

## AIF-C01<sup>Q&As</sup>

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.leads4pass.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 has installed a security camera. The company uses an ML model to evaluate the security camera footage for potential thefts. The company has discovered that the model disproportionately flags people who are members of a specific ethnic group. Which type of bias is affecting the model output?

- A. Measurement bias
- B. Sampling bias
- C. Observer bias
- D. Confirmation bias

Correct Answer: B

Sampling bias is the correct type of bias affecting the model output when it disproportionately flags people from a specific ethnic group.

---

**QUESTION 2**

Which option is a use case for generative AI models?

- A. Improving network security by using intrusion detection systems
- B. Creating photorealistic images from text descriptions for digital marketing
- C. Enhancing database performance by using optimized indexing
- D. Analyzing financial data to forecast stock market trends

Correct Answer: B

Generative AI models are used to create new content based on existing data. One common use case is generating photorealistic images from text descriptions, which is particularly useful in digital marketing, where visual content is key to

engaging potential customers.

Option B (Correct): "Creating photorealistic images from text descriptions for digital marketing": This is the correct answer because generative AI models, like those offered by Amazon Bedrock, can create images based on text descriptions,

making them highly valuable for generating marketing materials. Option A: "Improving network security by using intrusion detection systems" is incorrect because this is a use case for traditional machine learning models, not generative AI.

Option C: "Enhancing database performance by using optimized indexing" is incorrect as it is unrelated to generative AI.

Option D: "Analyzing financial data to forecast stock market trends" is incorrect because it typically involves predictive modeling rather than generative AI.

AWS AI Practitioner

**References:**

Use Cases for Generative AI Models on AWS: AWS highlights the use of generative AI for creative content generation, including image creation, text generation, and more, which is suited for digital marketing applications.

---

**QUESTION 3**

A company has built a chatbot that can respond to natural language questions with images. The company wants to ensure that the chatbot does not return inappropriate or unwanted images.

Which solution will meet these requirements?

- A. Implement moderation APIs.
- B. Retrain the model with a general public dataset.
- C. Perform model validation.
- D. Automate user feedback integration.

Correct Answer: A

Moderation APIs, such as Amazon Rekognition's Content Moderation API, can help filter and block inappropriate or unwanted images from being returned by a chatbot. These APIs are specifically designed to detect and manage undesirable

content in images. Option A (Correct): "Implement moderation APIs": This is the correct answer because moderation APIs are designed to identify and filter inappropriate content, ensuring the chatbot does not return unwanted images. Option

B: "Retrain the model with a general public dataset" is incorrect because retraining does not directly prevent inappropriate content from being returned. Option C: "Perform model validation" is incorrect as it ensures model correctness, not

content moderation.

Option D: "Automate user feedback integration" is incorrect because user feedback does not prevent inappropriate images in real-time.

AWS AI Practitioner

**References:**

AWS Content Moderation Services: AWS provides moderation APIs for filtering unwanted content from applications.

---

**QUESTION 4**

A social media company wants to use a large language model (LLM) for content moderation. The company wants to evaluate the LLM outputs for bias and potential discrimination against specific groups or individuals.

Which data source should the company use to evaluate the LLM outputs with the LEAST administrative effort?

- A. User-generated content

- B. Moderation logs
- C. Content moderation guidelines
- D. Benchmark datasets

Correct Answer: D

Benchmark datasets are pre-validated datasets specifically designed to evaluate machine learning models for bias, fairness, and potential discrimination. These datasets are the most efficient tool for assessing an LLM's performance against

known standards with minimal administrative effort.

Option D (Correct): "Benchmark datasets": This is the correct answer because using standardized benchmark datasets allows the company to evaluate model outputs for bias with minimal administrative overhead. Option A: "User-generated

content" is incorrect because it is unstructured and would require significant effort to analyze for bias. Option B: "Moderation logs" is incorrect because they represent historical data and do not provide a standardized basis for evaluating bias.

Option C: "Content moderation guidelines" is incorrect because they provide qualitative criteria rather than a quantitative basis for evaluation.

AWS AI Practitioner

References:

Evaluating AI Models for Bias on AWS: AWS supports using benchmark datasets to assess model fairness and detect potential bias efficiently.

---

## 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.

---

### QUESTION 6

A company is training a foundation model (FM). The company wants to increase the accuracy of the model up to a specific acceptance level.

Which solution will meet these requirements?

- A. Decrease the batch size.
- B. Increase the epochs.
- C. Decrease the epochs.
- D. Increase the temperature parameter.

Correct Answer: B

Increasing the number of epochs during model training allows the model to learn from the data over more iterations, potentially improving its accuracy up to a certain point. This is a common practice when attempting to reach a specific level of

accuracy. Option B (Correct): "Increase the epochs": This is the correct answer because increasing epochs allows the model to learn more from the data, which can lead to higher accuracy.

Option A: "Decrease the batch size" is incorrect as it mainly affects training speed and may lead to overfitting but does not directly relate to achieving a specific accuracy level.

Option C: "Decrease the epochs" is incorrect as it would reduce the training time, possibly preventing the model from reaching the desired accuracy. Option D: "Increase the temperature parameter" is incorrect because temperature affects the

randomness of predictions, not model accuracy.

AWS AI Practitioner

References:

Model Training Best Practices on AWS: AWS suggests adjusting training parameters, like the number of epochs, to improve model performance.

---

### QUESTION 7

An AI company periodically evaluates its systems and processes with the help of independent software vendors (ISVs). The company needs to receive email message notifications when an ISV's compliance reports become available.

Which AWS service can the company use to meet this requirement?

- A. AWS Audit Manager
- B. AWS Artifact
- C. AWS Trusted Advisor
- D. AWS Data Exchange

Correct Answer: D

AWS Data Exchange is a service that allows companies to securely exchange data with third parties, such as independent software vendors (ISVs). AWS Data Exchange can be configured to provide notifications, including email notifications,

when new datasets or compliance reports become available.

Option D (Correct): "AWS Data Exchange": This is the correct answer because it enables the company to receive notifications, including email messages, when ISVs' compliance reports are available.

Option A: "AWS Audit Manager" is incorrect because it focuses on assessing an organization's own compliance, not receiving third-party compliance reports. Option B: "AWS Artifact" is incorrect as it provides access to AWS's compliance

reports, not ISVs'.

Option C: "AWS Trusted Advisor" is incorrect as it offers optimization and best practices guidance, not compliance report notifications.

AWS AI Practitioner

References:

AWS Data Exchange Documentation: AWS explains how Data Exchange allows organizations to subscribe to third-party data and receive notifications when updates are available.

---

## QUESTION 8

A company is developing a new model to predict the prices of specific items. The model performed well on the training dataset. When the company deployed the model to production, the model's performance decreased significantly.

What should the company do to mitigate this problem?

- A. Reduce the volume of data that is used in training.
- B. Add hyperparameters to the model.
- C. Increase the volume of data that is used in training.
- D. Increase the model training time.

Correct Answer: C

When a model performs well on the training data but poorly in production, it is often due to overfitting. Overfitting occurs when a model learns patterns and noise specific to the training data, which does not generalize well to new, unseen data

in production. Increasing the volume of data used in training can help mitigate this problem by providing a more diverse and representative dataset, which helps the model generalize better. Option C (Correct): "Increase the volume of data

that is used in training":Increasing the data volume can help the model learn more generalized patterns rather than specific features of the training dataset, reducing overfitting and improving performance in production.

Option A:"Reduce the volume of data that is used in training" is incorrect, as reducing data volume would likely worsen the overfitting problem. Option B:"Add hyperparameters to the model" is incorrect because adding hyperparameters alone

does not address the issue of data diversity or model generalization.

Option D:"Increase the model training time" is incorrect because simply increasing training time does not prevent overfitting; the model needs more diverse data.

AWS AI Practitioner

References:

Best Practices for Model Training on AWS:AWS recommends using a larger and more diverse training dataset to improve a model's generalization capability and reduce the risk of overfitting.

---

## QUESTION 9

A company has thousands of customer support interactions per day and wants to analyze these interactions to identify frequently asked questions and develop insights.

Which AWS service can the company use to meet this requirement?

- A. Amazon Lex
- B. Amazon Comprehend
- C. Amazon Transcribe
- D. Amazon Translate

Correct Answer: B

Amazon Comprehend is the correct service to analyze customer support interactions and identify frequently asked questions and insights.

---

## QUESTION 10

A company needs to build its own large language model (LLM) based on only the company's private data. The company is concerned about the environmental effect of the training process.

Which Amazon EC2 instance type has the LEAST environmental effect when training LLMs?

- A. Amazon EC2 C series
- B. Amazon EC2 G series

C. Amazon EC2 P series

D. Amazon EC2 Trn series

Correct Answer: D

The Amazon EC2 Trn series (Trainium) instances are designed for high-performance, cost-effective machine learning training while being energy-efficient. AWS Trainium-powered instances are optimized for deep learning models and have

been developed to minimize environmental impact by maximizing energy efficiency. Option D (Correct): "Amazon EC2 Trn series": This is the correct answer because the Trn series is purpose-built for training deep learning models with lower

energy consumption, which aligns with the company's concern about environmental effects.

Option A: "Amazon EC2 C series" is incorrect because it is intended for compute-intensive tasks but not specifically optimized for ML training with environmental considerations.

Option B: "Amazon EC2 G series" (Graphics Processing Unit instances) is optimized for graphics-intensive applications but does not focus on minimizing environmental impact for training.

Option C: "Amazon EC2 P series" is designed for ML training but does not offer the same level of energy efficiency as the Trn series.

AWS AI Practitioner References:

AWS Trainium Overview: AWS promotes Trainium instances as their most energy-efficient and cost-effective solution for ML model training.

---

## QUESTION 11

A company has built an image classification model to predict plant diseases from photos of plant leaves. The company wants to evaluate how many images the model classified correctly.

Which evaluation metric should the company use to measure the model's performance?

A. R-squared score

B. Accuracy

C. Root mean squared error (RMSE)

D. Learning rate

Correct Answer: B

Accuracy is the most appropriate metric to measure the performance of an image classification model. It indicates the percentage of correctly classified images out of the total number of images. In the context of classifying plant diseases from

images, accuracy will help the company determine how well the model is performing by showing how many images were correctly classified.

Option B (Correct): "Accuracy": This is the correct answer because accuracy measures the proportion of correct



predictions made by the model, which is suitable for evaluating the performance of a classification model. Option A: "Root mean squared error (RMSE)" is incorrect as it is used for regression analysis, not classification tasks.

Option C: "Learning rate" is incorrect as it is a hyperparameter for training, not a performance metric.

Option D: "Accuracy" is incorrect because it is also used for regression tasks to measure prediction errors, not for classification accuracy. Option E: "Loss function" is incorrect as it is a performance metric, not a hyperparameter for training.

AWS AI Practitioner

References:

Evaluating Machine Learning Models on AWS: AWS documentation emphasizes the use of appropriate metrics, like accuracy, for classification tasks.

---

## QUESTION 12

A company wants to deploy a conversational chatbot to answer customer questions. The chatbot is based on a fine-tuned Amazon SageMaker JumpStart model. The application must comply with multiple regulatory frameworks.

Which capabilities can the company show compliance for? (Select TWO.)

- A. Auto scaling inference endpoints
- B. Threat detection
- C. Data protection
- D. Cost optimization
- E. Loosely coupled microservices

Correct Answer: BC

To comply with multiple regulatory frameworks, the company must ensure data protection and threat detection. Data protection involves safeguarding sensitive customer information, while threat detection identifies and mitigates security threats to the application.

Option C (Correct): "Data protection": This is correct because data protection is critical for compliance with privacy and security regulations. Option B (Correct): "Threat detection": This is correct because detecting and mitigating threats is essential to maintaining the security posture required for regulatory compliance.

Option A: "Auto scaling inference endpoints" is incorrect because auto-scaling does not directly relate to regulatory compliance.

Option D: "Cost optimization" is incorrect because it is focused on managing expenses, not compliance.

Option E: "Loosely coupled microservices" is incorrect because this architectural approach does not directly address compliance requirements.

AWS AI Practitioner

## References:

AWS Compliance Capabilities: AWS offers services and tools, such as data protection and threat detection, to help companies meet regulatory requirements for security and privacy.

---

**QUESTION 13**

A company is building an application that needs to generate synthetic data that is based on existing data.

Which type of model can the company use to meet this requirement?

- A. Generative adversarial network (GAN)
- B. XGBoost
- C. Residual neural network
- D. WaveNet

Correct Answer: A

Generative adversarial networks (GANs) are a type of deep learning model used for generating synthetic data based on existing datasets. GANs consist of two neural networks (a generator and a discriminator) that work together to create

realistic data. Option A (Correct): "Generative adversarial network (GAN)": This is the correct answer because GANs are specifically designed for generating synthetic data that closely resembles the real data they are trained on. Option

B: "XGBoost" is a gradient boosting algorithm for classification and regression tasks, not for generating synthetic data. Option C: "Residual neural network" is primarily used for improving the performance of deep networks, not for generating

synthetic data. Option D: "WaveNet" is a model architecture designed for generating raw audio waveforms, not synthetic data in general.

AWS AI Practitioner References:

GANs on AWS for Synthetic Data Generation: AWS supports the use of GANs for creating synthetic datasets, which can be crucial for applications like training machine learning models in environments where real data is scarce or sensitive.

---

**QUESTION 14**

A company has petabytes of unlabeled customer data to use for an advertisement campaign. The company wants to classify its customers into tiers to advertise and promote the company's products.

Which methodology should the company use to meet these requirements?

- A. Supervised learning
- B. Unsupervised learning
- C. Reinforcement learning
- D. Reinforcement learning from human feedback (RLHF)

Correct Answer: B

Unsupervised learning is the correct methodology for classifying customers into tiers when the data is unlabeled, as it does not require predefined labels or outputs.

---

## QUESTION 15

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.

[AIF-C01 VCE Dumps](#)

[AIF-C01 Exam Questions](#)

[AIF-C01 Braindumps](#)