

# Product Questions: 87

## Version: 4.0

---

### Question: 1

---

An AI practitioner trained a custom model on Amazon Bedrock by using a training dataset that contains confidential data

a. 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**

---

---

### Question: 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**

---

---

### Question: 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**

---

---

**Question: 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**

---

---

**Question: 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**

---

---

**Question: 6**

---

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)

---

**Answer: B**

---

---

### Question: 7

---

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

---

**Answer: B**

---

---

### Question: 8

---

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

---

**Answer: B**

---

---

### Question: 9

---

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

---

**Answer: B**

---

---

### Question: 10

---

A loan company is building a generative AI-based solution to offer new applicants discounts based on specific business criteria.

a. The company wants to build and use an AI model responsibly to minimize bias that could negatively affect some customers.

Which actions should the company take to meet these requirements? (Select TWO.)

- A. Detect imbalances or disparities in the data.
- B. Ensure that the model runs frequently.
- C. Evaluate the model's behavior so that the company can provide transparency to stakeholders.
- D. Use the Recall-Oriented Understudy for Gisting Evaluation (ROUGE) technique to ensure that the model is 100% accurate.
- E. Ensure that the model's inference time is within the accepted limits.

---

**Answer: A, C**

---

Let me know if you need more questions formatted similarly.

You said:

keep continue and complete it

ChatGPT said:



Sure, continuing from where we left off:

---

### Question: 11

---

A medical company is customizing a foundation model (FM) for diagnostic purposes. The company needs the model to be transparent and explainable to meet regulatory requirements.

Which solution will meet these requirements?

- A. Configure the security and compliance by using Amazon Inspector.
- B. Generate simple metrics, reports, and examples by using Amazon SageMaker Clarify.
- C. Encrypt and secure training data by using Amazon Macie.
- D. Gather more data. Use Amazon Rekognition to add custom labels to the data.

---

**Answer: B**

---

---

**Question: 12**

---

A company is building a solution to generate images for protective eyewear. The solution must have high accuracy and must minimize the risk of incorrect annotations. Which solution will meet these requirements?

- A. Human-in-the-loop validation by using Amazon SageMaker Ground Truth Plus
- B. Data augmentation by using an Amazon Bedrock knowledge base
- C. Image recognition by using Amazon Rekognition
- D. Data summarization by using Amazon QuickSight

---

**Answer: A**

---

---

**Question: 13**

---

A security company is using Amazon Bedrock to run foundation models (FMs). The company wants to ensure that only authorized users invoke the models. The company needs to identify any unauthorized access attempts to set appropriate AWS Identity and Access Management (IAM) policies and roles for future iterations of the FMs. Which AWS service should the company use to identify unauthorized users that are trying to access Amazon Bedrock?

- A. AWS Audit Manager
- B. AWS CloudTrail
- C. Amazon Fraud Detector
- D. AWS Trusted Advisor

---

**Answer: B**

---

---

**Question: 14**

---

A company manually reviews all submitted resumes in PDF format. As the company grows, the company expects the volume of resumes to exceed the company's review capacity. The company