 63
A Generative Al Engineer is working with a retail company that wants to enhance its customer experience by automatically handling common customer inquiries. They are working on an LLM-powered Al solution that should improve response times while maintaining a personalized interaction. They want to define the appropriate input and LLM task to do this.
Which input/output pair will do this?

- A. Input: Customer service chat logs; Output Group the chat logs by users, followed by summarizing each user's interactions, then respond
- B. Input: Customer reviews; Output Group the reviews by users and aggregate per-user average rating, then respond
- C. Input: Customer service chat logs; Output: Find the answers to similar questions and respond with a summary
- D. Input: Customer reviews: Output Classify review sentiment

정답：C

설명：
The task described in the question involves enhancing customer experience by automatically handling common customer inquiries using an LLM-powered AI solution. This requires the system to process input data (customer inquiries) and generate personalized, relevant responses efficiently. Let's evaluate the options step-by-step in the context of Databricks Generative AI Engineer principles, which emphasize leveraging LLMs for tasks like question answering, summarization, and retrieval-augmented generation (RAG).
Option A: Input: Customer reviews; Output: Group the reviews by users and aggregate per-user average rating, then respond This option focuses on analyzing customer reviews to compute average ratings per user. While this might be useful for sentiment analysis or user profiling, it does not directly address the goal of handling common customer inquiries or improving response times for personalized interactions. Customer reviews are typically feedback data, not real-time inquiries requiring immediate responses.
Databricks Reference: Databricks documentation on LLMs (e.g., "Building LLM Applications with Databricks") emphasizes that LLMs excel at tasks like question answering and conversational responses, not just aggregation or statistical analysis of reviews.
Option B: Input: Customer service chat logs; Output: Group the chat logs by users, followed by summarizing each user's interactions, then respond This option uses chat logs as input, which aligns with customer service scenarios. However, the output-grouping by users and summarizing interactions-focuses on user-specific summaries rather than directly addressing inquiries. While summarization is an LLM capability, this approach lacks the specificity of finding answers to common questions, which is central to the problem.
Databricks Reference: Per Databricks' "Generative AI Cookbook," LLMs can summarize text, but for customer service, the emphasis is on retrieval and response generation (e.g., RAG workflows) rather than user interaction summaries alone.
Option C: Input: Customer service chat logs; Output: Find the answers to similar questions and respond with a summary This option uses chat logs (real customer inquiries) as input and tasks the LLM with identifying answers to similar questions, then providing a summarized response. This directly aligns with the goal of handling common inquiries efficiently while maintaining personalization (by referencing past interactions or similar cases). It leverages LLM capabilities like semantic search, retrieval, and response generation,

which are core to Databricks' LLM workflows.
Databricks Reference: From Databricks documentation ("Building LLM-Powered Applications," 2023), an exact extract states: "For customer support use cases, LLMs can be used to retrieve relevant answers from historical data like chat logs and generate concise, contextually appropriate responses." This matches Option C's approach of finding answers and summarizing them.
Option D: Input: Customer reviews; Output: Classify review sentiment
This option focuses on sentiment classification of reviews, which is a valid LLM task but unrelated to handling customer inquiries or improving response times in a conversational context. It's more suited for feedback analysis than real-time customer service.
Databricks Reference: Databricks' "Generative AI Engineer Guide" notes that sentiment analysis is a common LLM task, but it's not highlighted for real-time conversational applications like customer support.
Conclusion: Option C is the best fit because it uses relevant input (chat logs) and defines an LLM task (finding answers and summarizing) that meets the requirements of improving response times and maintaining personalized interaction. This aligns with Databricks' recommended practices for LLM-powered customer service solutions, such as retrieval-augmented generation (RAG) workflows.

질문 # 64
A Generative Al Engineer is building a RAG application that answers questions about internal documents for the company SnoPen AI.
The source documents may contain a significant amount of irrelevant content, such as advertisements, sports news, or entertainment news, or content about other companies.
Which approach is advisable when building a RAG application to achieve this goal of filtering irrelevant information?

- A. Include in the system prompt that the application is not supposed to answer any questions unrelated to SnoPen Al.
- B. Keep all articles because the RAG application needs to understand non-company content to avoid answering questions about them.
- C. Consolidate all SnoPen AI related documents into a single chunk in the vector database.
- D. Include in the system prompt that any information it sees will be about SnoPenAI, even if no data filtering is performed.

정답：A

설명:
In a Retrieval-Augmented Generation (RAG) application built to answer questions about internal documents, especially when the dataset contains irrelevant content, it's crucial to guide the system to focus on the right information. The best way to achieve this is byincluding a clear instruction in the system prompt(option C).
* System Prompt as Guidance:The system prompt is an effective way to instruct the LLM to limit its focus to SnoPen AI-related content. By clearly specifying that the model should avoid answering questions unrelated to SnoPen AI, you add an additional layer of control that helps the model stay on- topic, even if irrelevant content is present in the dataset.
* Why This Approach Works:The prompt acts as a guiding principle for the model, narrowing its focus to specific domains. This prevents the model from generating answers based on irrelevant content, such as advertisements or news unrelated to SnoPen AI.
* Why Other Options Are Less Suitable:
* A (Keep All Articles): Retaining all content, including irrelevant materials, without any filtering makes the system prone to generating answers based on unwanted data.
* B (Include in the System Prompt about SnoPen AI): This option doesn't address irrelevant content directly, and without filtering, the model might still retrieve and use irrelevant data.
* D (Consolidating Documents into a Single Chunk): Grouping documents into a single chunk makes the retrieval process less efficient and won't help filter out irrelevant content effectively.
Therefore, instructing the system in the prompt not to answer questions unrelated to SnoPen AI (option C) is the best approach to ensure the system filters out irrelevant information.

질문 # 65
.....

Pass4Test에서 제공하는Databricks Databricks-Generative-AI-Engineer-Associate시험자료의 문제와 답은 실제시험의 문제와 답과 아주 비슷합니다, Databricks-Generative-AI-Engineer-Associate dumps를 데려가 주시면 기적을 안겨드릴게요, Databricks Databricks-Generative-AI-Engineer-Associate시험유형 주문번호와 불합격성적표를 메일로 보내오시면 바로 환불가능합니다, 저희 덤프만 공부하시면 시간도 절약하고 가격도 친근하며 시험준비로 인한 여러방면의 스트레스를 적게 받아Databricks인증 Databricks-Generative-AI-Engineer-Associate시험패스가 한결 쉬워집니다, Databricks-Generative-AI-Engineer-Associate덤프는 시험문제의 모든 범위가 포함되어 있어 시험적중율이 거의 100%에 달합니다, 시험신청하시는분들도 많아지고 또 많은 분들이 우리Pass4Test의Databricks Databricks-Generative-AI-Engineer-Associate자료로 시험을 패스했습니다.

누구보다 자신을 그렇게 만든 범인을 잡고 싶을 텐데.글쎄, 양 팔이 멋대로, 강일을 보듬어 안아버렸다.무서웠어요, Pass4Test에서 제공하는Databricks Databricks-Generative-AI-Engineer-Associate시험자료의 문제와 답은 실제시험의 문제와 답과 아주 비슷합니다.

시험준비에 가장 좋은 Databricks-Generative-AI-Engineer-Associate시험유형 최신버전 덤프샘플문제 다운로드

Databricks-Generative-AI-Engineer-Associate dumps를 데려가 주시면 기적을 안겨드릴게요, 주문번호와 불합격성적표를 메일로 보내오시면 바로 환불가능합니다, 저희 덤프만 공부하시면 시간도 절약하고 가격도 친근하며 시험준비로 인한 여러방면의 스트레스를 적게 받아Databricks인증 Databricks-Generative-AI-Engineer-Associate시험패스가 한결 쉬워집니다.

Databricks-Generative-AI-Engineer-Associate덤프는 시험문제의 모든 범위가 포함되어 있어 시험적중율이 거의 100%에 달합니다.

- Databricks-Generative-AI-Engineer-Associate시험대비 □ Databricks-Generative-AI-Engineer-Associate시험대비덤프 □ Databricks-Generative-AI-Engineer-Associate시험응시 □ 시험 자료를 무료로 다운로드하려면《 kr.fast2test.com 》을 통해⇛ Databricks-Generative-AI-Engineer-Associate ⇚를 검색하십시오Databricks-Generative-AI-Engineer-Associate시험대비
- Databricks-Generative-AI-Engineer-Associate시험유형 100% 시험패스 인증덤프문제 □ ☀ www.itdumpskr.com □☀□은□ Databricks-Generative-AI-Engineer-Associate □무료 다운로드를 받을 수 있는 최고의 사이트입니다Databricks-Generative-AI-Engineer-Associate인증덤프문제
- Databricks-Generative-AI-Engineer-Associate시험유형 시험 최신버전 자료 □ ▸ Databricks-Generative-AI-Engineer-Associate ◂를 무료로 다운로드하려면➥ www.passtip.net □웹사이트를 입력하세요Databricks-Generative-AI-Engineer-Associate높은 통과율 시험덤프공부
- 높은 통과율 Databricks-Generative-AI-Engineer-Associate시험유형 시험덤프 ↗ □ www.itdumpskr.com □에서 검색만 하면「 Databricks-Generative-AI-Engineer-Associate 」를 무료로 다운로드할 수 있습니다Databricks-Generative-AI-Engineer-Associate최신 덤프데모 다운
- Databricks-Generative-AI-Engineer-Associate최신 인증시험 덤프데모 □ Databricks-Generative-AI-Engineer-Associate높은 통과율 시험덤프공부 □ Databricks-Generative-AI-Engineer-Associate시험응시 □ 오픈 웹 사이트□ www.koreadumps.com □검색➠ Databricks-Generative-AI-Engineer-Associate □무료 다운로드Databricks-Generative-AI-Engineer-Associate높은 통과율 시험대비 덤프공부
- Databricks-Generative-AI-Engineer-Associate시험응시 □ Databricks-Generative-AI-Engineer-Associate최고품질 예상문제모음 □ Databricks-Generative-AI-Engineer-Associate높은 통과율 시험덤프공부 □ ➥ www.itdumpskr.com □은{ Databricks-Generative-AI-Engineer-Associate }무료 다운로드를 받을 수 있는 최고의 사이트입니다Databricks-Generative-AI-Engineer-Associate높은 통과율 시험대비 덤프공부
- Databricks Databricks-Generative-AI-Engineer-Associate 인증시험 □ ➽ www.koreadumps.com □을(를) 열고✔ Databricks-Generative-AI-Engineer-Associate □✔□를 입력하고 무료 다운로드를 받으십시오Databricks-Generative-AI-Engineer-Associate인기자격증 시험대비자료
- 시험대비 Databricks-Generative-AI-Engineer-Associate시험유형 최신 공부자료 □ □ www.itdumpskr.com □에서《 Databricks-Generative-AI-Engineer-Associate 》를 검색하고 무료 다운로드 받기Databricks-Generative-AI-Engineer-Associate높은 통과율 인기 시험자료
- Databricks-Generative-AI-Engineer-Associate인기시험자료 □ Databricks-Generative-AI-Engineer-Associate최고품질 인증시험자료 □ Databricks-Generative-AI-Engineer-Associate인기자격증 시험대비 공부자료 ◂➥ www.pass4test.net □을(를) 열고「 Databricks-Generative-AI-Engineer-Associate 」를 입력하고 무료 다운로드를 받으십시오Databricks-Generative-AI-Engineer-Associate높은 통과율 시험대비 덤프공부
- 높은 적중율을 자랑하는 Databricks-Generative-AI-Engineer-Associate시험유형 덤프샘플문제 □ □ www.itdumpskr.com □에서 검색만 하면" Databricks-Generative-AI-Engineer-Associate "를 무료로 다운로드할 수 있습니다Databricks-Generative-AI-Engineer-Associate시험대비
- Databricks-Generative-AI-Engineer-Associate인증시험덤프 □ Databricks-Generative-AI-Engineer-Associate인증덤프문제 □ Databricks-Generative-AI-Engineer-Associate인증시험덤프 □ 검색만 하면➠ www.itdumpskr.com □

에서➽ Databricks-Generative-AI-Engineer-Associate □무료 다운로드Databricks-Generative-AI-Engineer-Associate 시험응시

- tayawytv846317.blogsvirals.com, minajxjp878014.loginblogin.com, socialmarkz.com, honeyaqjj213604.wikifrontier.com, bookmarkangaroo.com, martinaaiyz230119.mywikiparty.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, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, harmonykbee217721.aboutyoublog.com, soulcreative.online, maciekpwa924225.wikigiogio.com, Disposable vapes

그 외, Pass4Test Databricks-Generative-AI-Engineer-Associate 시험 문제집 일부가 지금은 무료입니다: https://drive.google.com/open?id=1I8aytSzKMG5uh11TWlD8qC6qNKISvArH