

# Amazon AIF-C01 Exam Questions - Updated Frequently



## Amazon

### AIF-C01

**ExamName: AWS Certified AI Practitioner Exam**

**Questions & Answers Sample PDF**

(Preview content before you buy)

Check the full version using the link below.

<https://pass2certify.com/exam/aif-c01>

#### Unlock Full Features:

**Stay Updated:** 90 days of free exam updates

**Zero Risk:** 30-day money-back policy

**Instant Access:** Download right after purchase

**Always Here:** 24/7 customer support team

<https://pass2certify.com/exam/aif-c01>

Page 1 of 10

2026 Latest PDFVCE AIF-C01 PDF Dumps and AIF-C01 Exam Engine Free Share: [https://drive.google.com/open?id=1FPfW\\_91LpBUh4xYD0r\\_MhZts-nr\\_obS](https://drive.google.com/open?id=1FPfW_91LpBUh4xYD0r_MhZts-nr_obS)

The PDFVCE is one of the leading platforms that has been offering real and valid AWS Certified AI Practitioner (AIF-C01) exam practice test questions. These AWS Certified AI Practitioner (AIF-C01) exam questions are designed and verified by Amazon AIF-C01 subject matter experts. They work closely together and put all their expertise to check the Amazon AIF-C01 exam questions one by one.

The PDFVCE AIF-C01 PDF questions file, desktop practice test software, and web-based practice test software, all these three AIF-C01 practice test questions formats are ready for instant download. Just download any Amazon AIF-C01 Exam Questions format and start this journey with confidence.

>> **Instant AIF-C01 Discount** <<

## Pass Guaranteed Amazon - AIF-C01 - AWS Certified AI Practitioner Authoritative Instant Discount

Our AIF-C01 learning quiz has accompanied many people on their way to success and they will help you for sure. And you will learn about some of the advantages of our AIF-C01 training prep if you just free download the demos to have a check. You will understand that this is really a successful AIF-C01 Exam Questions that allows you to do more with less. With our AIF-C01 study materials for 20 to 30 hours, we can claim that you will pass the exam and get what you want.

### Amazon AIF-C01 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none"> <li>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.</li> </ul>
Topic 2	<ul style="list-style-type: none"> <li>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.</li> </ul>
Topic 3	<ul style="list-style-type: none"> <li>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.</li> </ul>
Topic 4	<ul style="list-style-type: none"> <li>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.</li> </ul>
Topic 5	<ul style="list-style-type: none"> <li>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.</li> </ul>

## Amazon AWS Certified AI Practitioner Sample Questions (Q87-Q92):

### NEW QUESTION # 87

Select the correct AI term from the following list for each statement. Each AI term should be selected one time. (Select THREE.)

\* AI

\* Deep learning

\* ML

Simulates human problem-solving capabilities

Select...
▼

Select...
▼

- AI
- Deep learning
- ML

Applies data-driven learning techniques to make predictions

Select...
▼

Select...
▼

- AI
- Deep learning
- ML

Focuses on processing data through intricate neural networks

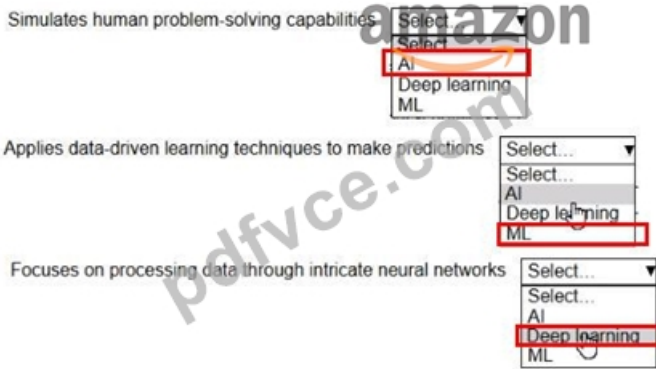
Select...
▼

Select...
▼

- AI
- Deep learning
- ML

**Answer:**

Explanation:



### NEW QUESTION # 88

An AI practitioner trained a custom model on Amazon Bedrock by using a training dataset that contains confidential data. The AI practitioner wants to ensure that the custom model does not generate inference responses based on confidential data. How should the AI practitioner prevent responses based on confidential data?

- A. Delete the custom model. Remove the confidential data from the training dataset. Retrain the custom model.
- B. Mask the confidential data in the inference responses by using dynamic data masking.
- C. Encrypt the confidential data in the inference responses by using Amazon SageMaker.
- D. Encrypt the confidential data in the custom model by using AWS Key Management Service (AWS KMS).

**Answer: A**

Explanation:

When a model is trained on a dataset containing confidential or sensitive data, the model may inadvertently learn patterns from this data, which could then be reflected in its inference responses. To ensure that a model does not generate responses based on confidential data, the most effective approach is to remove the confidential data from the training dataset and then retrain the model.

Explanation of Each Option:

\* Option A (Correct): "Delete the custom model. Remove the confidential data from the training dataset.

Retrain the custom model." This option is correct because it directly addresses the core issue: the model has been trained on confidential data. The only way to ensure that the model does not produce inferences based on this data is to remove the confidential information from the training dataset and then retrain the model from scratch. Simply deleting the model and retraining it ensures that no confidential data is learned or retained by the model. This approach follows the best practices recommended by AWS for handling sensitive data when using machine learning services like Amazon Bedrock.

\* Option B: "Mask the confidential data in the inference responses by using dynamic data masking." This option is incorrect because dynamic data masking is typically used to mask or obfuscate sensitive data in a database. It does not address the core problem of the model being trained on confidential data.

Masking data in inference responses does not prevent the model from using confidential data it learned during training.

\* Option C: "Encrypt the confidential data in the inference responses by using Amazon SageMaker." This option is incorrect because encrypting the inference responses does not prevent the model from generating outputs based on confidential data. Encryption only secures the data at rest or in transit but does not affect the model's underlying knowledge or training process.

\* Option D: "Encrypt the confidential data in the custom model by using AWS Key Management Service (AWS KMS)." This option is incorrect as well because encrypting the data within the model does not prevent the model from generating responses based on the confidential data it learned during training.

AWS KMS can encrypt data, but it does not modify the learning that the model has already performed.

AWS AI Practitioner References:

\* Data Handling Best Practices in AWS Machine Learning: AWS advises practitioners to carefully handle training data, especially when it involves sensitive or confidential information. This includes preprocessing steps like data anonymization or removal of sensitive data before using it to train machine learning models.

\* Amazon Bedrock and Model Training Security: Amazon Bedrock provides foundational models and customization capabilities, but any training involving sensitive data should follow best practices, such as removing or anonymizing confidential data to prevent unintended data leakage.

### NEW QUESTION # 89

A company wants to use Amazon Bedrock. The company needs to review which security aspects the company is responsible for when using Amazon Bedrock.

- A. Securing the company's data in transit and at rest
- B. Protecting the infrastructure that hosts Amazon Bedrock
- C. Patching and updating the versions of Amazon Bedrock
- D. Provisioning Amazon Bedrock within the company network

**Answer: A**

Explanation:

Comprehensive and Detailed

With Amazon Bedrock, AWS handles infrastructure security and patching (shared responsibility model).

Customers are responsible for securing their data (encryption, IAM, policies) both in transit and at rest.

Provisioning infrastructure (D) and platform patching (A, B) are AWS responsibilities.

Reference:

AWS Shared Responsibility Model

### NEW QUESTION # 90

A company trained an ML model on Amazon SageMaker to predict customer credit risk. The model shows 90% recall on training data and 40% recall on unseen testing data.

Which conclusion can the company draw from these results?

- A. The model has insufficient training data.
- B. The model has insufficient testing data.
- C. The model is overfitting on the training data.
- D. The model is underfitting on the training data.

**Answer: C**

Explanation:

The ML model shows 90% recall on training data but only 40% recall on unseen testing data, indicating a significant performance drop. This discrepancy suggests the model has learned the training data too well, including noise and specific patterns that do not generalize to new data, which is a classic sign of overfitting.

Exact Extract from AWS AI Documents:

From the Amazon SageMaker Developer Guide:

"Overfitting occurs when a model performs well on training data but poorly on unseen test data, as it has learned patterns specific to the training set, including noise, that do not generalize. A large gap between training and testing performance metrics, such as recall, is a common indicator of overfitting." (Source: Amazon SageMaker Developer Guide, Model Evaluation and Overfitting) Detailed Option A: The model is overfitting on the training data. This is the correct answer. The significant drop in recall from 90% (training) to 40% (testing) indicates the model is overfitting, as it performs well on training data but fails to generalize to unseen data.

Option B: The model is underfitting on the training data. Underfitting occurs when the model performs poorly on both training and testing data due to insufficient learning. With 90% recall on training data, the model is not underfitting.

Option C: The model has insufficient training data. Insufficient training data could lead to poor performance, but the high recall on training data (90%) suggests the model has learned the training data well, pointing to overfitting rather than a lack of data.

Option D: The model has insufficient testing data. Insufficient testing data might lead to unreliable test metrics, but it does not explain the large performance gap between training and testing, which is more indicative of overfitting.

Reference:

Amazon SageMaker Developer Guide: Model Evaluation and Overfitting (<https://docs.aws.amazon.com/sagemaker/latest/dg/model-evaluation.html>) AWS AI Practitioner Learning Path: Module on Model Performance and Evaluation AWS Documentation:

Understanding Overfitting and Underfitting (<https://aws.amazon.com/machine-learning/>)

### NEW QUESTION # 91

Which metric measures the runtime efficiency of operating AI models?

- A. Number of training instances
- B. Average response time
- C. Training time for each epoch
- D. Customer satisfaction score (CSAT)

**Answer: B**

