

# Databricks-Generative-AI-Engineer-Associate Exam Braindumps & Databricks-Generative-AI-Engineer-Associate Test Quiz & Databricks-Generative-AI-Engineer-Associate Practice Material

DarshanSathyanarayan/  
**Databricks\_Generative\_A...**

Databricks Certified Generative AI Engineer  
Associate Exam Question dumps and Practice Tests  
with Tips



0



0



1

Star



0

Forks



DOWNLOAD the newest TestValid Databricks-Generative-AI-Engineer-Associate PDF dumps from Cloud Storage for free:  
[https://drive.google.com/open?id=1dgESegKW4R8a\\_Dopj3Q8qfosJOv4qxK8](https://drive.google.com/open?id=1dgESegKW4R8a_Dopj3Q8qfosJOv4qxK8)

Our experts are well-aware of the problems of exam candidates particularly of those who can't manage to spare time to study the Databricks-Generative-AI-Engineer-Associate exam questions due to their heavy work pressure. Hence, our Databricks-Generative-AI-Engineer-Associate study materials have been developed into a simple content and language for our worthy customers all over the world. What is more, you will find there are only the keypoints in our Databricks-Generative-AI-Engineer-Associate learning guide.

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

|                              |                                                                                                                             |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| Certification Vendor:        | Databricks                                                                                                                  |
| Exam Name:                   | Databricks Certified Generative AI Engineer Associate                                                                       |
| Exam Number:                 | Databricks-Generative-AI-Engineer-Associate                                                                                 |
| Exam Format:                 | Multiple choice, Multiple select                                                                                            |
| Available Languages:         | English                                                                                                                     |
| Exam Duration:               | 90 minutes                                                                                                                  |
| Real Exam Qty:               | 45-60                                                                                                                       |
| Related Certifications:      | Databricks Certified Data Engineer Associate<br>Databricks Certified Machine Learning Associate                             |
| Certificate Validity Period: | 2 years                                                                                                                     |
| Exam Price:                  | USD 200                                                                                                                     |
| Recommended Training:        | Mosaic AI Documentation<br>Databricks Academy - Generative AI Courses                                                       |
| Exam Registration:           | Databricks Certification Portal                                                                                             |
| Sample Questions:            | Databricks Databricks-Generative-AI-Engineer-Associate Sample Questions                                                     |
| Exam Way:                    | Online proctored exam                                                                                                       |
| Pre Condition:               | No strict prerequisites, but experience with Python, machine learning fundamentals, and Databricks platform is recommended. |
| Official Syllabus URL:       | <a href="https://www.databricks.com/learn/certification">https://www.databricks.com/learn/certification</a>                 |

>> Latest Databricks-Generative-AI-Engineer-Associate Version <<

**Exam Databricks Databricks-Generative-AI-Engineer-Associate Objectives |  
Databricks-Generative-AI-Engineer-Associate Pass4sure Study Materials**

Our Databricks-Generative-AI-Engineer-Associate training materials have won great success in the market. Tens of thousands of the candidates are learning on our Databricks-Generative-AI-Engineer-Associate practice engine. First of all, our Databricks-Generative-AI-Engineer-Associate study dumps cover all related tests about computers. It will be easy for you to find your prepared learning material. If you are suspicious of our Databricks-Generative-AI-Engineer-Associate Exam Questions, you can download the free demo from our official websites.

## 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>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 3 | <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>                                                                      |
| Topic 4 | <ul style="list-style-type: none"><li>Governance: Generative AI Engineers who take the exam get knowledge about masking techniques, guardrail techniques, and legal</li><li>licensing requirements in this topic.</li></ul>                                                                                                                                                                                                                                               |
| Topic 5 | <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>                                                                                                                                         |

## Databricks Certified Generative AI Engineer Associate Sample Questions (Q63-Q68):

### NEW QUESTION # 63

A Generative AI Engineer has been reviewing issues with their company's LLM-based question-answering assistant and has determined that a technique called prompt chaining could help alleviate some performance concerns. However, to suggest this to their team, they have to clearly explain how it works and how it can benefit their question-answering assistant. Which explanation do they communicate to the team?

- A. It allows you to break down complex tasks into multiple independent subtasks. This enables the assistant to generate more comprehensive and accurate responses.
- B. It allows you to decrease the effort involved in crafting a prompt. Chains make it possible to reuse prompt text across multiple different use cases.
- C. It reduces the average cost of a typical request. Chains make more efficient use of the tokens produced to generate higher quality responses with fewer tokens.
- D. It allows you to reduce the latency of your applications. By having multiple chains participating in the response as a chain, you increase the rate at which the response is generated.

**Answer: A**

Explanation:

Prompt chaining is a fundamental design pattern in LLM application development used to handle complexity. Instead of sending a

single, massive, and highly complex prompt to an LLM-which often results in reasoning errors or hallucinations-chaining breaks the logic into a sequence of smaller, targeted steps. For example, a legal assistant might first chain a step to "identify the legal jurisdiction," followed by a step to "extract relevant statutes," and finally a step to "summarize the findings." This modularity improves reliability because each prompt has a narrower focus, making it easier for the model to follow instructions accurately. While it may actually increase latency (contradicting B) and cost (contradicting D) due to multiple API calls, the primary engineering benefit is the significant boost in the quality and robustness of the output. It also allows for intermediate validation and error handling between steps, which is impossible in a single-call architecture.

#### NEW QUESTION # 64

A Generative AI Engineer is experimenting with using parameters to configure an agent in Mosaic Agent Framework. However, they are struggling to get the agent to respond with relevant information with this configuration:

```
config = {"prompt_template": "You are a trivia bot. Generate a question based on the user's input: {user_input}", "input_vars": ["user_input"], "parameters": {"temperature": 0.01, "max_tokens": 500}} Which error is causing the problem?
```

- A. The prompt is not wrapped in ChatModel
- B. The prompt does not list available agents for the LLM to call
- C. The prompt does not set the retriever schema
- **D. The prompt does not parse the user's input vars**

**Answer: D**

Explanation:

In the Mosaic AI Agent Framework and underlying LangChain-based configurations, the "input\_vars" or "input\_variables" must be correctly mapped and referenced within the template. If the configuration dictionary identifies user\_input as the variable but the logic executing the chain does not correctly "inject" the runtime value into the {user\_input} placeholder, the LLM will receive a literal string (or an empty value) rather than the user's actual question. This results in the model failing to provide relevant information because it essentially doesn't know what the user asked. Engineering standards require ensuring that the key used in the input\_vars list matches the key in the JSON payload sent to the model serving endpoint. If there is a mismatch or a failure to parse, the prompt remains static, leading to generic or irrelevant responses.

#### NEW QUESTION # 65

A Generative AI 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 AI Engineer use to evaluate these two techniques?

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

**Answer: C**

Explanation:

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.

### NEW QUESTION # 66

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. Respond with "In Stock" if the customer asks for a product.
- C. 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.
- **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"}.**

**Answer: D**

Explanation:

\* 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.

### NEW QUESTION # 67

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

**Answer: A,B**

Explanation:

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.

## NEW QUESTION # 68

.....

**Exam Databricks-Generative-AI-Engineer-Associate Objectives:** <https://www.testvalid.com/Databricks-Generative-AI-Engineer-Associate-exam-collection.html>

- Databricks-Generative-AI-Engineer-Associate Reliable Exam Review □ Databricks-Generative-AI-Engineer-Associate Exam Duration □ Valid Databricks-Generative-AI-Engineer-Associate Test Syllabus □ The page for free download of □ Databricks-Generative-AI-Engineer-Associate □ on ⇒ [www.dumpsmaterials.com](http://www.dumpsmaterials.com) ⇐ will open immediately □ Test Databricks-Generative-AI-Engineer-Associate Questions Pdf
- New Databricks-Generative-AI-Engineer-Associate Test Preparation □ Databricks-Generative-AI-Engineer-Associate Latest Braindumps Ppt □ Databricks-Generative-AI-Engineer-Associate Braindumps Pdf □ Download ➡ Databricks-Generative-AI-Engineer-Associate □□□ for free by simply entering ➡ [www.pdfvce.com](http://www.pdfvce.com) □ website □ Databricks-Generative-AI-Engineer-Associate Braindumps Pdf
- Databricks-Generative-AI-Engineer-Associate Real Dump □ Valid Databricks-Generative-AI-Engineer-Associate Exam Tips □ Latest Databricks-Generative-AI-Engineer-Associate Test Blueprint □ Download ➡ Databricks-Generative-AI-Engineer-Associate □ for free by simply entering ➡ [www.vce4dumps.com](http://www.vce4dumps.com) ◀ website □ Databricks-Generative-AI-Engineer-Associate Learning Mode
- New Databricks-Generative-AI-Engineer-Associate Test Pattern □ Valid Databricks-Generative-AI-Engineer-Associate Exam Tips □ Databricks-Generative-AI-Engineer-Associate New Exam Braindumps □ Download ▷ Databricks-Generative-AI-Engineer-Associate ◁ for free by simply searching on ☀ [www.pdfvce.com](http://www.pdfvce.com) □ ☀ □ □ Valid Databricks-Generative-AI-Engineer-Associate Exam Tips
- Databricks-Generative-AI-Engineer-Associate Reliable Exam Review □ Free Databricks-Generative-AI-Engineer-Associate Exam Dumps □ Databricks-Generative-AI-Engineer-Associate Reliable Exam Review □ Search for ➡ Databricks-Generative-AI-Engineer-Associate □ and obtain a free download on ➡ [www.verifiedumps.com](http://www.verifiedumps.com) □ □ □ Databricks-Generative-AI-Engineer-Associate Interactive EBook
- Databricks-Generative-AI-Engineer-Associate Dumps □ Databricks-Generative-AI-Engineer-Associate Learning Mode □ Online Databricks-Generative-AI-Engineer-Associate Training Materials □ Enter ➤ [www.pdfvce.com](http://www.pdfvce.com) □ and search for □ Databricks-Generative-AI-Engineer-Associate □ to download for free □ Test Databricks-Generative-AI-Engineer-Associate Questions Pdf

- New Databricks-Generative-AI-Engineer-Associate Test Pattern ☐ Test Databricks-Generative-AI-Engineer-Associate Questions Pdf ☐ Free Databricks-Generative-AI-Engineer-Associate Exam Dumps ☐ Simply search for « Databricks-Generative-AI-Engineer-Associate » for free download on ➤ [www.prepawayete.com](http://www.prepawayete.com) ☐ ☐ Databricks-Generative-AI-Engineer-Associate Interactive EBook
- 100% Pass 2026 Databricks Databricks-Generative-AI-Engineer-Associate –High Hit-Rate Latest Version ☐ Open ✓ [www.pdfvce.com](http://www.pdfvce.com) ☐ ✓ ☐ and search for ✨ Databricks-Generative-AI-Engineer-Associate ✨ ☐ to download exam materials for free ☐ New Databricks-Generative-AI-Engineer-Associate Test Preparation
- Databricks-Generative-AI-Engineer-Associate Quiz Torrent - Databricks-Generative-AI-Engineer-Associate Exam Guide - Databricks-Generative-AI-Engineer-Associate Test Braindumps ☐ Download [ Databricks-Generative-AI-Engineer-Associate ] for free by simply entering ➤ [www.exam4labs.com](http://www.exam4labs.com) ☐ website ☐ Databricks-Generative-AI-Engineer-Associate Reliable Exam Review
- Databricks-Generative-AI-Engineer-Associate Quiz Torrent - Databricks-Generative-AI-Engineer-Associate Exam Guide - Databricks-Generative-AI-Engineer-Associate Test Braindumps ☐ ☐ [www.pdfvce.com](http://www.pdfvce.com) ☐ is best website to obtain ➡ Databricks-Generative-AI-Engineer-Associate ☐ for free download ☐ Databricks-Generative-AI-Engineer-Associate Learning Mode
- 100% Pass 2026 Databricks Databricks-Generative-AI-Engineer-Associate –High Hit-Rate Latest Version ☐ Download 「 Databricks-Generative-AI-Engineer-Associate 」 for free by simply entering ➡ [www.prepawayete.com](http://www.prepawayete.com) ☐ ☐ ☐ website ☐ ☐ Databricks-Generative-AI-Engineer-Associate Braindumps Pdf
- [giphy.com](https://giphy.com), [scalar.usc.edu](https://scalar.usc.edu), [quay.io](https://quay.io), [disqus.com](https://disqus.com), [www.slideshare.net](https://www.slideshare.net), [www.slideshare.net](https://www.slideshare.net), [disqus.com](https://disqus.com), [app.parler.com](https://app.parler.com), [www.shippingexplorer.net](https://www.shippingexplorer.net), [www.shippingexplorer.net](https://www.shippingexplorer.net), Disposable vapes

P.S. Free & New Databricks-Generative-AI-Engineer-Associate dumps are available on Google Drive shared by TestValid:  
[https://drive.google.com/open?id=1dgESegKW4R8a\\_Dopj3Q8qfosJOv4qxK8](https://drive.google.com/open?id=1dgESegKW4R8a_Dopj3Q8qfosJOv4qxK8)