


# Exam Databricks-Generative-AI-Engineer-Associate Quiz & Databricks-Generative-AI-Engineer-Associate Dumps Free



P.S. Free 2026 Databricks Databricks-Generative-AI-Engineer-Associate dumps are available on Google Drive shared by PrepAwayTest: <https://drive.google.com/open?id=1EYVRF82Fi-zjPV-keh-jqQLPvGRTEOhA>

These Databricks Certified Generative AI Engineer Associate (Databricks-Generative-AI-Engineer-Associate) exam questions help applicants prepare well prior to entering the actual Databricks Certified Generative AI Engineer Associate (Databricks-Generative-AI-Engineer-Associate) exam center. Due to our actual Databricks-Generative-AI-Engineer-Associate Exam Dumps, our valued customers always pass their Databricks Databricks-Generative-AI-Engineer-Associate exam on the very first try hence, saving their precious time and money too.

## Databricks Databricks-Generative-AI-Engineer-Associate Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"><li>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.</li></ul>
Topic 2	<ul style="list-style-type: none"><li>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.</li></ul>
Topic 3	<ul style="list-style-type: none"><li>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.</li></ul>
Topic 4	<ul style="list-style-type: none"><li>Data Preparation: Generative AI Engineers covers a chunking strategy for a given document structure and model constraints. The topic also focuses on filter extraneous content in source documents. Lastly, Generative AI Engineers also learn about extracting document content from provided source data and format.</li></ul>

Topic 5	<ul style="list-style-type: none"> <li>• Application Development: In this topic, Generative AI Engineers learn about tools needed to extract data, Langchain</li> <li>• 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.</li> </ul>
---------	---

>> Exam Databricks-Generative-AI-Engineer-Associate Quiz <<

## **Free PDF Quiz 2026 Trustable Databricks Databricks-Generative-AI-Engineer-Associate: Exam Databricks Certified Generative AI Engineer Associate Quiz**

Our company has the highly authoritative and experienced team. In order to let customers enjoy the best service, all Databricks-Generative-AI-Engineer-Associate exam prep of our company were designed by hundreds of experienced experts. Our Databricks-Generative-AI-Engineer-Associate test questions will help customers learn the important knowledge about exam. If you buy our products, it will be very easy for you to have the mastery of a core set of knowledge in the shortest time, at the same time, our Databricks-Generative-AI-Engineer-Associate Test Torrent can help you avoid falling into rote learning habits. You just need to spend 20 to 30 hours on study, and then you can take your exam. In addition, the authoritative production team of our Databricks-Generative-AI-Engineer-Associate exam prep will update the study system every day in order to make our customers enjoy the newest information.

### **Databricks Certified Generative AI Engineer Associate Sample Questions (Q57-Q62):**

#### **NEW QUESTION # 57**

A Generative AI Engineer is building a system which will answer questions on latest stock news articles. Which will NOT help with ensuring the outputs are relevant to financial news?

- A. Incorporate manual reviews to correct any problematic outputs prior to sending to the users
- B. Implement a comprehensive guardrail framework that includes policies for content filters tailored to the finance sector.
- C. Increase the compute to improve processing speed of questions to allow greater relevancy analysis C Implement a profanity filter to screen out offensive language

**Answer: C**

Explanation:

In the context of ensuring that outputs are relevant to financial news, increasing compute power (option B) does not directly improve the relevance of the LLM-generated outputs. Here's why:

\* Compute Power and Relevancy: Increasing compute power can help the model process inputs faster, but it does not inherently improve the relevance of the answers. Relevancy depends on the data sources, the retrieval method, and the filtering mechanisms in place, not on how quickly the model processes the query.

\* What Actually Helps with Relevance: Other methods, like content filtering, guardrails, or manual review, can directly impact the relevance of the model's responses by ensuring the model focuses on pertinent financial content. These methods help tailor the LLM's responses to the financial domain and avoid irrelevant or harmful outputs.

\* Why Other Options Are More Relevant:

\* A (Comprehensive Guardrail Framework): This will ensure that the model avoids generating content that is irrelevant or inappropriate in the finance sector.

\* C (Profanity Filter): While not directly related to financial relevancy, ensuring the output is clean and professional is still important in maintaining the quality of responses.

\* D (Manual Review): Incorporating human oversight to catch and correct issues with the LLM's output ensures the final answers are aligned with financial content expectations.

Thus, increasing compute power does not help with ensuring the outputs are more relevant to financial news, making option B the correct answer.

#### **NEW QUESTION # 58**

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_searchtest",
    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.get\_index()
- B. vsc.similarity\_search()
- C. vsc.create\_direct\_access\_index()
- D. vsc.create\_delta\_sync\_index()

**Answer: D**

Explanation:

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.

## NEW QUESTION # 59

A Generative AI Engineer is creating an LLM system that will retrieve news articles from the year 1918 and related to a user's query and summarize them. The engineer has noticed that the summaries are generated well but often also include an explanation of how the summary was generated, which is undesirable.

Which change could the Generative AI Engineer perform to mitigate this issue?

- A. Split the LLM output by newline characters to truncate away the summarization explanation.
- B. **Provide few shot examples of desired output format to the system and/or user prompt.**
- C. Tune the chunk size of news articles or experiment with different embedding models.
- D. Revisit their document ingestion logic, ensuring that the news articles are being ingested properly.

**Answer: B**

Explanation:

To mitigate the issue of the LLM including explanations of how summaries are generated in its output, the best approach is to adjust the training or prompt structure. Here's why Option D is effective:

- \* Few-shot Learning: By providing specific examples of how the desired output should look (i.e., just the summary without explanation), the model learns the preferred format. This few-shot learning approach helps the model understand not only what content to generate but also how to format its responses.
- \* Prompt Engineering: Adjusting the user prompt to specify the desired output format clearly can guide the LLM to produce summaries without additional explanatory text. Effective prompt design is crucial in controlling the behavior of generative models.

#### Why Other Options Are Less Suitable:

- \* A: While technically feasible, splitting the output by newline and truncating could lead to loss of important content or create awkward breaks in the summary.
- \* B: Tuning chunk sizes or changing embedding models does not directly address the issue of the model's tendency to generate explanations along with summaries.
- \* C: Revisiting document ingestion logic ensures accurate source data but does not influence how the model formats its output. By using few-shot examples and refining the prompt, the engineer directly influences the output format, making this approach the most targeted and effective solution.

#### NEW QUESTION # 60

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.get\_index()
- B. vsc.similarity\_search()
- C. vsc.create\_direct\_access\_index()
- D. vsc.create\_delta\_sync\_index()

#### Answer: D

##### Explanation:

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.

#### NEW QUESTION # 61

A Generative AI 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 AI.
- B. Include in the system prompt that any information it sees will be about SnoPenAI, even if no data filtering is performed.

- C. Keep all articles because the RAG application needs to understand non-company content to avoid answering questions about them.
- D. Consolidate all SnoPen AI related documents into a single chunk in the vector database.

#### Answer: A

Explanation:

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 by including 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.

#### NEW QUESTION # 62

.....

If you want to get promotions or high-paying jobs in the Databricks sector, then it is important for you to crack the Databricks Certified Generative AI Engineer Associate (Databricks-Generative-AI-Engineer-Associate) certification exam. The Databricks Databricks-Generative-AI-Engineer-Associate certification has become the best way to validate your skills and accelerate your tech career. Databricks-Generative-AI-Engineer-Associate Exam applicants who are doing jobs or busy with their other matters usually don't have enough time to study for the test.

**Databricks-Generative-AI-Engineer-Associate Dumps Free:** <https://www.prepawaytest.com/Databricks/Databricks-Generative-AI-Engineer-Associate-practice-exam-dumps.html>

- Exam Databricks-Generative-AI-Engineer-Associate Quiz 100% Pass | Trustable Databricks Databricks Certified Generative AI Engineer Associate Dumps Free Pass for sure  Download  Databricks-Generative-AI-Engineer-Associate  for free by simply entering ➡ [www.prep4away.com](http://www.prep4away.com)  website  Valid Databricks-Generative-AI-Engineer-Associate Exam Cost
- Free PDF 2026 Databricks-Generative-AI-Engineer-Associate Pass-Sure Exam Quiz  Go to website ➡ [www.pdfvce.com](http://www.pdfvce.com)  open and search for ( Databricks-Generative-AI-Engineer-Associate ) to download for free   Valid Databricks-Generative-AI-Engineer-Associate Dumps
- 100% Pass 2026 Databricks-Generative-AI-Engineer-Associate: Databricks Certified Generative AI Engineer Associate – High Hit-Rate Exam Quiz  Easily obtain free download of ➡ Databricks-Generative-AI-Engineer-Associate ↳ by searching on  [www.dumpsmaterials.com](http://www.dumpsmaterials.com)   Databricks-Generative-AI-Engineer-Associate Latest Exam Simulator
- Valid Databricks-Generative-AI-Engineer-Associate Dumps  Test Databricks-Generative-AI-Engineer-Associate Book  Databricks-Generative-AI-Engineer-Associate Practice Test Online  Copy URL ➡ [www.pdfvce.com](http://www.pdfvce.com)  open and search for [ Databricks-Generative-AI-Engineer-Associate ] to download for free  Authentic Databricks-Generative-AI-Engineer-Associate Exam Hub
- Databricks-Generative-AI-Engineer-Associate Sample Questions  Databricks-Generative-AI-Engineer-Associate Prepaway Dumps  Valid Databricks-Generative-AI-Engineer-Associate Dumps  Search for ✓ Databricks-Generative-AI-Engineer-Associate  ✓  and download it for free immediately on ➡ [www.torrentvce.com](http://www.torrentvce.com)      Databricks-Generative-AI-Engineer-Associate Reliable Braindumps Pdf
- Free PDF Quiz Databricks - High-quality Exam Databricks-Generative-AI-Engineer-Associate Quiz  Immediately open [ [www.pdfvce.com](http://www.pdfvce.com) ] and search for ➡ Databricks-Generative-AI-Engineer-Associate  to obtain a free download  Latest Databricks-Generative-AI-Engineer-Associate Exam Notes
- Authentic Databricks-Generative-AI-Engineer-Associate Exam Hub  Valid Databricks-Generative-AI-Engineer-Associate Test Sample  Study Databricks-Generative-AI-Engineer-Associate Plan  Copy URL ✓

www.practicevce.com ✓ open and search for ➤ Databricks-Generative-AI-Engineer-Associate to download for free  
Databricks-Generative-AI-Engineer-Associate Reliable Braindumps Pdf

P.S. Free 2026 Databricks-Generative-AI-Engineer-Associate dumps are available on Google Drive shared by PrepAwayTest: <https://drive.google.com/open?id=1EYVRF82Fi-zjPV-keh-jqQLPvGRtEoH>