

# **Databricks-Generative-AI-Engineer-Associate Valid Test Vce Free & Databricks-Generative-AI-Engineer-Associate Latest Dumps Ppt**



What's more, part of that Prep4sureGuide Databricks-Generative-AI-Engineer-Associate dumps now are free:  
<https://drive.google.com/open?id=1yOAIcYs-Pkx-Hj6LGbAxW9iUBFd4WQRM>

Nowadays in this information-based world the definition of the talents mean that the personnel boost both the knowledge in Databricks-Generative-AI-Engineer-Associate area and the practical abilities now. So if you want to be the talent the society actually needs you must apply your knowledge into the practical working and passing the test Databricks-Generative-AI-Engineer-Associate Certification can make you become the talent the society needs. If you buy our Databricks-Generative-AI-Engineer-Associate study materials you will pass the exam successfully and realize your goal to be the talent.

In every area, timing counts importantly. With the advantage of high efficiency, our Databricks-Generative-AI-Engineer-Associate practice materials help you avoid wasting time on selecting the important and precise content from the broad information. In such a way, you can confirm that you get the convenience and fast. By studying with our Databricks-Generative-AI-Engineer-Associate Real Exam for 20 to 30 hours, we can claim that you can get ready to attend the Databricks-Generative-AI-Engineer-Associate exam.

**>> Databricks-Generative-AI-Engineer-Associate Valid Test Vce Free <<**

## **Databricks-Generative-AI-Engineer-Associate Latest Dumps Ppt - Databricks-Generative-AI-Engineer-Associate New Study Guide**

Our goal is to help you save both time and money by providing you with the Databricks-Generative-AI-Engineer-Associate updated exam questions. Keep up the good work on preparing for the Databricks Databricks-Generative-AI-Engineer-Associate test with our actual Databricks Databricks-Generative-AI-Engineer-Associate Dumps. We are so confident that you will succeed on the first try that we will return your money according to the terms and conditions if you do not.

## **Databricks Certified Generative AI Engineer Associate Sample Questions (Q55-Q60):**

### **NEW QUESTION # 55**

A Generative AI Engineer developed an LLM application using the provisioned throughput Foundation Model API. Now that the application is ready to be deployed, they realize their volume of requests are not sufficiently high enough to create their own provisioned throughput endpoint. They want to choose a strategy that ensures the best cost-effectiveness for their application. What strategy should the Generative AI Engineer use?

- A. Deploy the model using pay-per-token throughput as it comes with cost guarantees
- B. Switch to using External Models instead
- C. Change to a model with a fewer number of parameters in order to reduce hardware constraint issues
- D. Throttle the incoming batch of requests manually to avoid rate limiting issues

### Answer: A

Explanation:

\* Problem Context: The engineer needs a cost-effective deployment strategy for an LLM application with relatively low request volume.

\* Explanation of Options:

\* Option A: Switching to external models may not provide the required control or integration necessary for specific application needs.

\* Option B: Using a pay-per-token model is cost-effective, especially for applications with variable or low request volumes, as it aligns costs directly with usage.

\* Option C: Changing to a model with fewer parameters could reduce costs, but might also impact the performance and capabilities of the application.

\* Option D: Manually throttling requests is a less efficient and potentially error-prone strategy for managing costs.

Option B is ideal, offering flexibility and cost control, aligning expenses directly with the application's usage patterns.

### NEW QUESTION # 56

A Generative AI Engineer has been asked to build an LLM-based question-answering application. The application should take into account new documents that are frequently published. The engineer wants to build this application with the least cost and least development effort and have it operate at the lowest cost possible.

Which combination of chaining components and configuration meets these requirements?

- A. For the application a prompt, an agent and a fine-tuned LLM are required. The agent is used by the LLM to retrieve relevant content that is inserted into the prompt which is given to the LLM to generate answers.
- B. For the application a prompt, a retriever, and an LLM are required. The retriever output is inserted into the prompt which is given to the LLM to generate answers.
- C. For the question-answering application, prompt engineering and an LLM are required to generate answers.
- D. The LLM needs to be frequently updated with new documents in order to provide most up-to-date answers.

### Answer: B

Explanation:

Problem Context: The task is to build an LLM-based question-answering application that integrates new documents frequently with minimal costs and development efforts.

Explanation of Options:

\* Option A: Utilizes a prompt and a retriever, with the retriever output being fed into the LLM. This setup is efficient because it dynamically updates the data pool via the retriever, allowing the LLM to provide up-to-date answers based on the latest documents without needing to frequently retrain the model. This method offers a balance of cost-effectiveness and functionality.

\* Option B: Requires frequent retraining of the LLM, which is costly and labor-intensive.

\* Option C: Only involves prompt engineering and an LLM, which may not adequately handle the requirement for incorporating new documents unless it's part of an ongoing retraining or updating mechanism, which would increase costs.

\* Option D: Involves an agent and a fine-tuned LLM, which could be overkill and lead to higher development and operational costs. Option A is the most suitable as it provides a cost-effective, minimal development approach while ensuring the application remains up-to-date with new information.

### 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. Implement a comprehensive guardrail framework that includes policies for content filters tailored to the finance sector.
- B. Incorporate manual reviews to correct any problematic outputs prior to sending to the users
- 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 has built an LLM-based system that will automatically translate user text between two languages. They now want to benchmark multiple LLM's on this task and pick the best one. They have an evaluation set with known high quality translation examples. They want to evaluate each LLM using the evaluation set with a performant metric.

Which metric should they choose for this evaluation?

- A. RECALL metric
- B. ROUGE metric
- C. NDCG metric
- D. BLEU metric

### Answer: D

Explanation:

The task is to benchmark LLMs for text translation using an evaluation set with known high-quality examples, requiring a performant metric. Let's evaluate the options.

\* Option A: ROUGE metric

\* ROUGE (Recall-Oriented Understudy for Gisting Evaluation) measures overlap between generated and reference texts, primarily for summarization. It's less suited for translation, where precision and word order matter more.

\* Databricks Reference: "ROUGE is commonly used for summarization, not translation evaluation" ("Generative AI Cookbook," 2023).

\* Option B: BLEU metric

\* BLEU (Bilingual Evaluation Understudy) evaluates translation quality by comparing n-gram overlap with reference translations, accounting for precision and brevity. It's widely used, performant, and appropriate for this task.

\* Databricks Reference: "BLEU is a standard metric for evaluating machine translation, balancing accuracy and efficiency" ("Building LLM Applications with Databricks").

\* Option C: NDCG metric

\* NDCG (Normalized Discounted Cumulative Gain) assesses ranking quality, not text generation.

It's irrelevant for translation evaluation.

\* Databricks Reference: "NDCG is suited for ranking tasks, not generative output scoring" ("Databricks Generative AI Engineer Guide").

\* Option D: RECALL metric

\* Recall measures retrieved relevant items but doesn't evaluate translation quality (e.g., fluency, correctness). It's incomplete for this use case.

\* Databricks Reference: No specific extract, but recall alone lacks the granularity of BLEU for text generation tasks.

Conclusion: Option B (BLEU) is the best metric for translation evaluation, offering a performant and standard approach, as endorsed by Databricks' guidance on generative tasks.

## NEW QUESTION # 59

A Generative AI Engineer is using an LLM to classify species of edible mushrooms based on text descriptions of certain features. The model is returning accurate responses in testing and the Generative AI Engineer is confident they have the correct list of possible labels, but the output frequently contains additional reasoning in the answer when the Generative AI Engineer only wants to return the label with no additional text.

Which action should they take to elicit the desired behavior from this LLM?

- A. Use zero shot chain-of-thought prompting to prevent a verbose output format
- **B. Use a system prompt to instruct the model to be succinct in its answer**
- C. Use zero shot prompting to instruct the model on expected output format
- D. Use few shot prompting to instruct the model on expected output format

**Answer: B**

Explanation:

The LLM classifies mushroom species accurately but includes unwanted reasoning text, and the engineer wants only the label. Let's assess how to control output format effectively.

- \* Option A: Use few shot prompting to instruct the model on expected output format
- \* Few-shot prompting provides examples (e.g., input: description, output: label). It can work but requires crafting multiple examples, which is effort-intensive and less direct than a clear instruction.
- \* Databricks Reference: "Few-shot prompting guides LLMs via examples, effective for format control but requires careful design" ("Generative AI Cookbook").
- \* Option B: Use zero shot prompting to instruct the model on expected output format
- \* Zero-shot prompting relies on a single instruction (e.g., "Return only the label") without examples. It's simpler than few-shot but may not consistently enforce succinctness if the LLM's default behavior is verbose.
- \* Databricks Reference: "Zero-shot prompting can specify output but may lack precision without examples" ("Building LLM Applications with Databricks").
- \* Option C: Use zero shot chain-of-thought prompting to prevent a verbose output format
- \* Chain-of-Thought (CoT) encourages step-by-step reasoning, which increases verbosity—opposite to the desired outcome. This contradicts the goal of label-only output.
- \* Databricks Reference: "CoT prompting enhances reasoning but often results in detailed responses" ("Databricks Generative AI Engineer Guide").
- \* Option D: Use a system prompt to instruct the model to be succinct in its answer

\* A system prompt (e.g., "Respond with only the species label, no additional text") sets a global instruction for the LLM's behavior. It's direct, reusable, and effective for controlling output style across queries.

\* Databricks Reference: "System prompts define LLM behavior consistently, ideal for enforcing concise outputs" ("Generative AI Cookbook," 2023).

Conclusion: Option D is the most effective and straightforward action, using a system prompt to enforce succinct, label-only responses, aligning with Databricks' best practices for output control.

## NEW QUESTION # 60

.....

We update our Databricks-Generative-AI-Engineer-Associate test prep within one year and you will download free which you need. After one year, we provide the client 50% discount benefit if buyers want to extend their service warranty so you can save much money. If you are the old client, you can enjoy some certain discount when buying Databricks-Generative-AI-Engineer-Associate Exam Torrent so you can enjoy more service and more benefits. Our update can provide the latest and most useful Databricks-Generative-AI-Engineer-Associate prep torrent to you and you can learn more and pass the Databricks-Generative-AI-Engineer-Associate exam successfully.

**Databricks-Generative-AI-Engineer-Associate Latest Dumps Ppt:** <https://www.prep4sureguide.com/Databricks-Generative-AI-Engineer-Associate-prep4sure-exam-guide.html>

If you study with our Databricks-Generative-AI-Engineer-Associate practice engine for 20 to 30 hours, we can claim that you can pass the exam as easy as a pie. The latest Databricks-Generative-AI-Engineer-Associate dumps pdf covers every topic of the certification exam and contains the latest test questions and answers. As the top company in IT field many companies regard Databricks-Generative-AI-Engineer-Associate certification as one of products manage elite standards in most of countries. Our Databricks-Generative-AI-Engineer-Associate exam torrent has three versions which people can choose according to their actual needs.

NavCreateGetFileDialog creates an Open dialog box, but it doesn't display Databricks-Generative-AI-Engineer-Associate or control it. According to the article, Rethink's robots simply didn't solve business problems as well as robots from competitors.

## Learning Material In 3 Different Formats for Databricks Databricks-Generative-AI-Engineer-Associate Exam Success

If you study with our Databricks-Generative-AI-Engineer-Associate Practice Engine for 20 to 30 hours, we can claim that you can pass the exam as easy as a pie. The latest Databricks-Generative-AI-Engineer-Associate dumps pdf covers every topic of the certification exam and contains the latest test questions and answers.

As the top company in IT field many companies regard Databricks-Generative-AI-Engineer-Associate certification as one of products manage elite standards in most of countries, Our Databricks-Generative-AI-Engineer-Associate exam torrent has three versions which people can choose according to their actual needs.

In the current era of rocketing development Databricks-Generative-AI-Engineer-Associate Valid Test Vce Free of the whole society, it's easy to be eliminated if people have just a single skill.

P.S. Free & New Databricks-Generative-AI-Engineer-Associate dumps are available on Google Drive shared by Prep4sureGuide: <https://drive.google.com/open?id=1yOAiCYs-Pkx-Hj6LGbAxW9iUBFd4WQRM>

