

KillTest

更に上のクオリティ 更に上のサービス



問題集

<http://www.killtest.jp>

1年で無料進級することに提供する

Exam : **AIF-C01**

Title : **AWS Certified AI
Practitioner**

Version : **DEMO**

1. 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:

A: Delete the custom model. Remove the confidential data from the training dataset. Retrain the custom model. Explanation: If the training dataset contains confidential data, the model may inadvertently learn and generate responses based on that data. The only way to ensure that the model does not generate responses based on the confidential data is to: Remove the confidential data from the training dataset. Retrain the custom model using the updated dataset. This process ensures that the model is not influenced by the sensitive information.

2. Which feature of Amazon OpenSearch Service gives companies the ability to build vector database applications?

- A. Integration with Amazon S3 for object storage
- B. Support for geospatial indexing and queries
- C. Scalable index management and nearest neighbor search capability
- D. Ability to perform real-time analysis on streaming data

Answer: C

Explanation:

The Amazon OpenSearch Service supports building vector database applications by enabling nearest neighbor search capability. This feature allows the service to efficiently perform similarity searches, which is crucial for applications that rely on vector embeddings (e.g., recommendation systems, image or text similarity searches). Combined with scalable index management, this makes OpenSearch an excellent choice for vector database applications.

3. A company wants to display the total sales for its top-selling products across various retail locations in the past 12 months.

Which AWS solution should the company use to automate the generation of graphs?

- A. Amazon Q in Amazon EC2
- B. Amazon Q Developer
- C. Amazon Q in Amazon QuickSight
- D. Amazon Q in AWS Chatbot

Answer: C

Explanation:

Amazon Q is a feature within Amazon QuickSight that allows users to ask questions about their data in

natural language and receive visualizations as responses. This functionality is particularly useful for generating graphs and visualizations based on specific queries regarding sales data.

4.A company wants to build an interactive application for children that generates new stories based on classic stories. The company wants to use Amazon Bedrock and needs to ensure that the results and topics are appropriate for children.

Which AWS service or feature will meet these requirements?

- A. Amazon Rekognition
- B. Amazon Bedrock playgrounds
- C. Guardrails for Amazon Bedrock
- D. Agents for Amazon Bedrock

Answer: C

Explanation:

C - Guardrails for Amazon Bedrock provides the necessary tools to ensure that the interactive story-generating application remains safe, appropriate, and engaging for children, making it the best choice for this scenario.

5.A company has developed an ML model for image classification. The company wants to deploy the model to production so that a web application can use the model.

The company needs to implement a solution to host the model and serve predictions without managing any of the underlying infrastructure.

Which solution will meet these requirements?

- A. Use Amazon SageMaker Serverless Inference to deploy the model.
- B. Use Amazon CloudFront to deploy the model.
- C. Use Amazon API Gateway to host the model and serve predictions.
- D. Use AWS Batch to host the model and serve predictions.

Answer: A

Explanation:

Amazon SageMaker Serverless Inference is a fully managed solution for deploying machine learning models without managing the underlying infrastructure. It automatically provisions compute capacity, scales based on request traffic, and serves predictions efficiently. This makes it an ideal choice for hosting a model and serving predictions for a web application with minimal management overhead. Why not the other options? B: Use Amazon CloudFront to deploy the model: Amazon CloudFront is a content delivery network (CDN) C: Use Amazon API Gateway to host the model and serve predictions: Amazon API Gateway is used to create APIs for accessing services. D: Use AWS Batch to host the model and serve predictions: AWS Batch is designed for batch processing and job scheduling, not for real-time inference or hosting ML models for web applications