

Latest AIF-C01 Mock Exam & Braindump AIF-C01 Free



Amazon (AWS) : AIF-C01

AWS Certified AI Practitioner

Latest Updated : Q & A

Get Full Pro PDF

Show More

Full Exam Practice

Amazon (AWS) : AIF-C01 AWS Certified AI Practitioner

P.S. Free 2026 Amazon AIF-C01 dumps are available on Google Drive shared by ActualPDF: https://drive.google.com/open?id=1qcFJLfm_5UU6w7PDU44ggxW6TGirT0yT

Do you want to pass your exam just one time? If you do, then you can choose us, we can help you pass the exam just one time. With experienced experts to compile and verify AIF-C01 training materials, the quality can be guaranteed. We also pass guarantee and money back guarantee if you fail to pass the exam. You can obtain the download link and password for AIF-C01 Exam Dumps within ten minutes, so that you can start your learning immediately. We have online and offline service, and the staff possess the professional knowledge for AIF-C01 exam dumps, if you have any questions, you can have a conversation with us.

Amazon AIF-C01 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none">• Fundamentals of Generative AI: This domain explores the basics of generative AI, focusing on techniques for creating new content from learned patterns, including text and image generation. It targets professionals interested in understanding generative models, such as developers and researchers in AI.
Topic 2	<ul style="list-style-type: none">• Applications of Foundation Models: This domain examines how foundation models, like large language models, are used in practical applications. It is designed for those who need to understand the real-world implementation of these models, including solution architects and data engineers who work with AI technologies to solve complex problems.
Topic 3	<ul style="list-style-type: none">• Fundamentals of AI and ML: This domain covers the fundamental concepts of artificial intelligence (AI) and machine learning (ML), including core algorithms and principles. It is aimed at individuals new to AI and ML, such as entry-level data scientists and IT professionals.
Topic 4	<ul style="list-style-type: none">• Guidelines for Responsible AI: This domain highlights the ethical considerations and best practices for deploying AI solutions responsibly, including ensuring fairness and transparency. It is aimed at AI practitioners, including data scientists and compliance officers, who are involved in the development and deployment of AI systems and need to adhere to ethical standards.
Topic 5	<ul style="list-style-type: none">• Security, Compliance, and Governance for AI Solutions: This domain covers the security measures, compliance requirements, and governance practices essential for managing AI solutions. It targets security professionals, compliance officers, and IT managers responsible for safeguarding AI systems, ensuring regulatory compliance, and implementing effective governance frameworks.

>> Latest AIF-C01 Mock Exam <<

Braindump AIF-C01 Free - Passing AIF-C01 Score

The price for AIF-C01 exam torrent is quite reasonable, you can afford it no matter you are a student or you are an employee in the company. You just need to spend some money, and you can get a certificate. In addition, AIF-C01 exam dumps are high-quality and accuracy, and you can pass the exam successfully by using them. We also pass guarantee and money back guarantee for your failure of the exam after using AIF-C01 Exam Dumps. We offer you free update for 365 days after purchasing, and the update version will be sent to your email address automatically.

Amazon AWS Certified AI Practitioner Sample Questions (Q154-Q159):

NEW QUESTION # 154

A company wants to create a chatbot that answers questions about human resources policies. The company is using a large language model (LLM) and has a large digital documentation base.

Which technique should the company use to optimize the generated responses?

- **A. Use Retrieval Augmented Generation (RAG).**
- B. Set the temperature to 1.
- C. Decrease the token size.
- D. Use few-shot prompting.

Answer: A

Explanation:

The company is building a chatbot using an LLM to answer questions about HR policies, with access to a large digital documentation base. Retrieval Augmented Generation (RAG) optimizes the LLM's responses by retrieving relevant information from the documentation base and using it to generate accurate, contextually grounded answers, reducing hallucinations and improving response quality.

Exact Extract from AWS AI Documents:

From the AWS Bedrock User Guide:

"Retrieval Augmented Generation (RAG) enhances the performance of large language models by retrieving relevant information from external knowledge bases, such as documentation or databases, and incorporating it into the generation process. This technique ensures responses are accurate and grounded in the provided data, making it ideal for applications like policy chatbots." (Source: AWS Bedrock User Guide, Retrieval Augmented Generation) Detailed Explanation:

- * Option A: Use Retrieval Augmented Generation (RAG). This is the correct answer. RAG leverages the documentation base to provide the LLM with relevant HR policy information, optimizing the chatbot's responses for accuracy and relevance.
- * Option B: Use few-shot prompting. Few-shot prompting provides a few examples in the prompt to guide the LLM, but it is less effective than RAG for large documentation bases, as it cannot dynamically retrieve specific policy details.
- * Option C: Set the temperature to 1. Setting the temperature to 1 controls the randomness of the LLM's output but does not optimize responses using external documentation. This option is unrelated to the documentation base.
- * Option D: Decrease the token size. Decreasing token size (likely referring to limiting input/output tokens) may reduce response length but does not optimize the quality of responses using the documentation base.

References:

AWS Bedrock User Guide: Retrieval Augmented Generation (<https://docs.aws.amazon.com/bedrock/latest/userguide/rag.html>)

AWS AI Practitioner Learning Path: Module on Generative AI Optimization Amazon Bedrock Developer Guide: Building Policy Chatbots (<https://aws.amazon.com/bedrock/>)

NEW QUESTION # 155

An AI practitioner has trained a model on a training dataset. The model performs well on the training data.

However, the model does not perform well on evaluation data. What is the MOST likely cause of this issue?

- **A. The model is overfit.**
- B. The model requires prompt engineering.
- C. The model is underfit.
- D. The model is biased.

Answer: A

Explanation:

When a model performs well on training data but poorly on evaluation/test data, it indicates overfitting.

Overfitting: The model memorizes the training data patterns instead of generalizing.

Underfitting (A) means the model performs poorly on both training and test data.

Bias (C) refers to systemic errors in predictions, not this training/test mismatch.
Prompt engineering (B) applies to generative AI, not general ML training models.
Reference:
AWS ML Glossary - Overfitting and Underfitting

NEW QUESTION # 156

A software company wants to use a large language model (LLM) for workflow automation. The application will transform user messages into JSON files. The company will use the JSON files as inputs for data pipelines. The company has a labeled dataset that contains user messages and output JSON files. Which solution will train the LLM for workflow automation?

- A. Fine-tuning
- B. Unsupervised learning
- C. Reinforcement learning from human feedback (RLHF)
- D. Continued pre-training

Answer: A

Explanation:

Fine-tuning is the process of training a pre-trained LLM with a labeled dataset specific to a desired task—in this case, mapping user messages to JSON outputs. Fine-tuning leverages supervised learning to specialize the model's outputs.

C is correct:

"Fine-tuning is a supervised learning approach in which a model is further trained on a custom, labeled dataset to adapt to a specific use case." (Reference: Amazon Bedrock Fine-Tuning, AWS Certified AI Practitioner Study Guide)

"Fine-tuning is a supervised learning approach in which a model is further trained on a custom, labeled dataset to adapt to a specific use case." (Reference: Amazon Bedrock Fine-Tuning, AWS Certified AI Practitioner Study Guide) A is incorrect—unsupervised learning does not use labeled data.

B (continued pre-training) uses unlabeled data.

D (RLHF) uses reward signals and human feedback, not direct labeled input/output pairs.

NEW QUESTION # 157

A company is exploring Amazon Nova models in Amazon Bedrock. The company needs a multimodal model that supports multiple languages.

- A. Nova Canvas
- B. Nova Lite
- C. Nova Reel
- D. Nova Pro

Answer: D

Explanation:

* Amazon Nova Pro is a multimodal foundation model in Amazon Bedrock that supports text, images, and multiple languages.

* Nova Lite is optimized for lightweight, faster inference at lower cost.

* Nova Canvas is a creative tool for visual design.

* Nova Reel is optimized for video-related use cases.

Reference:

AWS Documentation - Amazon Nova Models

NEW QUESTION # 158

An AI practitioner is using an Amazon Bedrock base model to summarize session chats from the customer service department. The AI practitioner wants to store invocation logs to monitor model input and output data.

- A. Configure AWS Audit Manager as the logs destination for the model.
- B. Configure AWS CloudTrail as the logs destination for the model.
- C. Enable model invocation logging in Amazon Bedrock.
- D. Configure model invocation logging in Amazon EventBridge.

