



# 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.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 financial institution is using Amazon Bedrock to develop an AI application. The application is hosted in a VPC. To meet regulatory compliance standards, the VPC is not allowed access to any internet traffic.

Which AWS service or feature will meet these requirements?

- A. AWS PrivateLink
- B. Amazon Macie
- C. Amazon CloudFront
- D. Internet gateway

Correct Answer: A

AWS PrivateLink enables private connectivity between VPCs and AWS services without exposing traffic to the public internet. This feature is critical for meeting regulatory compliance standards that require isolation from public internet traffic.

Option A (Correct): "AWS PrivateLink": This is the correct answer because it allows secure access to Amazon Bedrock and other AWS services from a VPC without internet access, ensuring compliance with regulatory standards. Option

B: "Amazon Macie" is incorrect because it is a security service for data classification and protection, not for managing private network traffic. Option C: "Amazon CloudFront" is incorrect because it is a content delivery network service and does

not provide private network connectivity. Option D: "Internet gateway" is incorrect as it enables internet access, which violates the VPC's no-internet-traffic policy.

AWS AI Practitioner References:

AWS PrivateLink Documentation: AWS highlights PrivateLink as a solution for connecting VPCs to AWS services privately, which is essential for organizations with strict regulatory requirements.

---

### QUESTION 2

A company is using the Generative AI Security Scoping Matrix to assess security responsibilities for its solutions. The company has identified four different solution scopes based on the matrix.

Which solution scope gives the company the MOST ownership of security responsibilities?

- A. Using a third-party enterprise application that has embedded generative AI features.
- B. Building an application by using an existing third-party generative AI foundation model (FM).
- C. Refining an existing third-party generative AI foundation model (FM) by fine-tuning the model by using data specific to the business.
- D. Building and training a generative AI model from scratch by using specific data that a customer owns.

Correct Answer: D



Building and training a generative AI model from scratch provides the company with the most ownership and control over security responsibilities. In this scenario, the company is responsible for all aspects of the security of the data, the

model, and the infrastructure. Option D (Correct): "Building and training a generative AI model from scratch by using specific data that a customer owns": This is the correct answer because it involves complete ownership of the model, data,

and infrastructure, giving the company the highest level of responsibility for security. Option A: "Using a third-party enterprise application that has embedded generative AI features" is incorrect as the company has minimal control over the

security of the AI features embedded within a third-party application. Option B: "Building an application using an existing third-party generative AI foundation model (FM)" is incorrect because security responsibilities are shared with the third-party model provider.

Option C: "Refining an existing third-party generative AI FM by fine-tuning the model with business-specific data" is incorrect as the foundation model and part of the security responsibilities are still managed by the third party.

AWS AI Practitioner

References:

Generative AI Security Scoping Matrix on AWS: AWS provides a security responsibility matrix that outlines varying levels of control and responsibility depending on the approach to developing and using AI models.

---

### QUESTION 3

A student at a university is copying content from generative AI to write essays.

Which challenge of responsible generative AI does this scenario represent?

- A. Toxicity
- B. Hallucinations
- C. Plagiarism
- D. Privacy

Correct Answer: C

The scenario where a student copies content from generative AI to write essays represents the challenge of plagiarism in responsible AI use.

---

### QUESTION 4

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 5

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.



VCE & PDF

GeekCert.com

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

2024 Latest geekcert AIF-C01 PDF and VCE dumps Download

---

[Latest AIF-C01 Dumps](#)

[AIF-C01 PDF Dumps](#)

[AIF-C01 Braindumps](#)