

# New Databricks-Generative-AI-Engineer-Associate Exam Papers & Reliable Databricks-Generative-AI-Engineer-Associate Test Cost



2025 Latest Pass4guide Databricks-Generative-AI-Engineer-Associate PDF Dumps and Databricks-Generative-AI-Engineer-Associate Exam Engine Free Share: [https://drive.google.com/open?id=1vrzp1CK7kG1m33hby9a\\_DnCw68BWdOd1](https://drive.google.com/open?id=1vrzp1CK7kG1m33hby9a_DnCw68BWdOd1)

Maybe you want to get the Databricks-Generative-AI-Engineer-Associate certification, but daily work and long-time traffic make you busier to improve yourself. Thanks to our Databricks-Generative-AI-Engineer-Associate training materials, you can learn for your certification anytime, everywhere. If you get our products, you will surely find a better self. As we all know, the best way to gain confidence is to do something successfully. With our Databricks-Generative-AI-Engineer-Associate Study Guide, you will easily pass the Databricks-Generative-AI-Engineer-Associate examination and gain more confidence.

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

Topic	Details
Topic 1	<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 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>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>
---------	---

>> New Databricks-Generative-AI-Engineer-Associate Exam Papers <<

## Get Special 30% EXTRA Discount on Databricks-Generative-AI-Engineer-Associate Dumps By Pass4guide

We provide first-rate service on the Databricks-Generative-AI-Engineer-Associate learning prep to the clients and they include the service before and after the sale, 24-hours online customer service and long-distance assistance, the refund service and the update service. The client can try out our and download Databricks-Generative-AI-Engineer-Associate Guide materials freely before the sale and if the client have problems about our Databricks-Generative-AI-Engineer-Associate study braindumps after the sale they can contact our customer service at any time.

### Databricks Certified Generative AI Engineer Associate Sample Questions (Q18-Q23):

#### NEW QUESTION # 18

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. 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.
- B. Respond with "In Stock" if the customer asks for a product.
- C. Respond with "Out of Stock" if the customer asks for a product.
- 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 # 19

A Generative AI Engineer has created a RAG application which can help employees retrieve answers from an internal knowledge base, such as Confluence pages or Google Drive. The prototype application is now working with some positive feedback from internal company testers. Now the Generative AI Engineer wants to formally evaluate the system's performance and understand where to focus their efforts to further improve the system.

How should the Generative AI Engineer evaluate the system?

- A. Use cosine similarity score to comprehensively evaluate the quality of the final generated answers.
- B. **Curate a dataset that can test the retrieval and generation components of the system separately. Use MLflow's built in evaluation metrics to perform the evaluation on the retrieval and generation components.**
- C. Use an LLM-as-a-judge to evaluate the quality of the final answers generated.
- D. Benchmark multiple LLMs with the same data and pick the best LLM for the job.

**Answer: B**

Explanation:

\* Problem Context: After receiving positive feedback for the RAG application prototype, the next step is to formally evaluate the system to pinpoint areas for improvement.

\* Explanation of Options:

\* Option A: While cosine similarity scores are useful, they primarily measure similarity rather than the overall performance of an RAG system.

\* Option B: This option provides a systematic approach to evaluation by testing both retrieval and generation components separately. This allows for targeted improvements and a clear understanding of each component's performance, using MLflow's metrics for a structured and standardized assessment.

\* Option C: Benchmarking multiple LLMs does not focus on evaluating the existing system's components but rather on comparing different models.

\* Option D: Using an LLM as a judge is subjective and less reliable for systematic performance evaluation.

OptionBis the most comprehensive and structured approach, facilitating precise evaluations and improvements on specific components of the RAG system.

## NEW QUESTION # 20

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:

```
{"error_code": "BAD_REQUEST", "message": "Bad request: rpc error: code = InvalidArgument desc = prompt token count (4595) cannot exceed 4096..."}  
AQUANTUM
```

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

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

**Answer: B,D**

Explanation:

\* 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.

OptionsCandDare 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.

## NEW QUESTION # 21

A Generative AI Engineer is building a system that will answer questions on currently unfolding news topics.

As such, it pulls information from a variety of sources including articles and social media posts. They are concerned about toxic posts on social media causing toxic outputs from their system.

Which guardrail will limit toxic outputs?

- A. Use only approved social media and news accounts to prevent unexpected toxic data from getting to the LLM.
- B. Log all LLM system responses and perform a batch toxicity analysis monthly.
- C. Reduce the amount of context Items the system will Include in consideration for its response.
- D. Implement rate limiting

**Answer: A**

Explanation:

The system answers questions on unfolding news topics using articles and social media, with a concern about toxic outputs from toxic inputs. A guardrail must limit toxicity in the LLM's responses. Let's evaluate the options.

\* Option A: Use only approved social media and news accounts to prevent unexpected toxic data from getting to the LLM  
\* Curating input sources (e.g., verified accounts) reduces exposure to toxic content at the data ingestion stage, directly limiting toxic outputs. This is a proactive guardrail aligned with data quality control.

\* Databricks Reference: "Control input data quality to mitigate unwanted LLM behavior, such as toxicity" ("Building LLM Applications with Databricks," 2023).

\* Option B: Implement rate limiting

\* Rate limiting controls request frequency, not content quality. It prevents overload but doesn't address toxicity in social media inputs or outputs.

\* Databricks Reference: Rate limiting is for performance, not safety: "Use rate limits to manage compute load" ("Generative AI Cookbook").

\* Option C: Reduce the amount of context items the system will include in consideration for its response

\* Reducing context might limit exposure to some toxic items but risks losing relevant information, and it doesn't specifically target toxicity. It's an indirect, imprecise fix.

\* Databricks Reference: Context reduction is for efficiency, not safety: "Adjust context size based on performance needs" ("Databricks Generative AI Engineer Guide").

\* Option D: Log all LLM system responses and perform a batch toxicity analysis monthly

\* Logging and analyzing responses is reactive, identifying toxicity after it occurs rather than preventing it. Monthly analysis doesn't limit real-time toxic outputs.

\* Databricks Reference: Monitoring is for auditing, not prevention: "Log outputs for post-hoc analysis, but use input filters for safety" ("Building LLM-Powered Applications").

Conclusion: Option A is the most effective guardrail, proactively filtering toxic inputs from unverified sources, which aligns with Databricks' emphasis on data quality as a primary safety mechanism for LLM systems.

## NEW QUESTION # 22

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. BGE-large
- D. Llama2-70b

**Answer: A**

Explanation:

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.

## NEW QUESTION # 23

.....

Pass4guide assists people in better understanding, studying, and passing more difficult certification exams. We take pride in successfully servicing industry experts by always delivering safe and dependable exam preparation materials. You will need authentic Databricks Databricks-Generative-AI-Engineer-Associate Exam Preparation material if you want to take the Databricks Certified Generative AI Engineer Associate exam to expand your career opportunities.

**Reliable Databricks-Generative-AI-Engineer-Associate Test Cost:** <https://www.pass4guide.com/Databricks-Generative-AI-Engineer-Associate-exam-guide-torrent.html>

- Excellect Databricks-Generative-AI-Engineer-Associate Pass Rate □ Latest Databricks-Generative-AI-Engineer-Associate Braindumps Pdf □ Databricks-Generative-AI-Engineer-Associate Exam Bible □ Immediately open ➔ [www.dumpsmaterials.com](http://www.dumpsmaterials.com) □□□ and search for { Databricks-Generative-AI-Engineer-Associate } to obtain a free download □ Databricks-Generative-AI-Engineer-Associate Reliable Braindumps Book
- Pass Guaranteed 2026 Authoritative Databricks New Databricks-Generative-AI-Engineer-Associate Exam Papers □ Search for ➤ Databricks-Generative-AI-Engineer-Associate ▲ and easily obtain a free download on ➔ [www.pdfvce.com](http://www.pdfvce.com) □□□ □ Databricks-Generative-AI-Engineer-Associate Exam Answers
- Latest Databricks-Generative-AI-Engineer-Associate Exam Vce □ Databricks-Generative-AI-Engineer-Associate Reliable Braindumps Book □ Exam Databricks-Generative-AI-Engineer-Associate Simulator Online □ Open □ [www.validtorrent.com](http://www.validtorrent.com) □ enter ➤ Databricks-Generative-AI-Engineer-Associate ▲ and obtain a free download □ Databricks-Generative-AI-Engineer-Associate Exam Answers
- Pass Guaranteed 2026 Authoritative Databricks New Databricks-Generative-AI-Engineer-Associate Exam Papers □ Search for ➤ Databricks-Generative-AI-Engineer-Associate ▲ and obtain a free download on □ [www.pdfvce.com](http://www.pdfvce.com) □□□ □ Databricks-Generative-AI-Engineer-Associate Training Online
- Databricks-Generative-AI-Engineer-Associate Exam Introduction □ Databricks-Generative-AI-Engineer-Associate Exam Introduction □ Databricks-Generative-AI-Engineer-Associate Exam Introduction □ Open website ➤ [www.vce4dumps.com](http://www.vce4dumps.com) □ and search for { Databricks-Generative-AI-Engineer-Associate } for free download □ Exam Databricks-Generative-AI-Engineer-Associate Simulator Online
- Trustworthy Databricks-Generative-AI-Engineer-Associate Source □ Databricks-Generative-AI-Engineer-Associate Training Online □ Databricks-Generative-AI-Engineer-Associate Reliable Study Plan □ Easily obtain free download of ➤ Databricks-Generative-AI-Engineer-Associate ▲ by searching on 「 [www.pdfvce.com](http://www.pdfvce.com) 」 ➔ Exam Databricks-Generative-AI-Engineer-Associate Simulator Online
- Exam Databricks-Generative-AI-Engineer-Associate Simulator Online □ New Databricks-Generative-AI-Engineer-Associate Dumps Book □ Databricks-Generative-AI-Engineer-Associate Exam Bible □ Download ➔ Databricks-Generative-AI-Engineer-Associate □ for free by simply searching on ➔ [www.practicevce.com](http://www.practicevce.com) □□□ □ Exam Databricks-Generative-AI-Engineer-Associate Simulator Online
- New Databricks-Generative-AI-Engineer-Associate Exam Papers - 2026 Databricks First-grade Reliable Databricks-Generative-AI-Engineer-Associate Test Cost □ Download ➔ Databricks-Generative-AI-Engineer-Associate □ for free by simply searching on “ [www.pdfvce.com](http://www.pdfvce.com) ” □ Exam Databricks-Generative-AI-Engineer-Associate Guide Materials
- Latest Databricks-Generative-AI-Engineer-Associate Exam Vce □ Valid Databricks-Generative-AI-Engineer-Associate Practice Materials □ Databricks-Generative-AI-Engineer-Associate Reliable Study Plan □ Search on ➔ [www.verifieddumps.com](http://www.verifieddumps.com) □ for □ Databricks-Generative-AI-Engineer-Associate □ to obtain exam materials for free download □ Exam Databricks-Generative-AI-Engineer-Associate Simulator Online
- Excellect Databricks-Generative-AI-Engineer-Associate Pass Rate □ Databricks-Generative-AI-Engineer-Associate Reliable Braindumps Book □ Databricks-Generative-AI-Engineer-Associate Exam Bible □ Search for ➔ Databricks-Generative-AI-Engineer-Associate □ and easily obtain a free download on 【 [www.pdfvce.com](http://www.pdfvce.com) 】 □ Databricks-Generative-AI-Engineer-Associate Exam Dumps Demo
- Exam Databricks-Generative-AI-Engineer-Associate Guide Materials □ Trustworthy Databricks-Generative-AI-Engineer-

Associate Source □ Databricks-Generative-AI-Engineer-Associate Reliable Braindumps Book □ Search for “Databricks-Generative-AI-Engineer-Associate” on ▷ [www.examcollectionpass.com](http://www.examcollectionpass.com) ↳ immediately to obtain a free download □ Latest Databricks-Generative-AI-Engineer-Associate Exam Vce

P.S. Free 2025 Databricks Databricks-Generative-AI-Engineer-Associate dumps are available on Google Drive shared by Pass4guide: [https://drive.google.com/open?id=1vrzplCK7kG1m33hb9y9a\\_DnCw68BWdOd1](https://drive.google.com/open?id=1vrzplCK7kG1m33hb9y9a_DnCw68BWdOd1)