# 覆蓋全面的Databricks Databricks-Generative-AI-Engineer-Associate考試大綱是行業領先材料和經過驗證的Databricks-Generative-AI-Engineer-Associate：Databricks Certified Generative AI Engineer Associate


P.S. VCESoft在Google Drive上分享了免費的2026 Databricks Databricks-Generative-AI-Engineer-Associate考試題庫：https://drive.google.com/open?id=1PGXX-jyac8z5gf1Z_loWuEzBsxPlFDhR

因為Databricks技術一直在快速發展，所以Databricks-Generative-AI-Engineer-Associate認證考試的試題也在不斷變化。因此， VCESoft的考古題也在一直更新。並且，如果你購買了VCESoft的資料，VCESoft將為你提供一年的免費更新服務。只要試題一更新，VCESoft馬上把最新版的資料發送給你。這樣就可以保證你隨時擁有最新版的資料。VCESoft不僅可以幫助你通過考試，還可以幫助你學習最新的知識。这样实惠的资料你千万不要错过。

作為一位 Databricks Databricks-Generative-AI-Engineer-Associate 考生而言，作好充分的準備可以幫助您通過考試。首先您必須去當地考試中心咨詢相關考試信息，然后挑選最新的 Databricks-Generative-AI-Engineer-Associate 考試題庫，因為擁有了最新的 Databricks-Generative-AI-Engineer-Associate 考試題庫可以有利的提高通過考試的機率。使用 VCESoft 的題庫可以節省您寶貴的時間，保證你順利通過 Databricks-Generative-AI-Engineer-Associate 考試。既能幫您節省時間，又可以順利幫助您通過考試，這將是您的最佳選擇。

**>> Databricks-Generative-AI-Engineer-Associate考試大綱 <<**

# 最新Databricks-Generative-AI-Engineer-Associate題庫，Databricks-Generative-AI-Engineer-Associate最新考題

根據過去的考試題和答案的研究，VCESoft提供的Databricks Databricks-Generative-AI-Engineer-Associate練習題和真實的考試試題有緊密的相似性。VCESoft是可以承諾您能100%通過你第一次參加的Databricks Databricks-Generative-AI-Engineer-Associate 認證考試。

## Databricks Databricks-Generative-AI-Engineer-Associate 考試大綱：

| 主題 | 簡介 |
|---|---|
| 主題 1 | • Design Applications: The topic focuses on designing a prompt that elicits a specifically formatted response. It also focuses on selecting model tasks to accomplish a given business requirement. Lastly, the topic covers chain components for a desired model input and output. |
| 主題 2 | • Application Development: In this topic, Generative AI Engineers learn about tools needed to extract data, Langchain<br>• similar tools, and assessing responses to identify common issues. Moreover, the topic includes questions about adjusting an LLM's response, LLM guardrails, and the best LLM based on the attributes of the application. |
| 主題 3 | • Evaluation and Monitoring: This topic is all about selecting an LLM choice and key metrics. Moreover, Generative AI Engineers learn about evaluating model performance. Lastly, the topic includes sub-topics about inference logging and usage of Databricks features. |
| 主題 4 | • Assembling and Deploying Applications: In this topic, Generative AI Engineers get knowledge about coding a chain using a pyfunc mode, coding a simple chain using langchain, and coding a simple chain according to requirements. Additionally, the topic focuses on basic elements needed to create a RAG application. Lastly, the topic addresses sub-topics about registering the model to Unity Catalog using MLflow. |
| 主題 5 | • Governance: Generative AI Engineers who take the exam get knowledge about masking techniques, guardrail techniques, and legal<br>• licensing requirements in this topic. |

## 最新的 Generative AI Engineer Databricks-Generative-AI-Engineer-Associate 免費考試真題 (Q18-Q23):

**問題 #18**
Which TWO chain components are required for building a basic LLM-enabled chat application that includes conversational capabilities, knowledge retrieval, and contextual memory?

- A. Chat loaders
- B. Conversation Buffer Memory
- C. Vector Stores
- D. External tools
- E. (Q)
- F. React Components

**答案：B,C**

解題說明：
Building a basic LLM-enabled chat application with conversational capabilities, knowledge retrieval, and contextual memory requires specific components that work together to process queries, maintain context, and retrieve relevant information. Databricks' Generative AI Engineer documentation outlines key components for such systems, particularly in the context of frameworks like LangChain or Databricks' MosaicML integrations. Let's evaluate the required components:
* Understanding the Requirements:
* Conversational capabilities: The app must generate natural, coherent responses.

* Knowledge retrieval: It must access external or domain-specific knowledge.
* Contextual memory: It must remember prior interactions in the conversation.
* Databricks Reference:"A typical LLM chat application includes a memory component to track conversation history and a retrieval mechanism to incorporate external knowledge"("Databricks Generative AI Cookbook," 2023).
* Evaluating the Options:
* A. (Q): This appears incomplete or unclear (possibly a typo). Without further context, it's not a valid component.
* B. Vector Stores: These store embeddings of documents or knowledge bases, enabling semantic search and retrieval of relevant information for the LLM. This is critical for knowledge retrieval in a chat application.
* Databricks Reference:"Vector stores, such as those integrated with Databricks' Lakehouse, enable efficient retrieval of contextual data for LLMs"("Building LLM Applications with Databricks").
* C. Conversation Buffer Memory: This component stores the conversation history, allowing the LLM to maintain context across multiple turns. It's essential for contextual memory.
* Databricks Reference:"Conversation Buffer Memory tracks prior user inputs and LLM outputs, ensuring context-aware responses"("Generative AI Engineer Guide").
* D. External tools: These (e.g., APIs or calculators) enhance functionality but aren't required for a basicchat app with the specified capabilities.
* E. Chat loaders: These might refer to data loaders for chat logs, but they're not a core chain component for conversational functionality or memory.
* F. React Components: These relate to front-end UI development, not the LLM chain's backend functionality.
* Selecting the Two Required Components:
* Forknowledge retrieval, Vector Stores (B) are necessary to fetch relevant external data, a cornerstone of Databricks' RAG-based chat systems.
* Forcontextual memory, Conversation Buffer Memory (C) is required to maintain conversation history, ensuring coherent and context-aware responses.
* While an LLM itself is implied as the core generator, the question asks for chain components beyond the model, making B and C the minimal yet sufficient pair for a basic application.
Conclusion: The two required chain components areB. Vector StoresandC. Conversation Buffer Memory, as they directly address knowledge retrieval and contextual memory, respectively, aligning with Databricks' documented best practices for LLM-enabled chat applications.

**問題 #19**
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. Decrease the chunk size of embedded documents
- B. Use a smaller embedding model to generate
- C. Reduce the maximum output tokens of the new model
- D. Retrain the response generating model using ALiBi
- E. Reduce the number of records retrieved from the vector database

**答案： A,E**

解題說明：
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.

**問題 #20**
A Generative Al Engineer is deciding between using LSH (Locality Sensitive Hashing) and HNSW (Hierarchical Navigable Small World) for indexing their vector database Their top priority is semantic accuracy Which approach should the Generative Al Engineer use to evaluate these two techniques?

- A. Compare the cosine similarities of the embeddings of returned results against those of a representative sample of test inputs
- B. Compare the Recall-Onented-Understudy for Gistmg Evaluation (ROUGE) scores of returned results for a representative sample of test inputs
- C. Compare the Levenshtein distances of returned results against a representative sample of test inputs
- D. Compare the Bilingual Evaluation Understudy (BLEU) scores of returned results for a representative sample of test inputs

**答案：A**

解題說明：
The task is to choose between LSH and HNSW for a vector database index, prioritizing semantic accuracy.
The evaluation must assess how well each method retrieves semantically relevant results. Let's evaluate the options.
* Option A: Compare the cosine similarities of the embeddings of returned results against those of a representative sample of test inputs
* Cosine similarity measures semantic closeness between vectors, directly assessing retrieval accuracy in a vector database. Comparing returned results' embeddings to test inputs' embeddings evaluates how well LSH or HNSW preserves semantic relationships, aligning with the priority.
* Databricks Reference:"Cosine similarity is a standard metric for evaluating vector search accuracy"("Databricks Vector Search Documentation," 2023).
* Option B: Compare the Bilingual Evaluation Understudy (BLEU) scores of returned results for a representative sample of test inputs
* BLEU evaluates text generation (e.g., translations), not vector retrieval accuracy. It's irrelevant for indexing performance.
* Databricks Reference:"BLEU applies to generative tasks, not retrieval"("Generative AI Cookbook").
* Option C: Compare the Recall-Oriented-Understudy for Gisting Evaluation (ROUGE) scores of returned results for a representative sample of test inputs
* ROUGE is for summarization evaluation, not vector search. It doesn't measure semantic accuracy in retrieval.
* Databricks Reference:"ROUGE is unsuited for vector database evaluation"("Building LLM Applications with Databricks").
* Option D: Compare the Levenshtein distances of returned results against a representative sample of test inputs
* Levenshtein distance measures string edit distance, not semantic similarity in embeddings. It's inappropriate for vector-based retrieval.
* Databricks Reference: No specific support for Levenshtein in vector search contexts.
Conclusion: Option A (cosine similarity) is the correct approach, directly evaluating semantic accuracy in vector retrieval, as recommended by Databricks for Vector Search assessments.

**問題 #21**
A Generative AI Engineer is developing a chatbot designed to assist users with insurance-related queries. The chatbot is built on a large language model (LLM) and is conversational. However, to maintain the chatbot's focus and to comply with company policy, it must not provide responses to questions about politics. Instead, when presented with political inquiries, the chatbot should respond with a standard message:
"Sorry, I cannot answer that. I am a chatbot that can only answer questions around insurance." Which framework type should be implemented to solve this?

- A. Security Guardrail
- B. Contextual Guardrail
- C. Compliance Guardrail
- D. Safety Guardrail

**答案：D**

解題說明：
In this scenario, the chatbot must avoid answering political questions and instead provide a standard message for such inquiries. Implementing a Safety Guardrail is the appropriate solution for this:
What is a Safety Guardrail?
Safety guardrails are mechanisms implemented in Generative AI systems to ensure the model behaves within specific bounds. In this case, it ensures the chatbot does not answer politically sensitive or irrelevant questions, which aligns with the business rules.

Preventing Responses to Political Questions:
The Safety Guardrail is programmed to detect specific types of inquiries (like political questions) and prevent the model from generating responses outside its intended domain. When such queries are detected, the guardrail intervenes and provides a pre-defined response: "Sorry, I cannot answer that. I am a chatbot that can only answer questions around insurance." How It Works in Practice:
The LLM system can include a classification layer or trigger rules based on specific keywords related to politics. When such terms are detected, the Safety Guardrail blocks the normal generation flow and responds with the fixed message.
Why Other Options Are Less Suitable:
B (Security Guardrail): This is more focused on protecting the system from security vulnerabilities or data breaches, not controlling the conversational focus.
C (Contextual Guardrail): While context guardrails can limit responses based on context, safety guardrails are specifically about ensuring the chatbot stays within a safe conversational scope.
D (Compliance Guardrail): Compliance guardrails are often related to legal and regulatory adherence, which is not directly relevant here.
Therefore, a Safety Guardrail is the right framework to ensure the chatbot only answers insurance-related queries and avoids political discussions.

**問題 #22**
A Generative Al Engineer has successfully ingested unstructured documents and chunked them by document sections. They would like to store the chunks in a Vector Search index. The current format of the dataframe has two columns: (i) original document file name (ii) an array of text chunks for each document.
What is the most performant way to store this dataframe?

- A. Split the data into train and test set, create a unique identifier for each document, then save to a Delta table
- B. Flatten the dataframe to one chunk per row, create a unique identifier for each row, and save to a Delta table
- C. First create a unique identifier for each document, then save to a Delta table
- D. Store each chunk as an independent JSON file in Unity Catalog Volume. For each JSON file, the key is the document section name and the value is the array of text chunks for that section

**答案：B**

解題說明：
* Problem Context: The engineer needs an efficient way to store chunks of unstructured documents to facilitate easy retrieval and search. The current dataframe consists of document filenames and associated text chunks.
* Explanation of Options:
* Option A: Splitting into train and test sets is more relevant for model training scenarios and not directly applicable to storage for retrieval in a Vector Search index.
* Option B: Flattening the dataframe such that each row contains a single chunk with a unique identifier is the most performant for storage and retrieval. This structure aligns well with how data is indexed and queried in vector search applications, making it easier to retrieve specific chunks efficiently.
* Option C: Creating a unique identifier for each document only does not address the need to access individual chunks efficiently, which is critical in a Vector Search application.
* Option D: Storing each chunk as an independent JSON file creates unnecessary overhead and complexity in managing and querying large volumes of files.
OptionBis the most efficient and practical approach, allowing for streamlined indexing and retrieval processes in a Delta table environment, fitting the requirements of a Vector Search index.

**問題 #23**
.....

- Databricks-Generative-AI-Engineer-Associate考試證照 👆 Databricks-Generative-AI-Engineer-Associate熱門認證 □ Databricks-Generative-AI-Engineer-Associate考試題庫 □ 打開➽ www.newdumpspdf.com □搜尋➡ Databricks-Generative-AI-Engineer-Associate □□□以免費下載考試資料Databricks-Generative-AI-Engineer-Associate證照資訊
- Databricks-Generative-AI-Engineer-Associate認證指南 □ Databricks-Generative-AI-Engineer-Associate考古題 □ Databricks-Generative-AI-Engineer-Associate真題材料 □ 來自網站➽ www.newdumpspdf.com □打開並搜索✔ Databricks-Generative-AI-Engineer-Associate □✔□免費下載Databricks-Generative-AI-Engineer-Associate資訊
- Databricks-Generative-AI-Engineer-Associate題庫資訊 ⧗ 最新Databricks-Generative-AI-Engineer-Associate考古題 □ Databricks-Generative-AI-Engineer-Associate軟件版 □ 在「 www.newdumpspdf.com 」網站上免費搜索{ Databricks-Generative-AI-Engineer-Associate }題庫Databricks-Generative-AI-Engineer-Associate認證指南
- Databricks-Generative-AI-Engineer-Associate考試大綱：Databricks Certified Generative AI Engineer Associate幫助您壹次通過考試，Databricks 最新Databricks-Generative-AI-Engineer-Associate題庫 □ ➠ www.newdumpspdf.com □ □上的免費下載□ Databricks-Generative-AI-Engineer-Associate □頁面立即打開Databricks-Generative-AI-Engineer-Associate考古題
- Databricks-Generative-AI-Engineer-Associate軟件版 □ Databricks-Generative-AI-Engineer-Associate考試 □ Databricks-Generative-AI-Engineer-Associate題庫資訊 □ 到✔ www.pdfexamdumps.com □✔□搜尋➠ Databricks-Generative-AI-Engineer-Associate □以獲取免費下載考試資料最新Databricks-Generative-AI-Engineer-Associate題庫資源
- Databricks Databricks-Generative-AI-Engineer-Associate考試大綱：Databricks Certified Generative AI Engineer Associate壹次通過考試 ℹ 免費下載【 Databricks-Generative-AI-Engineer-Associate 】只需進入[ www.newdumpspdf.com ]網站Databricks-Generative-AI-Engineer-Associate認證資料
- Databricks-Generative-AI-Engineer-Associate認證資料 □ 最新Databricks-Generative-AI-Engineer-Associate考古題 □ Databricks-Generative-AI-Engineer-Associate最新試題 □ ➠ www.pdfexamdumps.com □是獲取▷ Databricks-Generative-AI-Engineer-Associate ◁免費下載的最佳網站Databricks-Generative-AI-Engineer-Associate證照資訊
- 全面包括的Databricks-Generative-AI-Engineer-Associate考試大綱和資格考試中的領導者和無與倫比的最新Databricks-Generative-AI-Engineer-Associate題庫 □ 透過▷ www.newdumpspdf.com ◁輕鬆獲取[ Databricks-Generative-AI-Engineer-Associate ]免費下載Databricks-Generative-AI-Engineer-Associate熱門認證
- Databricks-Generative-AI-Engineer-Associate參考資料 □ Databricks-Generative-AI-Engineer-Associate證照信息 □ Databricks-Generative-AI-Engineer-Associate考古題 □ 透過“ www.pdfexamdumps.com ”輕鬆獲取《 Databricks-Generative-AI-Engineer-Associate 》免費下載Databricks-Generative-AI-Engineer-Associate認證資料
- Databricks-Generative-AI-Engineer-Associate真題材料 □ Databricks-Generative-AI-Engineer-Associate最新試題 □ Databricks-Generative-AI-Engineer-Associate考古題 □ 打開網站□ www.newdumpspdf.com □搜索【 Databricks-Generative-AI-Engineer-Associate 】免費下載Databricks-Generative-AI-Engineer-Associate證照信息
- Databricks-Generative-AI-Engineer-Associate資料 □ Databricks-Generative-AI-Engineer-Associate認證資料 □ Databricks-Generative-AI-Engineer-Associate考試 □ 免費下載▸ Databricks-Generative-AI-Engineer-Associate ◂只需在➽ tw.fast2test.com □上搜索Databricks-Generative-AI-Engineer-Associate證照信息
- www.stes.tyc.edu.tw, jadaakbv998401.blogthisbiz.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, myportal.utt.edu.tt, myportal.utt.edu.tt, montyzhca518377.answerblogs.com, geraldltfm487674.get-blogging.com, ronaldvsil768801.smblogsites.com, diegoskzm884960.spintheblog.com, fayjeiw179584.blogcudinti.com, kobibbdw896239.birderswiki.com, Disposable vapes

順便提一下，可以從雲存儲中下載VCESoft Databricks-Generative-AI-Engineer-Associate考試題庫的完整版：https://drive.google.com/open?id=1PGXX-jyac8z5gf1Z_loWuEzBsxPlFDhR