



# 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.pass4itsure.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 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 2

A company uses Amazon SageMaker for its ML pipeline in a production environment. The company has large input data sizes up to 1 GB and processing times up to 1 hour. The company needs near real-time latency.

Which SageMaker inference option meets these requirements?

- A. Real-time inference
- B. Serverless inference



C. Asynchronous inference

D. Batch transform

Correct Answer: A

Real-time inference is designed to provide immediate, low-latency predictions, which is necessary when the company requires near real-time latency for its ML models. This option is optimal when there is a need for fast responses, even with

large input data sizes and substantial processing times.

Option A (Correct): "Real-time inference": This is the correct answer because it supports low-latency requirements, which are essential for real-time applications where quick response times are needed.

Option B: "Serverless inference" is incorrect because it is more suited for intermittent, small-scale inference workloads, not for continuous, large-scale, low-latency needs.

Option C: "Asynchronous inference" is incorrect because it is used for workloads that do not require immediate responses.

Option D: "Batch transform" is incorrect as it is intended for offline, large-batch processing where immediate response is not necessary.

AWS AI Practitioner References:

Amazon SageMaker Inference Options: AWS documentation describes real-time inference as the best solution for applications that require immediate prediction results with low latency.

---

### QUESTION 3

A company wants to use AI to protect its application from threats. The AI solution needs to check if an IP address is from a suspicious source.

Which solution meets these requirements?

A. Build a speech recognition system.

B. Create a natural language processing (NLP) named entity recognition system.

C. Develop an anomaly detection system.

D. Create a fraud forecasting system.

Correct Answer: C

An anomaly detection system is suitable for identifying unusual patterns or behaviors, such as suspicious IP addresses, which might indicate a potential threat.

---

### QUESTION 4

An AI practitioner is using a large language model (LLM) to create content for marketing campaigns. The generated content sounds plausible and factual but is incorrect.



Which problem is the LLM having?

- A. Data leakage
- B. Hallucination
- C. Overfitting
- D. Underfitting

Correct Answer: B

In the context of AI, "hallucination" refers to the phenomenon where a model generates outputs that are plausible-sounding but are not grounded in reality or the training data. This problem often occurs with large language models (LLMs)

when they create information that sounds correct but is actually incorrect or fabricated. Option B (Correct):

"Hallucination": This is the correct answer because the problem described involves generating content that sounds factual but is

incorrect, which is characteristic of hallucination in generative AI models. Option A: "Data leakage" is incorrect as it involves the model accidentally learning from data it shouldn't have access to, which does not match the problem of generating

incorrect content.

Option C: "Overfitting" is incorrect because overfitting refers to a model that has learned the training data too well, including noise, and performs poorly on new data.

Option D: "Underfitting" is incorrect because underfitting occurs when a model is too simple to capture the underlying patterns in the data, which is not the issue here.

AWS AI Practitioner

References:

Large Language Models on AWS: AWS discusses the challenge of hallucination in large language models and emphasizes techniques to mitigate it, such as using guardrails and fine-tuning.

---

## QUESTION 5

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 6

A company wants to create an application by using Amazon Bedrock. The company has a limited budget and prefers flexibility without long-term commitment.

Which Amazon Bedrock pricing model meets these requirements?

- A. On-Demand
- B. Model customization
- C. Provisioned Throughput
- D. Spot Instance

Correct Answer: A

Amazon Bedrock offers an on-demand pricing model that provides flexibility without long- term commitments. This model allows companies to pay only for the resources they use, which is ideal for a limited budget and offers flexibility. Option

A (Correct): "On-Demand":This is the correct answer because on-demand pricing allows the company to use Amazon Bedrock without any long-term commitments and to manage costs according to their budget. Option B:"Model

customization" is a feature, not a pricing model. Option C:"Provisioned Throughput" involves reserving capacity ahead of time, which might not offer the desired flexibility and could lead to higher costs if the capacity is not fully used.

Option D:"Spot Instance" is a pricing model for EC2 instances and does not apply to Amazon Bedrock.

AWS AI Practitioner

References:

AWS Pricing Models for Flexibility:On-demand pricing is a key AWS model for services that require flexibility and no



long-term commitment, ensuring cost- effectiveness for projects with variable usage patterns.

---

### QUESTION 7

What are tokens in the context of generative AI models?

- A. Tokens are the basic units of input and output that a generative AI model operates on, representing words, subwords, or other linguistic units.
- B. Tokens are the mathematical representations of words or concepts used in generative AI models.
- C. Tokens are the pre-trained weights of a generative AI model that are fine-tuned for specific tasks.
- D. Tokens are the specific prompts or instructions given to a generative AI model to generate output.

Correct Answer: A

Tokens in generative AI models are the smallest units that the model processes, typically representing words, subwords, or characters. They are essential for the model to understand and generate language, breaking down text into

manageable parts for processing.

Option A (Correct): "Tokens are the basic units of input and output that a generative AI model operates on, representing words, subwords, or other linguistic units": This is the correct definition of tokens in the context of generative AI models.

Option B: "Mathematical representations of words" describes embeddings, not tokens.

Option C: "Pre-trained weights of a model" refers to the parameters of a model, not tokens.

Option D: "Prompts or instructions given to a model" refers to the queries or commands provided to a model, not tokens.

AWS AI Practitioner

References:

Understanding Tokens in NLP: AWS provides detailed explanations of how tokens are used in natural language processing tasks by AI models, such as in Amazon Comprehend and other AWS AI services.

---

### QUESTION 8

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 9

An e-commerce company wants to build a solution to determine customer sentiments based on written customer reviews of products.

Which AWS services meet these requirements? (Select TWO.)

- A. Amazon Lex
- B. Amazon Comprehend
- C. Amazon Polly
- D. Amazon Bedrock
- E. Amazon Rekognition

Correct Answer: BD

To determine customer sentiments based on written customer reviews, the company can use Amazon Comprehend and Amazon Bedrock.

---

#### QUESTION 10

Which strategy evaluates the accuracy of a foundation model (FM) that is used in image classification tasks?

- A. Calculate the total cost of resources used by the model.
- B. Measure the model's accuracy against a predefined benchmark dataset.
- C. Count the number of layers in the neural network.
- D. Assess the color accuracy of images processed by the model.

Correct Answer: B

Measuring the model's accuracy against a predefined benchmark dataset is the correct strategy to evaluate the accuracy of a foundation model (FM) used in image classification tasks.

---

#### QUESTION 11

A large retailer receives thousands of customer support inquiries about products every day. The customer support inquiries need to be processed and responded to quickly. The company wants to implement Agents for Amazon Bedrock.

What are the key benefits of using Amazon Bedrock agents that could help this retailer?



- A. Generation of custom foundation models (FMs) to predict customer needs
- B. Automation of repetitive tasks and orchestration of complex workflows
- C. Automatically calling multiple foundation models (FMs) and consolidating the results
- D. Selecting the foundation model (FM) based on predefined criteria and metrics

Correct Answer: B

Amazon Bedrock Agents provide the capability to automate repetitive tasks and orchestrate complex workflows using generative AI models. This is particularly beneficial for customer support inquiries, where quick and efficient processing is

crucial. Option B (Correct): "Automation of repetitive tasks and orchestration of complex workflows": This is the correct answer because Bedrock Agents can automate common customer service tasks and streamline complex processes,

improving response times and efficiency.

Option A: "Generation of custom foundation models (FMs) to predict customer needs" is incorrect as Bedrock agents do not create custom models. Option C: "Automatically calling multiple foundation models (FMs) and consolidating the results" is incorrect because Bedrock agents focus on task automation rather than combining model outputs.

Option D: "Selecting the foundation model (FM) based on predefined criteria and metrics" is incorrect as Bedrock agents are not designed for selecting models.

AWS AI Practitioner

References:

Amazon Bedrock Documentation: AWS explains that Bedrock Agents automate tasks and manage complex workflows, making them ideal for customer support automation.

---

## QUESTION 12

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 13

A company wants to build an ML model by using Amazon SageMaker. The company needs to share and manage variables for model development across multiple teams.

Which SageMaker feature meets these requirements?

- A. Amazon SageMaker Feature Store
- B. Amazon SageMaker Data Wrangler
- C. Amazon SageMaker Clarify
- D. Amazon SageMaker Model Cards

Correct Answer: A

Amazon SageMaker Feature Store is the correct solution for sharing and managing variables (features) across multiple teams during model development.

---

### QUESTION 14

A company built a deep learning model for object detection and deployed the model to production.

Which AI process occurs when the model analyzes a new image to identify objects?

- A. Training
- B. Inference
- C. Model deployment
- D. Bias correction



Correct Answer: B

Inference is the correct answer because it is the AI process that occurs when a deployed model analyzes new data (such as an image) to make predictions or identify objects.

---

### QUESTION 15

A company makes forecasts each quarter to decide how to optimize operations to meet expected demand. The company uses ML models to make these forecasts.

An AI practitioner is writing a report about the trained ML models to provide transparency and explainability to company stakeholders.

What should the AI practitioner include in the report to meet the transparency and explainability requirements?

- A. Code for model training
- B. Partial dependence plots (PDPs)
- C. Sample data for training
- D. Model convergence tables

Correct Answer: B

Partial dependence plots (PDPs) are visual tools used to show the relationship between a feature (or a set of features) in the data and the predicted outcome of a machine learning model. They are highly effective for providing transparency

and explainability of the model's behavior to stakeholders by illustrating how different input variables impact the model's predictions.

Option B (Correct): "Partial dependence plots (PDPs)": This is the correct answer because PDPs help to interpret how the model's predictions change with varying values of input features, providing stakeholders with a clearer understanding of

the model's decision-making process.

Option A: "Code for model training" is incorrect because providing the raw code for model training may not offer transparency or explainability to non-technical stakeholders.

Option C: "Sample data for training" is incorrect as sample data alone does not explain how the model works or its decision-making process. Option D: "Model convergence tables" is incorrect. While convergence tables can show the training

process, they do not provide insights into how input features affect the model's predictions.

AWS AI Practitioner

References:

Explainability in AWS Machine Learning: AWS provides various tools for model explainability, such as Amazon SageMaker Clarify, which includes PDPs to help explain the impact of different features on the model's predictions.



VCE & PDF

Pass4itSure.com

<https://www.pass4itsure.com/aif-c01.html>

2024 Latest pass4itsure AIF-C01 PDF and VCE dumps Download

---

[Latest AIF-C01 Dumps](#)

[AIF-C01 PDF Dumps](#)

[AIF-C01 VCE Dumps](#)