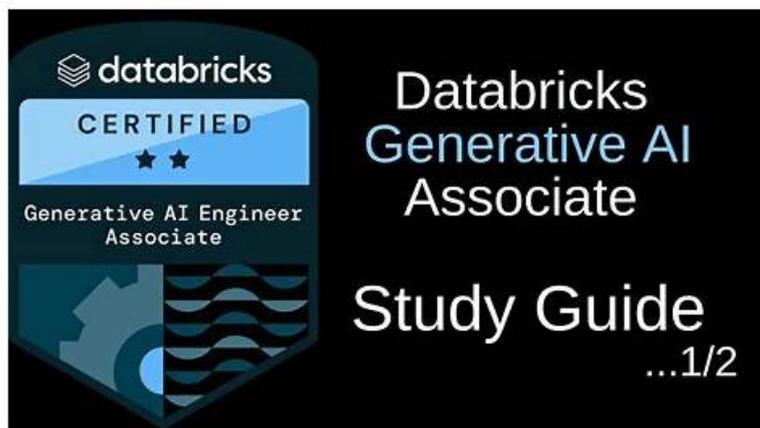


Databricks Databricks-Generative-AI-Engineer-Associate 認定テキスト & Databricks-Generative-AI-Engineer-Associate的中合格問題集



2026年Jpexam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=18q62aRQV-n6Z5p7TUa5m3-Yek5J2Vnbs>

Databricks-Generative-AI-Engineer-Associate試験に問題がある場合は、無料のデモを検討してください。弊社の最新のDatabricks-Generative-AI-Engineer-Associate試験トレントは、この業界では完璧な模範であり、さまざまな程度の試験受験者向けの明確なコンテンツに満ちています。最新のDatabricks-Generative-AI-Engineer-Associate試験トレントの結果は驚くほど驚くべきもので、試験受験者の98%以上が目標を無事に達成しました。また、Databricks-Generative-AI-Engineer-Associateテストダンプにより、あらゆる種類の教材の精度が非常に高いことが保証されました。

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 3	<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キ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 4	<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的中合格問題集 & Databricks-Generative-AI-Engineer-Associate難易度受験料

Jpexamは初めて試験を受けるあなたが一回で試験に合格して、認証資格を取ることを保証します。Jpexamが提供して差し上げたのは高品質のDatabricksのDatabricks-Generative-AI-Engineer-Associate「Databricks Certified Generative AI Engineer Associate」模擬問題集で、あなたがステップバイステップで試験に準備する手順を指導しています。JpexamのDatabricksのDatabricks-Generative-AI-Engineer-Associate試験問題集は絶対あなたに成功をもたらすことを保証します。JpexamのDatabricksのDatabricks-Generative-AI-Engineer-Associate認定試験に準備するために色々な方法があります。

Databricks Certified Generative AI Engineer Associate 認定 Databricks-Generative-AI-Engineer-Associate 試験問題 (Q37-Q42):

質問 # 37

After changing the response generating LLM in a RAG pipeline from GPT-4 to a model with a shorter context length that the company self-hosts, the Generative AI Engineer is getting the following error:

What TWO solutions should the Generative AI Engineer implement without changing the response generating model? (Choose two.)

- A. Reduce the maximum output tokens of the new model
- B. Use a smaller embedding model to generate
- C. Retrain the response generating model using ALiBi
- **D. Reduce the number of records retrieved from the vector database**
- **E. Decrease the chunk size of embedded documents**

正解: D、E

解説:

Problem Context: After switching to a model with a shorter context length, the error message indicating that the prompt token count has exceeded the limit suggests that the input to the model is too large.

Explanation of Options:

Option A: Use a smaller embedding model to generate - This wouldn't necessarily address the issue of prompt size exceeding the model's token limit.

Option B: Reduce the maximum output tokens of the new model - This option affects the output length, not the size of the input being too large.

Option C: Decrease the chunk size of embedded documents - This would help reduce the size of each document chunk fed into the model, ensuring that the input remains within the model's context length limitations.

Option D: Reduce the number of records retrieved from the vector database - By retrieving fewer records, the total input size to the model can be managed more effectively, keeping it within the allowable token limits.

Option E: Retrain the response generating model using ALiBi - Retraining the model is contrary to the stipulation not to change the response generating model.

Options C and D are the most effective solutions to manage the model's shorter context length without changing the model itself, by adjusting the input size both in terms of individual document size and total documents retrieved.

質問 # 38

A Generative AI Engineer is using an LLM to classify species of edible mushrooms based on text descriptions of certain features.

The model is returning accurate responses in testing and the Generative AI Engineer is confident they have the correct list of possible labels, but the output frequently contains additional reasoning in the answer when the Generative AI Engineer only wants to return the label with no additional text.

Which action should they take to elicit the desired behavior from this LLM?

- A. Use zero shot chain-of-thought prompting to prevent a verbose output format
- **B. Use a system prompt to instruct the model to be succinct in its answer**
- C. Use few shot prompting to instruct the model on expected output format
- D. Use zero shot prompting to instruct the model on expected output format

正解: B

解説:

The LLM classifies mushroom species accurately but includes unwanted reasoning text, and the engineer wants only the label. Let's assess how to control output format effectively.

Option A: Use few shot prompting to instruct the model on expected output format Few-shot prompting provides examples (e.g., input: description, output: label). It can work but requires crafting multiple examples, which is effort-intensive and less direct than a clear instruction.

Databricks Reference: "Few-shot prompting guides LLMs via examples, effective for format control but requires careful design" ("Generative AI Cookbook").

Option B: Use zero shot prompting to instruct the model on expected output format Zero-shot prompting relies on a single instruction (e.g., "Return only the label") without examples. It's simpler than few-shot but may not consistently enforce succinctness if the LLM's default behavior is verbose.

Databricks Reference: "Zero-shot prompting can specify output but may lack precision without examples" ("Building LLM Applications with Databricks").

Option C: Use zero shot chain-of-thought prompting to prevent a verbose output format Chain-of-Thought (CoT) encourages step-by-step reasoning, which increases verbosity-opposite to the desired outcome. This contradicts the goal of label-only output.

Databricks Reference: "CoT prompting enhances reasoning but often results in detailed responses" ("Databricks Generative AI Engineer Guide").

Option D: Use a system prompt to instruct the model to be succinct in its answer A system prompt (e.g., "Respond with only the species label, no additional text") sets a global instruction for the LLM's behavior. It's direct, reusable, and effective for controlling output style across queries.

Databricks Reference: "System prompts define LLM behavior consistently, ideal for enforcing concise outputs" ("Generative AI Cookbook," 2023).

Conclusion: Option D is the most effective and straightforward action, using a system prompt to enforce succinct, label-only responses, aligning with Databricks' best practices for output control.

質問 # 39

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

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

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

正解: C

解説:

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, hence Option B.

質問 # 40

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. Throttle the incoming batch of requests manually to avoid rate limiting issues
- C. Switch to using External Models instead
- D. Change to a model with a fewer number of parameters in order to reduce hardware constraint issues

正解: A

解説:

* 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.

質問 # 41

A team wants to serve a code generation model as an assistant for their software developers. It should support multiple programming languages. Quality is the primary objective.

Which of the Databricks Foundation Model APIs, or models available in the Marketplace, would be the best fit?

- A. CodeLlama-34B
- B. MPT-7b
- C. Llama2-70b
- D. BGE-large

正解: A

解説:

For a code generation model that supports multiple programming languages and where quality is the primary objective, CodeLlama-34B is the most suitable choice. Here's the reasoning:

Specialization in Code Generation:

CodeLlama-34B is specifically designed for code generation tasks. This model has been trained with a focus on understanding and generating code, which makes it particularly adept at handling various programming languages and coding contexts.

Capacity and Performance:

The "34B" indicates a model size of 34 billion parameters, suggesting a high capacity for handling complex tasks and generating high-quality outputs. The large model size typically correlates with better understanding and generation capabilities in diverse scenarios.

Suitability for Development Teams:

Given that the model is optimized for code, it will be able to assist software developers more effectively than general-purpose models. It understands coding syntax, semantics, and the nuances of different programming languages.

Why Other Options Are Less Suitable:

A (Llama2-70b): While also a large model, it's more general-purpose and may not be as fine-tuned for code generation as CodeLlama.

B (BGE-large): This model may not specifically focus on code generation.

C (MPT-7b): Smaller than CodeLlama-34B and likely less capable in handling complex code generation tasks at high quality. Therefore, for a high-quality, multi-language code generation application, CodeLlama-34B (option D) is the best fit.

質問 # 42

.....

あなたはこのような人々の一人ですか。さまざまな資料とトレーニング授業を前にして、どれを選ぶか本当に困っているのです。もしそうだったら、これ以上困ることはないです。Jpexamはあなたにとって最も正確な選

扱いですから。我々はあなたに試験問題と解答に含まれている全面的な試験資料を提供することができます。Jpexamの解答は最も正確な解釈ですから、あなたがより良い知識を身につけることに助けになります。Jpexamを利用したら、DatabricksのDatabricks-Generative-AI-Engineer-Associate認定試験に受かることを信じています。それも我々が全てのお客様に対する約束です。

Databricks-Generative-AI-Engineer-Associate的中合格問題集: https://www.jpexam.com/Databricks-Generative-AI-Engineer-Associate_exam.html

- 正確なDatabricks-Generative-AI-Engineer-Associate認定テキスト | 素晴らしい合格率のDatabricks-Generative-AI-Engineer-Associate Exam | 無料Databricks-Generative-AI-Engineer-Associate: Databricks Certified Generative AI Engineer Associate □ □ www.jpexam.com □ は、☀ Databricks-Generative-AI-Engineer-Associate □ ☀ □ を無料でダウンロードするのに最適なサイトですDatabricks-Generative-AI-Engineer-Associate模擬体験
- Databricks-Generative-AI-Engineer-Associate試験の準備方法 | 100% 合格率のDatabricks-Generative-AI-Engineer-Associate認定テキスト試験 | 検証するDatabricks Certified Generative AI Engineer Associate的中合格問題集 (M) ➤ 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学習範囲
- Databricks-Generative-AI-Engineer-Associate専門知識内容 □ Databricks-Generative-AI-Engineer-Associate合格体験談 □ Databricks-Generative-AI-Engineer-Associate資格難易度 □ □ www.jpexam.com □ 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日本語版試験勉強法
- 実際のDatabricks-Generative-AI-Engineer-Associate認定テキスト - 合格スムーズDatabricks-Generative-AI-Engineer-Associate的中合格問題集 | 一生懸命にDatabricks-Generative-AI-Engineer-Associate難易度受験料Databricks Certified Generative AI Engineer Associate □ ⇒ www.goshiken.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模擬体験
- 正確なDatabricks-Generative-AI-Engineer-Associate認定テキスト | 素晴らしい合格率のDatabricks-Generative-AI-Engineer-Associate Exam | 無料Databricks-Generative-AI-Engineer-Associate: Databricks Certified Generative AI Engineer Associate □ (www.passtest.jp) 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テキスト
- Databricks-Generative-AI-Engineer-Associate資格難易度 □ Databricks-Generative-AI-Engineer-Associate問題無料 □ Databricks-Generative-AI-Engineer-Associate出題範囲 □ 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対策学習
- 試験の準備方法-一番優秀なDatabricks-Generative-AI-Engineer-Associate認定テキスト試験-最高のDatabricks-Generative-AI-Engineer-Associate的中合格問題集 □ サイト ➡ www.passtest.jp □ □ □ で { Databricks-Generative-AI-Engineer-Associate } 問題集をダウンロードDatabricks-Generative-AI-Engineer-Associate日本語pdf問題
- 実際のDatabricks-Generative-AI-Engineer-Associate認定テキスト - 合格スムーズDatabricks-Generative-AI-Engineer-Associate的中合格問題集 | 一生懸命にDatabricks-Generative-AI-Engineer-Associate難易度受験料Databricks Certified Generative AI Engineer Associate □ ➤ 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合格内容
- 試験Databricks-Generative-AI-Engineer-Associate認定テキスト - 検証するDatabricks-Generative-AI-Engineer-Associate的中合格問題集 | 大人気Databricks-Generative-AI-Engineer-Associate難易度受験料 □ 最新 ➡ Databricks-Generative-AI-Engineer-Associate □ 問題集ファイルは ✓ www.passtest.jp □ ✓ □ にて検索Databricks-Generative-AI-Engineer-Associateテキスト
- 正確なDatabricks-Generative-AI-Engineer-Associate認定テキスト | 素晴らしい合格率のDatabricks-Generative-AI-Engineer-Associate Exam | 無料Databricks-Generative-AI-Engineer-Associate: 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問題無料
- 検証するDatabricks-Generative-AI-Engineer-Associate認定テキスト - 合格スムーズDatabricks-Generative-AI-Engineer-Associate的中合格問題集 | 完璧なDatabricks-Generative-AI-Engineer-Associate難易度受験料 □ 検索するだけで ➡ www.it-passports.com □ から ✓ Databricks-Generative-AI-Engineer-Associate □ ✓ □ を無料でダウンロードDatabricks-Generative-AI-Engineer-Associateテキスト
- bd.enrollbusiness.com, gettr.com, www.stes.tyc.edu.tw, 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, pixabay.com, myportal.utt.edu.tt, www.stes.tyc.edu.tw, wibki.com, Disposable vapes

2026年Jpexam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=18q62aRQV-n6Z5p7TUa5m3-Yek5J2Vnbs>