



CertsHero

Amazon

AIF-C01 Exam

AWS Certified AI Practitioner

Full version

Questions and answers

**Thank You for
Purchasing**



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

needs an automated system to convert the **PDF resumes into plain text format** for additional processing.

Which AWS service meets this requirement?

- A. **Amazon Textract**
- B. Amazon Personalize
- C. Amazon Lex
- D. Amazon Transcribe

Answer: A

Question: 15

A company wants to use large language models (LLMs) with Amazon Bedrock to develop a **chat interface** for the company's product manuals. The manuals are stored as PDF files.

Which solution meets these requirements MOST cost-effectively?

- A. Use prompt engineering to add one PDF file as context to the user prompt when the prompt is submitted to Amazon Bedrock.
- B. Use prompt engineering to add all the PDF files as context to the user prompt when the prompt is submitted to Amazon Bedrock.
- C. Use all the PDF documents to fine-tune a model with Amazon Bedrock. Use the fine-tuned model to process user prompts.
- D. **Upload PDF documents to an Amazon Bedrock knowledge base.** Use the knowledge base to provide context when users submit prompts to Amazon Bedrock.

Answer: D

Question: 16

Which term describes the **numerical representations of real-world objects** and concepts that AI and natural language processing (NLP) models use to improve understanding of textual information?

- A. **Embeddings**
- B. Tokens
- C. Models
- D. Binaries

Answer: A

Question: 17

A company is building an application that needs to **generate synthetic data** that is based on existing data.

Which type of model can the company use to meet this requirement?

- A. Generative adversarial network (**GAN**)
- B. XGBoost
- C. Residual neural network
- D. WaveNet

Answer: A

Question: 18

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.

Answer: C

Question: 19

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

Answer: A

Question: 20

A digital devices company wants to predict customer demand for memory hardware. The company **does not have coding experience** or knowledge of ML algorithms and needs to develop a data-driven predictive model. The company needs to perform analysis on internal data and external data. Which solution will meet these requirements?

- A. Store the data in Amazon S3. Create ML models and demand forecast predictions by using Amazon SageMaker built-in algorithms that use the data from Amazon S3.
- B. Import the data into Amazon SageMaker Data Wrangler. Create ML models and demand forecast predictions by using SageMaker built-in algorithms.
- C. Import the data into Amazon SageMaker Data Wrangler. Build ML models and demand forecast predictions by using an Amazon Personalize Trending-Now recipe.
- D. Import the data into **Amazon SageMaker Canvas**. Build ML models and demand forecast predictions by selecting the values in the data from SageMaker Canvas.

Answer: D

I'll continue to format the rest. Let me know if you would like me to provide them all in one go or in parts.

You said:

yes keep continue

ChatGPT said:



Sure! Continuing from where we left off:

Question: 21

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.

Answer: A

Question: 22

An AI practitioner is using an **Amazon Bedrock** base model to summarize session chats from the customer service department. The AI practitioner wants to store **invocation logs** to monitor model input and output data.

Which strategy should the AI practitioner use?

- A. Configure AWS CloudTrail as the logs destination for the model.
- B. Enable **invocation logging** in **Amazon Bedrock**.
- C. Configure AWS Audit Manager as the logs destination for the model.
- D. Configure model invocation logging in Amazon EventBridge.

Answer: B

Question: 23

A company needs to build its own large language model (LLM) based on only the company's private data.

a. The company is concerned about the environmental effect of the training process.

Which Amazon EC2 instance type has the LEAST **environmental effect** when training LLMs?

- A. Amazon EC2 C series
- B. Amazon EC2 G series
- C. Amazon EC2 P series
- D. Amazon **EC2 Trn series**

Answer: D

Question: 24

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

Answer: A

Question: 25

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

Answer: B

6 Question: 26

A company is using Amazon SageMaker Studio notebooks to build and train ML models. The company **stores the data in an Amazon S3 bucket**. The company needs to manage the flow of data from Amazon S3 to SageMaker Studio notebooks.

Which solution will meet this requirement?

- A. Use Amazon Inspector to monitor SageMaker Studio.
- B. Use Amazon Macie to monitor SageMaker Studio.
- C. Configure **SageMaker to use a VPC with an S3 endpoint**.
- D. Configure SageMaker to use S3 Glacier Deep Archive.

Answer: C

6 Question: 27

A company is using domain-specific models. The company wants to avoid creating new models from the beginning. The company instead wants to adapt **pre-trained models** to create models for new, related tasks.

Which ML strategy meets these requirements?

- A. ~~Increase the number of epochs.~~
- B. Use **transfer learning**.
- C. Decrease the number of epochs.
- D. Use unsupervised learning.

Answer: B

Question: 28

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.

Answer: C

Question: 29

A company wants to use a large language model (LLM) to develop a conversational agent. The company needs to **prevent** the LLM from being manipulated with common prompt engineering techniques to perform **undesirable actions or expose sensitive information**.

Which action will reduce these risks?

- A. Create a **prompt template** that teaches the LLM to detect attack patterns.
- B. Increase the temperature parameter on invocation requests to the LLM.
- C. Avoid using LLMs that are not listed in Amazon SageMaker.
- D. Decrease the number of input tokens on invocations of the LLM.

Answer: A

Question: 30

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.

Answer: C

Question: 31

A company wants to create a chatbot by using a foundation model (FM) on **Amazon Bedrock**. The FM needs to access encrypted data that is stored in an Amazon S3 bucket.

The data is encrypted with Amazon S3 managed keys (SSE-S3).

The **FM encounters a failure when attempting to access the S3 bucket data**.

Which solution will meet these requirements?

- A. Ensure that the role that **Amazon Bedrock** assumes has permission to decrypt data with the **correct encryption key**.
- B. Set the access permissions for the S3 buckets to allow public access to enable access over the internet.
- C. Use prompt engineering techniques to tell the model to look for information in Amazon S3.
- D. Ensure that the S3 data does not contain sensitive information.

Answer: A

Question: 32

A company has a foundation model (FM) that was customized by using Amazon Bedrock to answer customer queries about products. The company wants to validate the model's responses to new types of queries. The company needs to upload a new dataset that Amazon Bedrock can use for validation.

Which AWS service meets these requirements?

- A. Amazon S3
- B. Amazon Elastic Block Store (Amazon EBS)
- C. Amazon Elastic File System (Amazon EFS)
- D. AWS Snowcone

Answer: A

I'll continue to format the remaining questions in the same format. Stay tuned!

You said:

yes continue

ChatGPT said:



Sure! Continuing from where we left off:

Question: 33

A company wants to assess the costs that are associated with using a large language model (LLM) to generate inferences. The company wants to use Amazon Bedrock to build generative AI applications. Which factor will ~~drive the inference costs~~?

- A. Number of tokens consumed
- B. Temperature value
- C. Amount of data used to train the LLM
- D. Total training time

Answer: A

Question: 34

An AI company periodically evaluates its systems and processes with the help of independent software vendors (ISVs). The company needs to receive email message notifications when an ISV's compliance reports become available.

Which AWS service can the company use to meet this requirement?

- A. AWS Audit Manager
- B. AWS Artifact
- C. AWS Trusted Advisor
- D. AWS Data Exchange

Answer: B

6 Question: 35

A company wants to use a large language model (LLM) on Amazon Bedrock for sentiment analysis. The company needs the LLM to produce more consistent responses to the same input prompt. Which adjustment to an inference parameter should the company make to meet these requirements?

- A. Decrease the temperature value
- B. Increase the temperature value
- C. Decrease the length of output tokens
- D. Increase the maximum generation length

Answer: A

Question: 36

A company is implementing the Amazon Titan foundation model (FM) by using Amazon Bedrock. The company needs to supplement the model by using relevant data from the company's private data sources.

Which solution will meet this requirement?

- A. Use a different FM
- B. Choose a lower temperature value
- C. Create an Amazon Bedrock knowledge base
- D. Enable model invocation logging

Answer: C

Question: 37

A company wants to develop an educational game where users answer questions such as the following: "A jar contains six red, four green, and three yellow marbles. What is the probability of choosing a green marble from the jar?"

Which solution meets these requirements with the LEAST operational overhead?

- A. Use supervised learning to create a regression model that will predict probability.
- B. Use reinforcement learning to train a model to return the probability.
- C. Use code that will calculate probability by using simple rules and computations.
- D. Use unsupervised learning to create a model that will estimate probability density.

Answer: C

Question: 38

Which functionality does Amazon SageMaker Clarify provide?

- A. Integrates a Retrieval Augmented Generation (RAG) workflow
- B. Monitors the quality of ML models in production
- C. Documents critical details about ML models
- D. Identifies potential bias during data preparation

Answer: D

Question: 39

A company is building an ML model. The company collected new data and analyzed the data by creating a correlation matrix, calculating statistics, and visualizing the data. Which stage of the ML pipeline is the company currently in?

- A. Data pre-processing
- B. Feature engineering
- C. Exploratory data analysis
- D. Hyperparameter tuning

Answer: C

Question: 40

A company has documents that are missing some words because of a database error. The company wants to build an ML model that can suggest potential words to fill in the missing text. Which type of model meets this requirement?

- A. Topic modeling
- B. Clustering models
- C. Prescriptive ML models
- D. BERT-based models

Answer: D

Question: 41

A company is building a chatbot to improve user experience. The company is using a large language model (LLM) from Amazon Bedrock for intent detection. The company wants to use few-shot learning to improve intent detection accuracy. Which additional data does the company need to meet these requirements?

- A. Pairs of chatbot responses and correct user intents
- B. Pairs of user messages and correct chatbot responses
- C. Pairs of user messages and correct user intents
- D. Pairs of user intents and correct chatbot responses

Answer: C

Question: 42

A company is building a large language model (LLM) question answering chatbot. The company wants to decrease the number of actions call center employees need to take to respond to customer questions.

Which business objective should the company use to evaluate the effect of the LLM chatbot?

- A. Website engagement rate
- B. Average call duration
- C. Corporate social responsibility
- D. Regulatory compliance

Answer: B

Question: 43

A company is using few-shot prompting on a base model that is hosted on Amazon Bedrock. The model currently uses 10 examples in the prompt. The model is invoked once daily and is performing well. The company wants to lower the monthly **cost**.

Which solution will meet these requirements?

- A. Customize the model by using fine-tuning.
- B. **Decrease** the number of **tokens** in the prompt.
- C. Increase the number of tokens in the prompt.
- D. Use Provisioned Throughput.

Answer: B

Question: 44

An accounting firm wants to implement a large language model (LLM) to automate document processing. The firm must proceed responsibly to avoid potential harms.

What should the firm do when developing and deploying the LLM? (Select TWO.)

- A. Include fairness metrics for **model evaluation**.
- B. Adjust the temperature parameter of the model.
- C. Modify the training **data** to **mitigate bias**.
- D. Avoid overfitting on the training data.
- E. Apply prompt engineering techniques.

Answer: A, C

Explanation:

I'll continue with more questions. Stay tuned!

You said:

yes continue

ChatGPT said:



Sure! Continuing from where we left off:

Question: 45

A company has built an image classification model to predict plant diseases from photos of plant leaves. The company wants to evaluate how many images the model classified correctly. Which evaluation metric should the company use to measure the model's **performance**?

- A. R-squared score
- B. **Accuracy**
- C. Root mean squared error (RMSE)
- D. Learning rate

Answer: B

Question: 46

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

Answer: B

Question: 47

A company is training a foundation model (FM). The company wants to **increase the accuracy** of the

model up to a specific acceptance level.
Which solution will meet these requirements?

- A. Decrease the batch size.
- B. **Increase the epochs.**
- C. Decrease the epochs.
- D. Increase the temperature parameter.

Answer: B

Question: 48

A company has built a chatbot that can respond to natural language questions with images. The company wants to ensure that the chatbot **does not return inappropriate or unwanted images**. Which solution will meet these requirements?

- A. **Implement moderation APIs.**
- B. Retrain the model with a general public dataset.
- C. Perform model validation.
- D. Automate user feedback integration.

Answer: A

Question: 49

A law firm wants to build an AI application by using large language models (LLMs). The application will read legal documents and extract key points from the documents. Which solution meets these requirements?

- A. Build an automatic named entity recognition system.
- B. Create a recommendation engine.
- C. **Develop a summarization chatbot.**
- D. Develop a multi-language translation system.

Answer: C

Question: 50

A company wants to classify human genes into 20 categories based on gene characteristics. The company **needs an ML algorithm to document how the inner mechanism of the model** affects the

output.

Which ML algorithm meets these requirements?

- A. Decision trees
- B. Linear regression
- C. Logistic regression
- D. Neural networks

Answer: A

Question: 51

A company wants to develop a large language model (LLM) application by using Amazon Bedrock and customer data that is uploaded to Amazon S3. The company's security policy states that **each team can access data for only the team's own customers**.

Which solution will meet these requirements?

- A. Create an Amazon **Bedrock custom service role** for each team that has access to only the team's customer data.
- B. Create a custom service role that has Amazon S3 access. Ask teams to specify the customer name on each Amazon Bedrock request.
- C. Redact personal data in Amazon S3. Update the S3 bucket policy to allow team access to customer data.
- D. Create one Amazon Bedrock role that has full Amazon S3 access. Create IAM roles for each team that have access to only each team's customer folders.

Answer: A

Question: 52

A company wants to use a **large language model (LLM)** on Amazon Bedrock for **sentiment analysis**. The company wants to know **how much information can fit into one prompt**.

Which consideration will inform the company's decision?

- A. Temperature
- B. Context window
- C. Batch size
- D. Model size

Answer: B

Question: 53

An AI practitioner has built a deep learning model to **classify** the types of materials in **images**. The AI practitioner now wants to measure the model performance. Which metric will help the AI practitioner evaluate the performance of the model?

- A. **Confusion** matrix
- B. Correlation matrix
- C. R2 score
- D. Mean squared error (MSE)

Answer: A

Question: 54

An AI practitioner is building a model to generate images of humans in various professions. The AI practitioner discovered that the input data is biased and that specific attributes affect the image generation and create bias in the model. Which technique will solve the problem?

- A. Data augmentation for imbalanced classes
- B. Model monitoring for class distribution
- C. Retrieval Augmented Generation (RAG)
- D. Watermark detection for images

Answer: A

Question: 55

A company is building an ML model to analyze archived data
a. The company must perform inference on large datasets that are multiple GBs in size. The company does not need to access the model predictions immediately.
Which Amazon SageMaker inference option will meet these requirements?

- A. Batch transform
- B. Real-time inference
- C. Serverless inference
- D. Asynchronous inference

Answer: A

Explanation:

I'll continue with more questions shortly. Stay tuned!

You said:

continue

ChatGPT said:



Sure! Continuing with the next set of questions:

Question: 56

A company needs to choose a model from Amazon Bedrock to use internally. The company must identify a model that generates responses in a style that the company's employees prefer. What should the company do to meet these requirements?

- A. Evaluate the models by using built-in prompt datasets.
- B. Evaluate the models by using a human workforce and custom prompt datasets.
- C. Use public model leaderboards to identify the model.
- D. Use the model InvocationLatency runtime metrics in Amazon CloudWatch when trying models.

Answer: B

Question: 57

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.

Answer: D

Question: 58

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

Answer: A

Question: 59

A company wants to use language models to create an application for inference on edge devices. The inference must have the lowest latency possible.

Which solution will meet these requirements?

- A. Deploy optimized small language models (SLMs) on edge devices.
- B. Deploy optimized large language models (LLMs) on edge devices.
- C. Incorporate a centralized small language model (SLM) API for asynchronous communication with edge devices.
- D. Incorporate a centralized large language model (LLM) API for asynchronous communication with edge devices.

Answer: A

Question: 60

A company is building a contact center application and wants to gain insights from customer conversations. The company wants to analyze and extract key information from the audio of the customer calls.

Which solution meets these requirements?

- A. Build a conversational chatbot by using Amazon Lex.
- B. Transcribe call recordings by using Amazon Transcribe.
- C. Extract information from call recordings by using Amazon SageMaker Model Monitor.
- D. Create classification labels by using Amazon Comprehend.

Answer: B

Question: 61

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

Answer: A

Question: 62

A company is using a pre-trained large language model (LLM) to build a chatbot for product recommendations. The company needs the LLM outputs to be short and written in a specific language.

Which solution will align the LLM response quality with the company's expectations?

- A. Adjust the prompt.
- B. Choose an LLM of a different size.
- C. Increase the temperature.
- D. Increase the Top K value.

Answer: A

Question: 63

A company uses a foundation model (FM) from Amazon Bedrock for an AI search tool. The company wants to fine-tune the model to be more accurate by using the company's data.

Which strategy will successfully fine-tune the model?

- A. Provide labeled data with the prompt field and the completion field.
- B. Prepare the training dataset by creating a .txt file that contains multiple lines in .csv format.
- C. Purchase Provisioned Throughput for Amazon Bedrock.
- D. Train the model on journals and textbooks.

Answer: A

Question: 64

An AI practitioner has a database of animal photos. The AI practitioner wants to automatically identify and categorize the animals in the photos without manual human effort. Which strategy meets these requirements?

- A. Object detection
- B. Anomaly detection
- C. Named entity recognition
- D. Inpainting

Answer: A

Question: 65

A research company implemented a chatbot by using a foundation model (FM) from Amazon Bedrock. The chatbot searches for answers to questions from a large database of research papers. After multiple prompt engineering attempts, the company notices that the FM is performing poorly because of the complex scientific terms in the research papers. How can the company improve the performance of the chatbot?

- A. Use few-shot prompting to define how the FM can answer the questions.
- B. Use domain adaptation fine-tuning to adapt the FM to complex scientific terms.
- C. Change the FM inference parameters.
- D. Clean the research paper data to remove complex scientific terms.

Answer: B

Question: 66

A medical company deployed a disease detection model on Amazon Bedrock. To comply with privacy policies, the company wants to prevent the model from including personal patient information in its responses. The company also wants to receive notification when policy violations occur. Which solution meets these requirements?

- A. Use Amazon Macie to scan the model's output for sensitive data and set up alerts for potential violations.
- B. Configure AWS CloudTrail to monitor the model's responses and create alerts for any detected personal information.
- C. Use Guardrails for Amazon Bedrock to filter content. Set up Amazon CloudWatch alarms for notification of policy violations.
- D. Implement Amazon SageMaker Model Monitor to detect data drift and receive alerts when model quality degrades.

Answer: C

Explanation:

I'll continue with more questions shortly. Stay tuned!

You said:

continue

ChatGPT said:



Sure! Continuing with the next set of questions:

Question: 67

An education provider is building a question and answer application that uses a generative AI model to explain complex concepts. The education provider wants to automatically change the style of the model response depending on who is asking the question. The education provider will give the model the age range of the user who has asked the question.

Which solution meets these requirements with the LEAST implementation effort?

- A. Fine-tune the model by using additional training data that is representative of the various age ranges that the application will support.
- B. Add a role description to the prompt context that instructs the model of the age range that the response should target.
- C. Use chain-of-thought reasoning to deduce the correct style and complexity for a response suitable for that user.
- D. Summarize the response text depending on the age of the user so that younger users receive shorter responses.

Answer: B

Question: 68

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

Answer: D

Question: 69

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.

Answer: B

Question: 70

A company has terabytes of data in a database that the company can use for business analysis. The company wants to build an AI-based application that can build a **SQL query** from **input text** that employees provide. The employees have minimal experience with technology. Which solution meets these requirements?

- A. **Generative pre-trained transformers (GPT)**
- B. Residual neural network
- C. Support vector machine
- D. WaveNet

Answer: A

Question: 71

Which metric measures the **runtime efficiency** of operating AI models?

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

Answer: C

Question: 72

Which option is a benefit of ongoing pre-training when fine-tuning a foundation model (FM)?

- A. Helps decrease the model's complexity
- B. Improves model performance over time
- C. Decreases the training time requirement
- D. Optimizes model inference time

Answer: B

Question: 73

An AI practitioner wants to use a foundation model (FM) to design a search application. The search application must handle queries that have text and images.

Which type of FM should the AI practitioner use to power the search application?

- A. Multi-modal embedding model
- B. Text embedding model
- C. Multi-modal generation model
- D. Image generation model

Answer: A

Question: 74

A company is using an Amazon Bedrock base model to summarize documents for an internal use case. The company trained a custom model to improve the summarization quality.

Which action must the company take to use the custom model through Amazon Bedrock?

- A. Purchase Provisioned Throughput for the custom model.
- B. Deploy the custom model in an Amazon SageMaker endpoint for real-time inference.
- C. Register the model with the Amazon SageMaker Model Registry.
- D. Grant access to the custom model in Amazon Bedrock.

Answer: B

Question: 75

A company has built a solution by using generative AI. The solution uses large language models (LLMs) to translate training manuals from English into other languages. The company wants to evaluate the accuracy of the solution by examining the text generated for the manuals. Which model evaluation strategy meets these requirements?

- A. Bilingual Evaluation Understudy (BLEU)
- B. Root mean squared error (RMSE)
- C. Recall-Oriented Understudy for Gisting Evaluation (ROUGE)
- D. F1 score

Answer: A

Question: 76

How can companies use large language models (LLMs) securely on Amazon Bedrock?

- A. Design clear and specific prompts. Configure AWS Identity and Access Management (IAM) roles and policies by using least privilege access.
- B. Enable AWS Audit Manager for automatic model evaluation jobs.
- C. Enable Amazon Bedrock automatic model evaluation jobs.
- D. Use Amazon CloudWatch Logs to make models explainable and to monitor for bias.

Answer: A

Question: 77

A company is building a customer service chatbot. The company wants the chatbot to improve its responses by learning from past interactions and online resources. Which AI learning strategy provides this self-improvement capability?

- A. Supervised learning with a manually curated dataset of good responses and bad responses
- B. Reinforcement learning with rewards for positive customer feedback
- C. Unsupervised learning to find clusters of similar customer inquiries
- D. Supervised learning with a continuously updated FAQ database

Answer: B

Question: 78

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

Answer: B, C

Explanation:

Let me know if you'd like to continue with any more questions or if you need further assistance!

You said:

yes, continue till end

ChatGPT said:



Sure! Continuing with the final set of questions:

Question: 79

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

Answer: B, D

Question: 80

A company wants to use a pre-trained generative AI model to generate content for its marketing campaigns. The company needs to ensure that the generated content aligns with the company's brand voice and messaging requirements.

Which solution meets these requirements?

- A. Optimize the model's architecture and hyperparameters to improve the model's overall performance.
- B. Increase the model's complexity by adding more layers to the model's architecture.
- C. Create effective prompts that provide clear instructions and context to guide the model's generation.
- D. Select a large, diverse dataset to pre-train a new generative model.

Answer: C

Question: 81

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

Answer: C

Question: 82

A company wants to make a chatbot to help customers. The chatbot will help solve technical problems without human intervention. The company chose a foundation model (FM) for the chatbot. The chatbot needs to produce responses that adhere to company tone. Which solution meets these requirements?

- A. Set a low limit on the number of tokens the FM can produce.
- B. Use batch inferencing to process detailed responses.
- C. Experiment and refine the prompt until the FM produces the desired responses.
- D. Define a higher number for the temperature parameter.

Answer: C

Question: 83

A company has installed a security camera

a. The company uses an ML model to evaluate the security camera footage for potential thefts. The company has discovered that the model disproportionately flags people who are members of a specific ethnic group.

Which type of bias is affecting the model output?

A. Measurement bias

B. Sampling bias

C. Observer bias

D. Confirmation bias

Answer: B

Question: 84

A company wants to use a large language model (LLM) on Amazon Bedrock for sentiment analysis. The company wants to classify the sentiment of text passages as positive or negative.

Which prompt engineering strategy meets these requirements?

A. Provide examples of text passages with corresponding positive or negative labels in the prompt followed by the new text passage to be classified.

B. Provide a detailed explanation of sentiment analysis and how LLMs work in the prompt.

C. Provide the new text passage to be classified without any additional context or examples.

D. Provide the new text passage with a few examples of unrelated tasks, such as text summarization or question answering.

Answer: A

Question: 85

Which AWS service or feature can help an AI development team quickly deploy and consume a foundation model (FM) within the team's VPC?

A. Amazon Personalize

B. Amazon SageMaker JumpStart

C. PartyRock, an Amazon Bedrock Playground

D. Amazon SageMaker endpoints

Answer: D

Question: 86

A company has a database of petabytes of unstructured data from internal sources. The company wants to transform this data into a structured format so that its data scientists can perform machine learning (ML) tasks.

Which service will meet these requirements?

- A. Amazon Lex
- B. Amazon Rekognition
- C. Amazon Kinesis Data Streams
- D. AWS Glue

Answer: D

Question: 87

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

Answer: B

Thank You for Purchasing AIF-C01 PDF

Download Free Practice Test Demo from Here:

<https://www.certshero.com/AIF-C01.html>

Test Your Preparation with Practice Exam Software

Use Coupon “**20OFF**” for extra 20% discount on purchase of Practice Test Software. Practice Exam Software helps you validate your preparation in simulated exam environment.

<https://www.certshero.com>