

# **BraindumpQuiz Databricks-Generative-AI-Engineer-Associate Exam Questions are Verified by Subject Matter Experts**



DOWNLOAD the newest BraindumpQuiz Databricks-Generative-AI-Engineer-Associate PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1DZ2mZaYl7I6Bo3F-Xdw4YXhhZxBZMzV7>

For any candidate, choosing the Databricks-Generative-AI-Engineer-Associate question torrent material is the key to passing the exam. Our study materials can fully meet all your needs: Avoid wasting your time and improve your learning efficiency. Spending little hours per day within one week, you can pass the exam easily. You will don't take any risks and losses if you purchase and learn our Databricks-Generative-AI-Engineer-Associate Latest Exam Dumps, do you?

We have always set great store by superior after sale service, since we all tend to take responsibility for our customers who decide to choose our Databricks-Generative-AI-Engineer-Associate training materials. We pride ourselves on our industry-leading standards of customer care. Our worldwide after sale staffs will provide the most considerate after-sale service for you in twenty four hours a day, seven days a week, that is to say, no matter you are or whenever it is, as long as you have any question about our Databricks-Generative-AI-Engineer-Associate Exam Torrent or about the exam or even about the related certification, you can feel free to contact our after sale service staffs who will always waiting for you on the internet.

**>> Exam Databricks-Generative-AI-Engineer-Associate Braindumps <<**

**Databricks-Generative-AI-Engineer-Associate Valid Test Fee, Original**

## Databricks-Generative-AI-Engineer-Associate Questions

When you first contacted us with Databricks-Generative-AI-Engineer-Associate quiz torrent, you may be confused about our Databricks-Generative-AI-Engineer-Associate exam question and would like to learn more about our products to confirm our claims. We have a trial version for you to experience. If you choose to purchase our Databricks-Generative-AI-Engineer-Associate quiz torrent, you will have the right to get the update system and the update system is free of charge. We do not charge any additional fees. Once our Databricks-Generative-AI-Engineer-Associate Learning Materials are updated, we will automatically send you the latest information about our Databricks-Generative-AI-Engineer-Associate exam question. We assure you that our company will provide customers with a sustainable update system.

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

Topic	Details
Topic 1	<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 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>• 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 4	<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 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>

## Databricks Certified Generative AI Engineer Associate Sample Questions (Q22-Q27):

### NEW QUESTION # 22

A Generative AI Engineer is creating an LLM-based application. The documents for its retriever have been chunked to a maximum of 512 tokens each. The Generative AI Engineer knows that cost and latency are more important than quality for this application. They have several context length levels to choose from.

Which will fulfill their need?

- A. context length 32768; smallest model is 14GB and embedding dimension 4096
- B. context length 2048; smallest model is 11GB and embedding dimension 2560
- **C. context length 512; smallest model is 0.13GB and embedding dimension 384**
- D. context length 514; smallest model is 0.44GB and embedding dimension 768

**Answer: C**

Explanation:

When prioritizing cost and latency over quality in a Large Language Model (LLM)-based application, it is crucial to select a

configuration that minimizes both computational resources and latency while still providing reasonable performance. Here's why Dis is the best choice:

- \* Context length: The context length of 512 tokens aligns with the chunk size used for the documents (maximum of 512 tokens per chunk). This is sufficient for capturing the needed information and generating responses without unnecessary overhead.

- \* Smallest model size: The model with a size of 0.13GB is significantly smaller than the other options.

This small footprint ensures faster inference times and lower memory usage, which directly reduces both latency and cost.

- \* Embedding dimension: While the embedding dimension of 384 is smaller than the other options, it is still adequate for tasks where cost and speed are more important than precision and depth of understanding.

This setup achieves the desired balance between cost-efficiency and reasonable performance in a latency-sensitive, cost-conscious application.

### NEW QUESTION # 23

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. Increase the compute to improve processing speed of questions to allow greater relevancy analysis**
- C. Implement a profanity filter to screen out offensive language
- D. Implement a comprehensive guardrail framework that includes policies for content filters tailored to the finance sector.

**Answer: B**

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 # 24

A Generative AI Engineer received the following business requirements for an external chatbot.

The chatbot needs to know what types of questions the user asks and routes to appropriate models to answer the questions. For example, the user might ask about upcoming event details. Another user might ask about purchasing tickets for a particular event.

What is an ideal workflow for such a chatbot?

- A. There should be two different chatbots handling different types of user queries.
- B. The chatbot should only process payments
- **C. The chatbot should be implemented as a multi-step LLM workflow. First, identify the type of question asked, then route the question to the appropriate model. If it's an upcoming event question, send the query to a text-to-SQL model. If it's about ticket purchasing, the customer should be redirected to a payment platform.**
- D. The chatbot should only look at previous event information

**Answer: C**

Explanation:

- \* Problem Context: The chatbot must handle various types of queries and intelligently route them to the appropriate responses or systems.

\* Explanation of Options:

\* Option A: Limiting the chatbot to only previous event information restricts its utility and does not meet the broader business requirements.

\* Option B: Having two separate chatbots could unnecessarily complicate user interaction and increase maintenance overhead.

\* Option C: Implementing a multi-step workflow where the chatbot first identifies the type of question and then routes it accordingly is the most efficient and scalable solution. This approach allows the chatbot to handle a variety of queries dynamically, improving user experience and operational efficiency.

\* Option D: Focusing solely on payments would not satisfy all the specified user interaction needs, such as inquiring about event details.

Option C offers a comprehensive workflow that maximizes the chatbot's utility and responsiveness to different user needs, aligning perfectly with the business requirements.

## NEW QUESTION # 25

A Generative AI Engineer is creating an agent-based LLM system for their favorite monster truck team. The system can answer text based questions about the monster truck team, lookup event dates via an API call, or query tables on the team's latest standings. How could the Generative AI Engineer best design these capabilities into their system?

- A. Build a system prompt with all possible event dates and table information in the system prompt. Use a RAG architecture to lookup generic text questions and otherwise leverage the information in the system prompt.
- B. Ingest PDF documents about the monster truck team into a vector store and query it in a RAG architecture.
- C. Write a system prompt for the agent listing available tools and bundle it into an agent system that runs a number of calls to solve a query.
- D. Instruct the LLM to respond with "RAG", "API", or "TABLE" depending on the query, then use text parsing and conditional statements to resolve the query.

**Answer: C**

Explanation:

In this scenario, the Generative AI Engineer needs to design a system that can handle different types of queries about the monster truck team. The queries may involve text-based information, API lookups for event dates, or table queries for standings. The best solution is to implement a tool-based agent system.

Here's how option B works, and why it's the most appropriate answer:

\* System Design Using Agent-Based Model: In modern agent-based LLM systems, you can design a system where the LLM (Large Language Model) acts as a central orchestrator. The model can "decide" which tools to use based on the query. These tools can include API calls, table lookups, or natural language searches. The system should contain a system prompt that informs the LLM about the available tools.

\* System Prompt Listing Tools: By creating a well-crafted system prompt, the LLM knows which tools are at its disposal. For instance, one tool may query an external API for event dates, another might look up standings in a database, and a third may involve searching a vector database for general text-based information. The agent will be responsible for calling the appropriate tool depending on the query.

\* Agent Orchestration of Calls: The agent system is designed to execute a series of steps based on the incoming query. If a user asks for the next event date, the system will recognize this as a task that requires an API call. If the user asks about standings, the agent might query the appropriate table in the database. For text-based questions, it may call a search function over ingested data. The agent orchestrates this entire process, ensuring the LLM makes calls to the right resources dynamically.

\* Generative AI Tools and Context: This is a standard architecture for integrating multiple functionalities into a system where each query requires different actions. The core design in option B is efficient because it keeps the system modular and dynamic by leveraging tools rather than overloading the LLM with static information in a system prompt (like option D).

\* Why Other Options Are Less Suitable:

\* A (RAG Architecture): While relevant, simply ingesting PDFs into a vector store only helps with text-based retrieval. It wouldn't help with API lookups or table queries.

\* C (Conditional Logic with RAG/API/TABLE): Although this approach works, it relies heavily on manual text parsing and might introduce complexity when scaling the system.

\* D (System Prompt with Event Dates and Standings): Hardcoding dates and table information into a system prompt isn't scalable. As the standings or events change, the system would need constant updating, making it inefficient.

By bundling multiple tools into a single agent-based system (as in option B), the Generative AI Engineer can best handle the diverse requirements of this system.

## NEW QUESTION # 26

A Generative AI Engineer is developing a patient-facing healthcare-focused chatbot. If the patient's question is not a medical emergency, the chatbot should solicit more information from the patient to pass to the doctor's office and suggest a few relevant pre-approved medical articles for reading. If the patient's question is urgent, direct the patient to calling their local emergency services.

Given the following user input:

"I have been experiencing severe headaches and dizziness for the past two days." Which response is most appropriate for the chatbot to generate?

- **A. Please call your local emergency services.**
- B. Please provide your age, recent activities, and any other symptoms you have noticed along with your headaches and dizziness.
- C. Headaches can be tough. Hope you feel better soon!
- D. Here are a few relevant articles for your browsing. Let me know if you have questions after reading them.

**Answer: A**

Explanation:

\* Problem Context: The task is to design responses for a healthcare-focused chatbot that appropriately addresses the urgency of a patient's symptoms.

\* Explanation of Options:

\* Option A: Suggesting articles might be suitable for less urgent inquiries but is inappropriate for symptoms that could indicate a serious condition.

\* Option B: Given the description of severe symptoms like headaches and dizziness, directing the patient to emergency services is prudent. This aligns with medical guidelines that recommend immediate professional attention for such severe symptoms.

\* Option C: Offering well-wishes does not address the potential seriousness of the symptoms and lacks appropriate action.

\* Option D: While gathering more information is part of a detailed assessment, the immediate need here suggests a more urgent response.

Given the potential severity of the described symptoms, Option B is the most appropriate, ensuring the chatbot directs patients to seek urgent care when needed, potentially saving lives.


## NEW QUESTION # 27

.....

Getting tired of humdrum life, you may want to get some successful feeling or try something different instead. We all know that is of important to pass the Databricks-Generative-AI-Engineer-Associate exam and get the Databricks-Generative-AI-Engineer-Associate certification for someone who wants to find a good job in internet area, and it is not a simple thing to prepare for exam. So you are in the right place now. The Databricks-Generative-AI-Engineer-Associate practice materials are a great beginning to prepare your exam. Actually, just think of our Databricks practice materials as the best way to pass the exam is myopic. They can not only achieve this, but ingeniously help you remember more content at the same time.

**Databricks-Generative-AI-Engineer-Associate Valid Test Fee:** <https://www.braindumpquiz.com/Databricks-Generative-AI-Engineer-Associate-exam-material.html>

- Databricks-Generative-AI-Engineer-Associate Practice Test Engine □ Databricks-Generative-AI-Engineer-Associate Practice Test Engine □ Databricks-Generative-AI-Engineer-Associate Exam Cram □ Open website ⇒ [www.prep4pass.com](http://www.prep4pass.com) ⇐ and search for ( Databricks-Generative-AI-Engineer-Associate ) for free download □ □Databricks-Generative-AI-Engineer-Associate Exam Topics Pdf
- Databricks-Generative-AI-Engineer-Associate Practice Test Engine □ New Databricks-Generative-AI-Engineer-Associate Test Fee □ Databricks-Generative-AI-Engineer-Associate Latest Dumps Files □ Enter { [www.pdfvce.com](http://www.pdfvce.com) } and search for ► Databricks-Generative-AI-Engineer-Associate ◄ to download for free □Databricks-Generative-AI-Engineer-Associate Valid Dumps Demo
- Free PDF Databricks - Databricks-Generative-AI-Engineer-Associate - Databricks Certified Generative AI Engineer Associate Pass-Sure Exam Braindumps □ Simply search for □ Databricks-Generative-AI-Engineer-Associate □ for free download on ☼ [www.prep4away.com](http://www.prep4away.com) □☼ □ □Certification Databricks-Generative-AI-Engineer-Associate Exam
- Databricks-Generative-AI-Engineer-Associate Valid Dumps Demo □ PDF Databricks-Generative-AI-Engineer-Associate VCE □ Databricks-Generative-AI-Engineer-Associate Braindumps Pdf □ Easily obtain “ Databricks-Generative-AI-Engineer-Associate ” for free download through ⇒ [www.pdfvce.com](http://www.pdfvce.com) ⇐ □Valid Databricks-Generative-AI-Engineer-Associate Study Materials
- Get instant Success With Databricks Databricks-Generative-AI-Engineer-Associate Exam Questions [2025] □ Copy URL ☼ [www.testsdumps.com](http://www.testsdumps.com) □☼ □ open and search for □ Databricks-Generative-AI-Engineer-Associate □ to download for

free  Databricks-Generative-AI-Engineer-Associate Exam

- Databricks-Generative-AI-Engineer-Associate Latest Dumps Files □ Download Databricks-Generative-AI-Engineer-Associate Pdf & Databricks-Generative-AI-Engineer-Associate Exam Discount Voucher □ Simply search for 「 Databricks-Generative-AI-Engineer-Associate 」 for free download on ⇒ [www.pdfvce.com](http://www.pdfvce.com) ⇐ □ PDF Databricks-Generative-AI-Engineer-Associate VCE
- Databricks Databricks-Generative-AI-Engineer-Associate Practice Test - The Secret To Overcome Exam Anxiety □ □ [www.prep4away.com](http://www.prep4away.com) □ is best website to obtain 《 Databricks-Generative-AI-Engineer-Associate 》 for free download □ □ Databricks-Generative-AI-Engineer-Associate Exam Topics Pdf
- Databricks Databricks-Generative-AI-Engineer-Associate Questions - Pass Exam With Ease (2025) □ Search for ( Databricks-Generative-AI-Engineer-Associate ) on ► [www.pdfvce.com](http://www.pdfvce.com) ◀ immediately to obtain a free download □ □ Databricks-Generative-AI-Engineer-Associate Exam Topics Pdf
- Certification Databricks-Generative-AI-Engineer-Associate Exam ↗ Databricks-Generative-AI-Engineer-Associate Braindumps Pdf □ Valid Databricks-Generative-AI-Engineer-Associate Study Materials □ Search for ➡ Databricks-Generative-AI-Engineer-Associate □ and download exam materials for free through ➡ [www.testsimulate.com](http://www.testsimulate.com) □ □ □ Certification Databricks-Generative-AI-Engineer-Associate Exam
- Printable Databricks-Generative-AI-Engineer-Associate PDF ♥ □ Valid Databricks-Generative-AI-Engineer-Associate Study Materials □ Certification Databricks-Generative-AI-Engineer-Associate Exam □ Search for ➡ Databricks-Generative-AI-Engineer-Associate □ and easily obtain a free download on ⇒ [www.pdfvce.com](http://www.pdfvce.com) ⇐ □ Guaranteed Databricks-Generative-AI-Engineer-Associate Passing
- Top Exam Databricks-Generative-AI-Engineer-Associate Braindumps | Valid Databricks-Generative-AI-Engineer-Associate: Databricks Certified Generative AI Engineer Associate 100% Pass □ Enter □ [www.itcerttest.com](http://www.itcerttest.com) □ and search for [ Databricks-Generative-AI-Engineer-Associate ] to download for free □ Databricks-Generative-AI-Engineer-Associate Valid Dumps Demo
- [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [myportal.utt.edu.tt](http://myportal.utt.edu.tt), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [fxsensei.top](http://fxsensei.top), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), [www.stes.tyc.edu.tw](http://www.stes.tyc.edu.tw), Disposable vapes

2025 Latest BraindumpQuiz Databricks-Generative-AI-Engineer-Associate PDF Dumps and Databricks-Generative-AI-Engineer-Associate Exam Engine Free Share: <https://drive.google.com/open?id=1DZ2mZaYl7I6B03F-Xdw4YXhhZxBZMzV7>