



AIF-C01^{Q&As}

Amazon AWS Certified AI Practitioner

Pass Amazon AIF-C01 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.geekcert.com/aif-c01.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

A company wants to use generative AI to increase developer productivity and software development. The company wants to use Amazon Q Developer.

What can Amazon Q Developer do to help the company meet these requirements?

- A. Create software snippets, reference tracking, and open-source license tracking.
- B. Run an application without provisioning or managing servers.
- C. Enable voice commands for coding and providing natural language search.
- D. Convert audio files to text documents by using ML models.

Correct Answer: A

Amazon Q Developer is a tool designed to assist developers in increasing productivity by generating code snippets, managing reference tracking, and handling open-source license tracking. These features help developers by automating

parts of the software development process.

Option A (Correct): "Create software snippets, reference tracking, and open- source license tracking":This is the correct answer because these are key features that help developers streamline and automate tasks, thus improving productivity.

Option B:"Run an application without provisioning or managing servers" is incorrect as it refers to AWS Lambda or AWS Fargate, not Amazon Q Developer. Option C:"Enable voice commands for coding and providing natural language

search" is incorrect because this is not a function of Amazon Q Developer. Option D:"Convert audio files to text documents by using ML models" is incorrect as this refers to Amazon Transcribe, not Amazon Q Developer.

AWS AI Practitioner References:

Amazon Q Developer Features:AWS documentation outlines how Amazon Q Developer supports developers by offering features that reduce manual effort and improve efficiency.

QUESTION 2

An AI practitioner has built a deep learning model to classify the types of materials in images. The AI practitioner now wants to measure the model performance.

Which metric will help the AI practitioner evaluate the performance of the model?

- A. Confusion matrix
- B. Correlation matrix
- C. R2 score
- D. Mean squared error (MSE)



Correct Answer: A

A confusion matrix is the correct metric for evaluating the performance of a classification model, such as the deep learning model built to classify types of materials in images.

QUESTION 3

A company has a database of petabytes of unstructured data from internal sources. The company wants to transform this data into a structured format so that its data scientists can perform machine learning (ML) tasks.

Which service will meet these requirements?

- A. Amazon Lex
- B. Amazon Rekognition
- C. Amazon Kinesis Data Streams
- D. AWS Glue

Correct Answer: D

AWS Glue is the correct service for transforming petabytes of unstructured data into a structured format suitable for machine learning tasks.

QUESTION 4

A company wants to classify human genes into 20 categories based on gene characteristics. The company needs an ML algorithm to document how the inner mechanism of the model affects the output.

Which ML algorithm meets these requirements?

- A. Decision trees
- B. Linear regression
- C. Logistic regression
- D. Neural networks

Correct Answer: A

Decision trees are an interpretable machine learning algorithm that clearly documents the decision-making process by showing how each input feature affects the output. This transparency is particularly useful when explaining how the model

arrives at a certain decision, making it suitable for classifying genes into categories. Option A (Correct): "Decision trees": This is the correct answer because decision trees provide a clear and interpretable representation of how input features

influence the model's output, making it ideal for understanding the inner mechanisms affecting predictions.

Option B: "Linear regression" is incorrect because it is used for regression tasks, not classification.



Option C: "Logistic regression" is incorrect as it does not provide the same level of interpretability in documenting decision-making processes. Option D: "Neural networks" is incorrect because they are often considered "black boxes" and do not

easily explain how they arrive at their outputs.

AWS AI Practitioner References:

Interpretable Machine Learning Models on AWS: AWS supports using interpretable models, such as decision trees, for tasks that require clear documentation of how input data affects output decisions.

QUESTION 5

A company wants to use language models to create an application for inference on edge devices. The inference must have the lowest latency possible.

Which solution will meet these requirements?

- A. Deploy optimized small language models (SLMs) on edge devices.
- B. Deploy optimized large language models (LLMs) on edge devices.
- C. Incorporate a centralized small language model (SLM) API for asynchronous communication with edge devices.
- D. Incorporate a centralized large language model (LLM) API for asynchronous communication with edge devices.

Correct Answer: A

To achieve the lowest latency possible for inference on edge devices, deploying optimized small language models (SLMs) is the most effective solution. SLMs require fewer resources and have faster inference times, making them ideal for deployment on edge devices where processing power and memory are limited. Option A (Correct): "Deploy optimized small language models (SLMs) on edge devices": This is the correct answer because SLMs provide fast inference with low latency, which is crucial for edge deployments. Option B: "Deploy optimized large language models (LLMs) on edge devices" is incorrect because LLMs are resource-intensive and may not perform well on edge devices due to their size and computational demands. Option C: "Incorporate a centralized small language model (SLM) API for asynchronous communication with edge devices" is incorrect because it introduces network latency due to the need for communication with a centralized server. Option D: "Incorporate a centralized large language model (LLM) API for asynchronous communication with edge devices" is incorrect for the same reason, with even greater latency due to the larger model size. AWS AI Practitioner References: Optimizing AI Models for Edge Devices on AWS: AWS recommends using small, optimized models for edge deployments to ensure minimal latency and efficient performance.