

よくできたDatabricks Databricks-Generative-AI-Engineer-Associateオンライン試験は主要材料 & 正確的なDatabricks-Generative-AI-Engineer-Associateテスト難易度



ちなみに、It-Passports Databricks-Generative-AI-Engineer-Associateの一部をクラウドストレージからダウンロードできます: https://drive.google.com/open?id=1Ocahtj4oYBW_O0hsbTI5XkK4BykWZPyg

お客様は弊社のDatabricks-Generative-AI-Engineer-Associate問題集を購入するかどうかと判断する前に、我が社は無料で提供するサンプルをダウンロードして試すことができます。それで、不必要な損失を避けられます。お客様はDatabricks-Generative-AI-Engineer-Associate問題集を購入してから、勉強中で何の質問があると、行き届いたサービスを得られています。お客様はDatabricks-Generative-AI-Engineer-Associate資格認証試験に失敗したら、弊社は全額返金できます。その他、Databricks-Generative-AI-Engineer-Associate問題集の更新版を無料で提供します。

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

Certification Vendor:	Databricks
Exam Name:	Databricks Certified Generative AI Engineer Associate
Exam Number:	Generative AI Engineer Associate
Exam Format:	Multiple Select, Multiple Choice
Certificate Validity Period:	2 years
Exam Price:	\$200 USD
Real Exam Qty:	45
Exam Duration:	90 minutes
Related Certifications:	Databricks Certified Machine Learning Associate Databricks Certified Data Engineer Associate
Available Languages:	Japanese, Korean, English, Portuguese (Brazil)
Passing Score:	Scaled Score (approx. 70%)

Sample Questions:	Databricks Databricks-Generative-AI-Engineer-Associate Sample Questions
Exam Way:	Online (Proctored) or Test Center
Pre Condition:	None (Recommended: 6+ months hands-on experience)
Official Syllabus URL:	https://www.databricks.com/learn/certification/genai-engineer-associate

>> Databricks-Generative-AI-Engineer-Associate オンライン試験 <<

Databricks Databricks-Generative-AI-Engineer-Associate テスト難易度、 Databricks-Generative-AI-Engineer-Associate テストトレーニング

当社It-Passportsのすべての専門家および教授の唯一の目標は、すべての人々に最適で適切なDatabricks-Generative-AI-Engineer-Associate学習教材を設計することです。多くの顧客のさまざまな要求に応じて、彼らはすべての顧客向けに3つの異なるバージョンのDatabricks-Generative-AI-Engineer-Associate認定試験ガイド資料を設計しました：PDF、ソフト、およびAPPバージョン。弊社のDatabricks-Generative-AI-Engineer-Associate試験問題を使用するすべての人がDatabricks-Generative-AI-Engineer-Associate試験に合格し、関連する認定資格を取得できることを心から願っています。そして、Databricks-Generative-AI-Engineer-Associate試験問題の合格率は98%以上です。

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	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 2	- 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	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	アプリケーションの設計: このトピックでは、特定の形式の応答を引き出すプロンプトの設計に焦点を当てています。また、特定のビジネス要件を達成するためのモデルタスクの選択にも焦点を当てています。最後に、このトピックでは、必要なモデル入力と出力のチェーンコンポーネントについて説明します。

Databricks Certified Generative AI Engineer Associate 認定 Databricks-Generative-AI-Engineer-Associate 試験問題 (Q39-Q44):

質問 # 39

A Generative AI Engineer is deploying a customer-facing, fine-tuned LLM on their public website. Given the large investment the company put into fine-tuning this model, and the proprietary nature of the tuning data, they are concerned about model inversion attacks. Which of the following Databricks AI Security Framework (DASF) risk mitigation strategies are most relevant to this use case?

- A. Use secure model features with Databricks Feature Store
- B. Apply attribute-based access controls (ABAC) to limit unauthorized access
- C. Leverage Databricks access control lists (ACLs) to configure permissions for accessing models
- D. Implement AI guardrails to allow users to configure and enforce compliance

正解: D

解説:

Model inversion attacks occur when an attacker uses the model's outputs to reconstruct the sensitive training data used during the fine-tuning process. To mitigate this in a public-facing application, implementing AI Guardrails is the most relevant strategy. Guardrails act as a programmable "filter" between the LLM and the end-user. They can be configured to detect if a model's response contains patterns that look like proprietary training data or PII (Personally Identifiable Information). While ACLs (B) and ABAC (D) protect the model's infrastructure (who can invoke the API), they do not inspect the content of the output, which is where the inversion attack actually manifests. Databricks provides integrated guardrail capabilities (via Mosaic AI Gateway) specifically to enforce compliance and prevent the leakage of sensitive internal knowledge that may have been baked into the model weights during fine-tuning.

質問 # 40

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

```
from databricks.vector_search.client import VectorSearchClient

vsc = VectorSearchClient()

vsc.create_endpoint(
    name="vector_search_test",
    endpoint_type="STANDARD"
)
```

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

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

正解: D

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 # 41

A Generative AI Engineer is ready to deploy an LLM application written using Foundation Model APIs. They want to follow security best practices for production scenarios Which authentication method should they choose?

- A. Use an access token belonging to service principals
- B. Use a frequently rotated access token belonging to either a workspace user or a service principal
- C. Use OAuth machine-to-machine authentication
- D. Use an access token belonging to any workspace user

正解: A

解説:

The task is to deploy an LLM application using Foundation Model APIs in a production environment while adhering to security best practices. Authentication is critical for securing access to Databricks resources, such as the Foundation Model API. Let's evaluate the options based on Databricks' security guidelines for production scenarios.

Option A: Use an access token belonging to service principals

Service principals are non-human identities designed for automated workflows and applications in Databricks. Using an access token tied to a service principal ensures that the authentication is scoped to the application, follows least-privilege principles (via role-based access control), and avoids reliance on individual user credentials. This is a security best practice for production deployments.

Databricks Reference: "For production applications, use service principals with access tokens to authenticate securely, avoiding user-specific credentials" ("Databricks Security Best Practices," 2023). Additionally, the "Foundation Model API Documentation" states: "Service principal tokens are recommended for programmatic access to Foundation Model APIs." Option B: Use a frequently rotated access token belonging to either a workspace user or a service principal Frequent rotation enhances security by limiting token exposure, but tying the token to a workspace user introduces risks (e.g., user account changes, broader permissions). Including both user and service principal options dilutes the focus on application-specific security, making this less ideal than a service-principal-only approach. It also adds operational overhead without clear benefits over Option A.

Databricks Reference: "While token rotation is a good practice, service principals are preferred over user accounts for application authentication" ("Managing Tokens in Databricks," 2023).

Option C: Use OAuth machine-to-machine authentication

OAuth M2M (e.g., client credentials flow) is a secure method for application-to-service communication, often using service principals under the hood. However, Databricks' Foundation Model API primarily supports personal access tokens (PATs) or service principal tokens over full OAuth flows for simplicity in production setups. OAuth M2M adds complexity (e.g., managing refresh tokens) without a clear advantage in this context.

Databricks Reference: "OAuth is supported in Databricks, but service principal tokens are simpler and sufficient for most API-based workloads" ("Databricks Authentication Guide," 2023).

Option D: Use an access token belonging to any workspace user

Using a user's access token ties the application to an individual's identity, violating security best practices. It risks exposure if the user leaves, changes roles, or has overly broad permissions, and it's not scalable or auditable for production.

Databricks Reference: "Avoid using personal user tokens for production applications due to security and governance concerns" ("Databricks Security Best Practices," 2023).

Conclusion: Option A is the best choice, as it uses a service principal's access token, aligning with Databricks' security best practices for production LLM applications. It ensures secure, application-specific authentication with minimal complexity, as explicitly recommended for Foundation Model API deployments.

質問 # 42

A Generative AI Engineer has created a RAG application to look up answers to questions about a series of fantasy novels that are being asked on the author's web forum. The fantasy novel texts are chunked and embedded into a vector store with metadata (page number, chapter number, book title), retrieved with the user's query, and provided to an LLM for response generation. The Generative AI Engineer used their intuition to pick the chunking strategy and associated configurations but now wants to more methodically choose the best values.

Which TWO strategies should the Generative AI Engineer take to optimize their chunking strategy and parameters? (Choose two.)

- A. Pass known questions and best answers to an LLM and instruct the LLM to provide the best token count. Use a summary statistic (mean, median, etc.) of the best token counts to choose chunk size.
- B. Create an LLM-as-a-judge metric to evaluate how well previous questions are answered by the most appropriate chunk. Optimize the chunking parameters based upon the values of the metric.
- C. Add a classifier for user queries that predicts which book will best contain the answer. Use this to filter retrieval.
- D. Change embedding models and compare performance.
- E. Choose an appropriate evaluation metric (such as recall or NDCG) and experiment with changes in the chunking strategy, such as splitting chunks by paragraphs or chapters. Choose the strategy that gives the best performance metric.

正解: B、E

質問 # 43

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. Wrap the LLM's prediction function into a Flask application and serve using Gunicorn
- 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. 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

正解: C

解説:

* Problem Context: The goal is to deploy a trained LLM on Databricks in the simplest and most integrated manner.

* Explanation of Options:

* Option A: This method involves unnecessary steps like logging the model as a pickle object, which is not the most efficient path in a Databricks environment.

* Option B: Logging the model with MLflow during training and then using MLflow's API to register and start serving the model is straightforward and leverages Databricks' built-in functionalities for seamless model deployment.

* Option C: Building and running a Docker container is a complex and less integrated approach within the Databricks ecosystem.

* Option D: Using Flask and Gunicorn is a more manual approach and less integrated compared to the native capabilities of Databricks and MLflow.

Option B provides the most straightforward and efficient process, utilizing Databricks' ecosystem to its full advantage for deploying models.

質問 # 44

.....

Databricks-Generative-AI-Engineer-Associate テスト難易度: <https://www.it-passports.com/Databricks-Generative-AI-Engineer-Associate.html>

- Databricks-Generative-AI-Engineer-Associate 日本語認定対策 □ Databricks-Generative-AI-Engineer-Associate 日本語受験教科書 □ Databricks-Generative-AI-Engineer-Associate 実際試験 □ ☀ www.shikenpass.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 最新受験攻略
- Databricks-Generative-AI-Engineer-Associate 最新受験攻略 □ Databricks-Generative-AI-Engineer-Associate 実際試験 □ Databricks-Generative-AI-Engineer-Associate 的中問題集 □ [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 日本語版問題集
- 公認された Databricks-Generative-AI-Engineer-Associate オンライン試験 - 資格試験のリーダー - 高品質 Databricks-Generative-AI-Engineer-Associate: Databricks Certified Generative AI Engineer Associate □ ➡ www.passtest.jp □ 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 学習範囲
- Databricks-Generative-AI-Engineer-Associate 認証 pdf 資料 □ Databricks-Generative-AI-Engineer-Associate 最新受験攻略 □ Databricks-Generative-AI-Engineer-Associate 受験内容 □ ➡ www.goshiken.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 認証 pdf 資料
- 信頼できる Databricks-Generative-AI-Engineer-Associate オンライン試験と便利な Databricks-Generative-AI-Engineer-Associate テスト難易度 □ □ Databricks-Generative-AI-Engineer-Associate □ の試験問題は《 www.it-passports.com 》で無料配信中 Databricks-Generative-AI-Engineer-Associate 問題サンプル
- Databricks-Generative-AI-Engineer-Associate 復習内容 □ Databricks-Generative-AI-Engineer-Associate 問題サンプル □ Databricks-Generative-AI-Engineer-Associate 模試エンジン □ 「 www.goshiken.com 」には無料の“Databricks-Generative-AI-Engineer-Associate”問題集があります Databricks-Generative-AI-Engineer-Associate 模試エンジン
- Databricks-Generative-AI-Engineer-Associate 模試エンジン □ Databricks-Generative-AI-Engineer-Associate 日本語認定対策 □ Databricks-Generative-AI-Engineer-Associate 資格受験料 □ “ www.japancert.com ”の無料ダウンロード

ド【 Databricks-Generative-AI-Engineer-Associate 】ページが開きますDatabricks-Generative-AI-Engineer-Associate 対策学習

- 公認されたDatabricks-Generative-AI-Engineer-Associateオンライン試験 - 資格試験のリーダー - 高品質
Databricks-Generative-AI-Engineer-Associate: Databricks Certified Generative AI Engineer Associate □ 今すぐ➡
www.goshiken.com □□□を開き、□ 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
- Databricks-Generative-AI-Engineer-Associate合格内容 □ Databricks-Generative-AI-Engineer-Associate復習内容 □
Databricks-Generative-AI-Engineer-Associate関連問題資料 □ { www.mogixam.com }に移動し、➡ 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的中問題集
- 信頼的なDatabricks-Generative-AI-Engineer-Associateオンライン試験 - 合格スムーズDatabricks-Generative-AI-
Engineer-Associateテスト難易度 | 効果的なDatabricks-Generative-AI-Engineer-Associateテストトレーニング □
➤ www.goshiken.com □から ➡ Databricks-Generative-AI-Engineer-Associate □□□を検索して、試験資料を無料で
ダウンロードしてくださいDatabricks-Generative-AI-Engineer-Associate認証pdf資料
- Databricks Databricks-Generative-AI-Engineer-Associate認定試験に合格できる参考書を紹介 □ 《
www.mogixam.com》には無料の{ Databricks-Generative-AI-Engineer-Associate }問題集がありますDatabricks-
Generative-AI-Engineer-Associate資格受験料
- aprenderfotografia.online, 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, 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, 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, myportal.utt.edu.tt, www.stes.tyc.edu.tw, 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, myportal.utt.edu.tt,
myportal.utt.edu.tt, myportal.utt.edu.tt, Disposable vapes

無料でクラウドストレージから最新のIt-Passports Databricks-Generative-AI-Engineer-Associate PDFダンプをダウンロードする: https://drive.google.com/open?id=1Ocahtj4oYBW_O0hsbTI5XkK4BykWZPyg