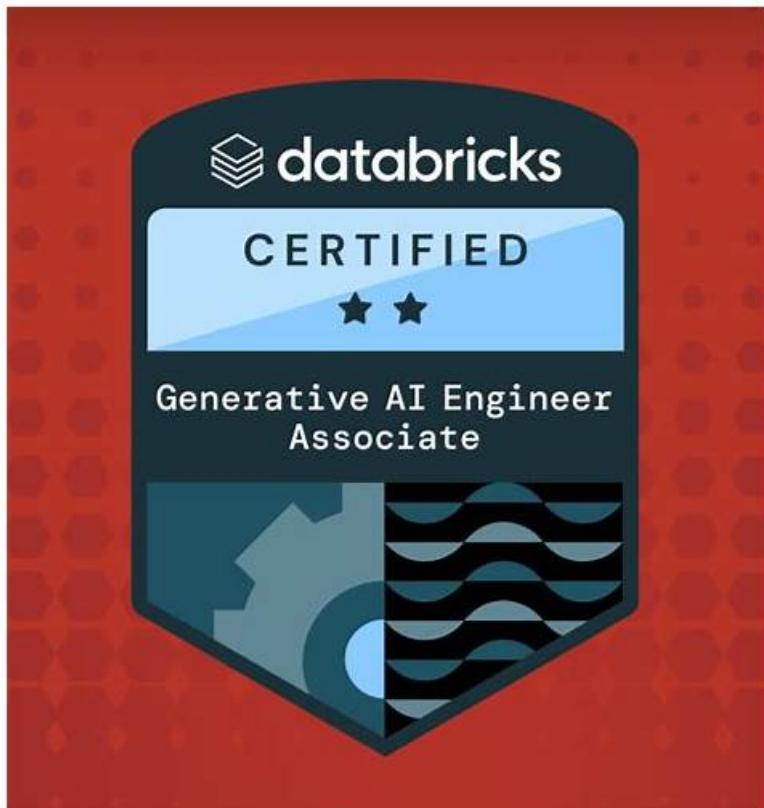


Databricks-Generative-AI-Engineer-Associate試験の準備方法 | 正確的なDatabricks-Generative-AI-Engineer-Associate日本語版対策ガイド試験 | ユニークなDatabricks Certified Generative AI Engineer Associate試験解説



2025年ShikenPASSの最新Databricks-Generative-AI-Engineer-Associate PDFダンプおよびDatabricks-Generative-AI-Engineer-Associate試験エンジンの無料共有: https://drive.google.com/open?id=1MA_YoYYMiLNJ3mlib_vzRqqdHOzgr0pS

21世紀には、{Examcode}認定は受験者の特定の能力を表すため、社会でますます認知されるようになります。ただし、{Examcode}認定を取得するには、Databricks-Generative-AI-Engineer-Associate試験の準備に多くの時間を費やす必要があります。Databricks-Generative-AI-Engineer-Associate模擬試験を購入すると、当社のウェブサイトはプロの技術を使用してすべてのユーザーのプライバシーを暗号化し、ハッカーの盗用を防ぎます。私たちは、ビジネスがお客様のために十分に考慮された場合にのみ継続できると考えているため、当社の評判を損なうような行為は一切行いません。Databricks-Generative-AI-Engineer-Associate試験問題に完全な信頼を寄せていただければ幸いです。失望することはありません。

Databricks Databricks-Generative-AI-Engineer-Associate 認定試験の出題範囲:

トピック	出題範囲
トピック 1	<ul style="list-style-type: none">ガバナンス: 試験を受けるジェネレーティブ AI エンジニアは、このトピックのマスキング手法、ガードレール手法、および法的ライセンス要件に関する知識を習得します。

トピック 2	<ul style="list-style-type: none"> アプリケーション開発: このトピックでは、Generative AI エンジニアは、データの抽出に必要なツール、Langchain 類似ツール、一般的な問題を特定するための応答の評価について学習します。さらに、このトピックには、LLM の応答の調整、LLM ガードレール、およびアプリケーションの属性に基づいた最適な LLM に関する質問が含まれています。
トピック 3	<ul style="list-style-type: none"> アプリケーションの組み立てとデプロイ: このトピックでは、Generative AI エンジニアは、pyfunc モードを使用してチェーンをコーディングする方法、langchain を使用してシンプルなチェーンをコーディングする方法、要件に従ってシンプルなチェーンをコーディングする方法を学びます。さらに、このトピックでは、RAG アプリケーションを作成するためには必要な基本要素に焦点を当てています。最後に、このトピックでは、MLflow を使用してモデルを Unity Catalog に登録する方法に関するサブトピックを取り上げます。
トピック 4	<ul style="list-style-type: none"> データ準備: Generative AI エンジニアは、特定のドキュメント構造とモデル制約のチャンギング戦略について説明します。このトピックでは、ソースドキュメント内の不要なコンテンツのフィルター処理にも重点を置いています。最後に、Generative AI エンジニアは、提供されたソースデータと形式からドキュメントコンテンツを抽出する方法についても学習します。
トピック 5	<ul style="list-style-type: none"> 評価と監視: このトピックでは、LLM の選択と主要なメトリックについて説明します。さらに、Generative AI エンジニアはモデルのパフォーマンスの評価について学習します。最後に、このトピックには推論ログと Databricks 機能の使用に関するサブトピックが含まれています。

>> Databricks-Generative-AI-Engineer-Associate日本語版対策ガイド <<

よくできたDatabricks-Generative-AI-Engineer-Associate日本語版対策ガイド & 資格試験のリーダー & 信頼できるDatabricks-Generative-AI-Engineer-Associate試験解説

Databricks-Generative-AI-Engineer-Associate勉強のトレントを購入すると、24時間オンラインの効率的なサービスを提供します。Databricks-Generative-AI-Engineer-Associate学習資料に関するご質問はいつでもお問い合わせいただけます。また、いつでもご連絡いただけます。もちろん、忙しくてオンラインで連絡する時間がない場合は、心配しないで、いつでもDatabricks-Generative-AI-Engineer-Associateガイド資料に関する問題をメールでお知らせください。カスタマーサービスからすぐにメールが届きます。一言で言えば、24時間オンラインの効率的なサービスは、すべての問題を解決して試験に合格するのに役立つと考えています。

Databricks Certified Generative AI Engineer Associate 認定 Databricks-Generative-AI-Engineer-Associate 試験問題 (Q12-Q17):

質問 #12

A Generative AI Engineer is developing a RAG system for their company to perform internal document Q&A for structured HR policies, but the answers returned are frequently incomplete and unstructured. It seems that the retriever is not returning all relevant context. The Generative AI Engineer has experimented with different embedding and response generating LLMs but that did not improve results.

Which TWO options could be used to improve the response quality?

Choose 2 answers

- A. Increase the document chunk size
- B. Fine tune the response generation model
- C. Add the section header as a prefix to chunks
- D. Split the document by sentence
- E. Use a larger embedding model

正解: A、C

解説:

The problem describes a Retrieval-Augmented Generation (RAG) system for HR policy Q&A where responses are incomplete and unstructured due to the retriever failing to return sufficient context. The engineer has already tried different embedding and response-generating LLMs without success, suggesting the issue lies in the retrieval process-specifically, how documents are chunked and indexed. Let's evaluate the options.

* Option A: Add the section header as a prefix to chunks

* Adding section headers provides additional context to each chunk, helping the retriever understand the chunk's relevance within the document structure (e.g., "Leave Policy: Annual Leave" vs. just "Annual Leave"). This can improve retrieval precision for structured HR policies.

* Databricks Reference: "Metadata, such as section headers, can be appended to chunks to enhance retrieval accuracy in RAG systems" ("Databricks Generative AI Cookbook," 2023).

* Option B: Increase the document chunk size

* Larger chunks include more context per retrieval, reducing the chance of missing relevant information split across smaller chunks. For structured HR policies, this can ensure entire sections or rules are retrieved together.

* Databricks Reference: "Increasing chunk size can improve context completeness, though it may trade off with retrieval specificity" ("Building LLM Applications with Databricks").

* Option C: Split the document by sentence

* Splitting by sentence creates very small chunks, which could exacerbate the problem by fragmenting context further. This is likely why the current system fails-it retrieves incomplete snippets rather than cohesive policy sections.

* Databricks Reference: No specific extract opposes this, but the emphasis on context completeness in RAG suggests smaller chunks worsen incomplete responses.

* Option D: Use a larger embedding model

* A larger embedding model might improve vector quality, but the question states that experimenting with different embedding models didn't help. This suggests the issue isn't embedding quality but rather chunking/retrieval strategy.

* Databricks Reference: Embedding models are critical, but not the focus when retrieval context is the bottleneck.

* Option E: Fine tune the response generation model

* Fine-tuning the LLM could improve response coherence, but if the retriever doesn't provide complete context, the LLM can't generate full answers. The root issue is retrieval, not generation.

* Databricks Reference: Fine-tuning is recommended for domain-specific generation, not retrieval fixes ("Generative AI Engineer Guide").

Conclusion: Options A and B address the retrieval issue directly by enhancing chunk context-either through metadata (A) or size (B)-aligning with Databricks' RAG optimization strategies. C would worsen the problem, while D and E don't target the root cause given prior experimentation.

質問 # 13

A Generative AI Engineer is building an LLM-based application that has an important transcription (speech-to-text) task. Speed is essential for the success of the application. Which open Generative AI models should be used?

- A. Llama-2-70b-chat-hf
- B. DBRX
- C. MPT-30B-Instruct
- D. **whisper-large-v3 (1.6B)**

正解: D

解説:

The task requires an open generative AI model for a transcription (speech-to-text) task where speed is essential. Let's assess the options based on their suitability for transcription and performance characteristics, referencing Databricks' approach to model selection.

* Option A: Llama-2-70b-chat-hf

* Llama-2 is a text-based LLM optimized for chat and text generation, not speech-to-text. It lacks transcription capabilities.

* Databricks Reference: "Llama models are designed for natural language generation, not audio processing" ("Databricks Model Catalog").

* Option B: MPT-30B-Instruct

* MPT-30B is another text-based LLM focused on instruction-following and text generation, not transcription. It's irrelevant for speech-to-text tasks.

* Databricks Reference: No specific mention, but MPT is categorized under text LLMs in Databricks' ecosystem, not audio models.

* Option C: DBRX

* DBRX, developed by Databricks, is a powerful text-based LLM for general-purpose generation.

It doesn't natively support speech-to-text and isn't optimized for transcription.

* Databricks Reference: "DBRX excels at text generation and reasoning tasks" ("Introducing DBRX," 2023)-no mention of audio capabilities.

* Option D: whisper-large-v3 (1.6B)

* Whisper, developed by OpenAI, is an open-source model specifically designed for speech-to-text transcription. The "large-v3" variant (1.6 billion parameters) balances accuracy and efficiency, with optimizations for speed via quantization or deployment on GPUs-key for the application's requirements.

* Databricks Reference: "For audio transcription, models like Whisper are recommended for their speed and accuracy" ("Generative AI Cookbook," 2023). Databricks supports Whisper integration in its MLflow or Lakehouse workflows.

Conclusion: Only D. whisper-large-v3 is a speech-to-text model, making it the sole suitable choice. Its design prioritizes transcription, and its efficiency (e.g., via optimized inference) meets the speed requirement, aligning with Databricks' model deployment best practices.

質問 # 14

A Generative AI Engineer is building a system that will answer questions on currently unfolding news topics.

As such, it pulls information from a variety of sources including articles and social media posts. They are concerned about toxic posts on social media causing toxic outputs from their system.

Which guardrail will limit toxic outputs?

- A. Implement rate limiting
- B. Log all LLM system responses and perform a batch toxicity analysis monthly.
- **C. Use only approved social media and news accounts to prevent unexpected toxic data from getting to the LLM.**
- D. Reduce the amount of context items the system will include in consideration for its response.

正解: C

解説:

The system answers questions on unfolding news topics using articles and social media, with a concern about toxic outputs from toxic inputs. A guardrail must limit toxicity in the LLM's responses. Let's evaluate the options.

* Option A: Use only approved social media and news accounts to prevent unexpected toxic data from getting to the LLM

* Curating input sources (e.g., verified accounts) reduces exposure to toxic content at the data ingestion stage, directly limiting toxic outputs. This is a proactive guardrail aligned with data quality control.

* Databricks Reference: "Control input data quality to mitigate unwanted LLM behavior, such as toxicity" ("Building LLM Applications with Databricks," 2023).

* Option B: Implement rate limiting

* Rate limiting controls request frequency, not content quality. It prevents overload but doesn't address toxicity in social media inputs or outputs.

* Databricks Reference: Rate limiting is for performance, not safety: "Use rate limits to manage compute load" ("Generative AI Cookbook").

* Option C: Reduce the amount of context items the system will include in consideration for its response

* Reducing context might limit exposure to some toxic items but risks losing relevant information, and it doesn't specifically target toxicity. It's an indirect, imprecise fix.

* Databricks Reference: Context reduction is for efficiency, not safety: "Adjust context size based on performance needs" ("Databricks Generative AI Engineer Guide").

* Option D: Log all LLM system responses and perform a batch toxicity analysis monthly

* Logging and analyzing responses is reactive, identifying toxicity after it occurs rather than preventing it. Monthly analysis doesn't limit real-time toxic outputs.

* Databricks Reference: Monitoring is for auditing, not prevention: "Log outputs for post-hoc analysis, but use input filters for safety" ("Building LLM-Powered Applications").

Conclusion: Option A is the most effective guardrail, proactively filtering toxic inputs from unverified sources, which aligns with Databricks' emphasis on data quality as a primary safety mechanism for LLM systems.

質問 # 15

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



databricks

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

正解: B

解説:

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.

質問 # 16

A Generative AI Engineer interfaces with an LLM with prompt/response behavior that has been trained on customer calls inquiring about product availability. The LLM is designed to output "In Stock" if the product is available or only the term "Out of Stock" if not. Which prompt will work to allow the engineer to respond to call classification labels correctly?

- A. Respond with "Out of Stock" if the customer asks for a product.
- B. You will be given a customer call transcript where the customer inquires about product availability. Respond with "In Stock" if the product is available or "Out of Stock" if not.
- C. Respond with "In Stock" if the customer asks for a product.
- D. You will be given a customer call transcript where the customer asks about product availability. The outputs are either "In Stock" or "Out of Stock". Format the output in JSON, for example: `{"call_id": "123", "label": "In Stock"}`.

正解: D

解説:

* Problem Context: The Generative AI Engineer needs a prompt that will enable an LLM trained on customer call transcripts to classify and respond correctly regarding product availability. The desired response should clearly indicate whether a product is "In Stock" or "Out of Stock," and it should be formatted in a way that is structured and easy to parse programmatically, such as JSON.

* Explanation of Options:

* Option A: Respond with "In Stock" if the customer asks for a product. This prompt is too generic and does not specify how to handle the case when a product is not available, nor does it provide a structured output format.

* Option B: This option is correctly formatted and explicit. It instructs the LLM to respond based on the availability mentioned in the customer call transcript and to format the response in JSON.

This structure allows for easy integration into systems that may need to process this information automatically, such as customer service dashboards or databases.

* Option C: Respond with "Out of Stock" if the customer asks for a product. Like option A, this prompt is also insufficient as it only covers the scenario where a product is unavailable and does not provide a structured output.

* Option D: While this prompt correctly specifies how to respond based on product availability, it lacks the structured output format, making it less suitable for systems that require formatted data for further processing.

Given the requirements for clear, programmatically usable outputs, Option B is the optimal choice because it provides precise instructions on how to respond and includes a JSON format example for structuring the output, which is ideal for automated systems or further data handling.

質問 #17

.....

今日では、柔軟な学習方法が電子製品の開発でますます一般的になっています。最新の技術は、同様に、我々はこの分野で最も主導的な地位にあることから、当社DatabricksのDatabricks-Generative-AI-Engineer-Associate実際の試験に適用されています。また、あなたは私たちのDatabricks-Generative-AI-Engineer-Associate練習材料の3つのバージョンが存在するために多様な選択肢があります。同時に、Databricks-Generative-AI-Engineer-Associate試験に合格し、Databricks-Generative-AI-Engineer-Associate学習教材の有効性と正確性について希望のDatabricks-Generative-AI-Engineer-Associate認定を取得する必要があります。

Databricks-Generative-AI-Engineer-Associate試験解説: <https://www.shikenpass.com/Databricks-Generative-AI-Engineer-Associate-shiken.html>

- Databricks-Generative-AI-Engineer-Associate最新テスト □ Databricks-Generative-AI-Engineer-Associate資格専門知識 □ Databricks-Generative-AI-Engineer-Associate最新対策問題 □ ▶ www.passtest.jp ▶ にて限定無料の＊ Databricks-Generative-AI-Engineer-Associate □＊□問題集をダウンロードせよ Databricks-Generative-AI-Engineer-Associate参考書内容
- 権威のあるDatabricks Databricks-Generative-AI-Engineer-Associate | 効率的なDatabricks-Generative-AI-Engineer-Associate日本語版対策ガイド試験 | 試験の準備方法Databricks Certified Generative AI Engineer Associate試験解説 □ ▶ www.goshiken.com ▶ で ✓ Databricks-Generative-AI-Engineer-Associate □✓□を検索して、無料でダウンロードしてください Databricks-Generative-AI-Engineer-Associate日本語版参考資料
- 初段Databricks-Generative-AI-Engineer-Associate日本語版対策ガイド - 資格試験のリーダー - 完璧なDatabricks-Generative-AI-Engineer-Associate試験解説 □＊ www.jpshiken.com □＊□サイトにて最新＊ Databricks-Generative-AI-Engineer-Associate □＊□問題集をダウンロード Databricks-Generative-AI-Engineer-Associate受験対策
- 信頼できる Databricks-Generative-AI-Engineer-Associate日本語版対策ガイド - 合格スムーズDatabricks-Generative-AI-Engineer-Associate試験解説 | 正確的なDatabricks-Generative-AI-Engineer-Associate関連試験 Databricks Certified Generative AI Engineer Associate □ 今すぐ ▶ www.goshiken.com □□□で □ Databricks-Generative-AI-Engineer-Associate □を検索して、無料でダウンロードしてください Databricks-Generative-AI-Engineer-Associate資格認定試験
- Databricks-Generative-AI-Engineer-Associate日本語版復習資料 □ Databricks-Generative-AI-Engineer-Associate日本語版参考資料 □ Databricks-Generative-AI-Engineer-Associate最新対策問題 □ 今すぐ □ www.jptestking.com □で 「 Databricks-Generative-AI-Engineer-Associate 」を検索し、無料でダウンロードしてください Databricks-Generative-AI-Engineer-Associate問題と解答
- 効率的なDatabricks-Generative-AI-Engineer-Associate日本語版対策ガイド - 合格スムーズDatabricks-Generative-AI-Engineer-Associate試験解説 | 素晴らしい Databricks-Generative-AI-Engineer-Associate関連試験 □ ウェブサイト ▶ www.goshiken.com □□□から ▶ Databricks-Generative-AI-Engineer-Associate □を開いて検索し、無料でダウンロードしてください Databricks-Generative-AI-Engineer-Associate日本語版復習資料
- 信頼できる Databricks-Generative-AI-Engineer-Associate日本語版対策ガイド - 合格スムーズDatabricks-Generative-AI-Engineer-Associate試験解説 | 正確的なDatabricks-Generative-AI-Engineer-Associate関連試験 Databricks Certified Generative AI Engineer Associate □ 《 www.passtest.jp 》で 「 Databricks-Generative-AI-Engineer-Associate 」を検索し、無料でダウンロードしてください Databricks-Generative-AI-Engineer-Associate日本語版復習資料
- 試験の準備方法-権威のあるDatabricks-Generative-AI-Engineer-Associate日本語版対策ガイド試験-信頼できる Databricks-Generative-AI-Engineer-Associate試験解説 □ ▶ www.goshiken.com □に移動し、「 Databricks-Generative-AI-Engineer-Associate 」を検索して、無料でダウンロード可能な試験資料を探します Databricks-

Generative-AI-Engineer-Associate受験対策

- Databricks-Generative-AI-Engineer-Associate受験内容 □ Databricks-Generative-AI-Engineer-Associate参考書内容 □ Databricks-Generative-AI-Engineer-Associate日本語版復習資料 □ 【 jp.fast2test.com 】で使える無料オンライン版 ➔ Databricks-Generative-AI-Engineer-Associate □□□ の試験問題Databricks-Generative-AI-Engineer-Associate資格準備
- 効率的なDatabricks-Generative-AI-Engineer-Associate日本語版対策ガイド - 合格スムーズDatabricks-Generative-AI-Engineer-Associate試験解説 | 素晴らしいDatabricks-Generative-AI-Engineer-Associate関連試験 □ { www.goshiken.com }で使える無料オンライン版《 Databricks-Generative-AI-Engineer-Associate 》の試験問題Databricks-Generative-AI-Engineer-Associate日本語版参考資料
- Databricks-Generative-AI-Engineer-Associate関連資格試験対応 □ Databricks-Generative-AI-Engineer-Associate日本語版復習資料 □ Databricks-Generative-AI-Engineer-Associate復習過去問 □ ➤ www.shikenpass.com □に移動し、□ Databricks-Generative-AI-Engineer-Associate □を検索して、無料でダウンロード可能な試験資料を探しますDatabricks-Generative-AI-Engineer-Associate最新対策問題
- elearning.eauqardho.edu.so, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, myportal.utt.edu.tt, rdguitar.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, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, Disposable vapes

さらに、ShikenPASS Databricks-Generative-AI-Engineer-Associateダンプの一部が現在無料で提供されています：https://drive.google.com/open?id=1MA_YoYYMiLNJ3m1b_vzRqqdHOzgr0pS