Exam Code: AIF-C01

Exam Name: AWS Certified AI Practitioner

Vdumps

Exam A

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
Section:

QUESTION 2
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

Correct Answer: B
Section:

QUESTION 3
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.

Correct Answer: C
Section:

QUESTION 4
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.

Correct Answer: B
Section:

QUESTION 5
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.

Correct Answer: C
Section:

QUESTION 6
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
Section:

QUESTION 7
A company is building an ML model to analyze archived data. 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

Correct Answer: A
Section:
Explanation:

QUESTION 8
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.

Correct Answer: B
Section:

QUESTION 9
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
Section:

QUESTION 10
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).

Correct Answer: A
Section:

QUESTION 11
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

Correct Answer: C
Section:

QUESTION 12
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

Correct Answer: C
Section:

QUESTION 13
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

Correct Answer: C
Section:

QUESTION 14
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.

Correct Answer: A
Section:

QUESTION 15
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

Correct Answer: B
Section:

QUESTION 16
A loan company is building a generative AI-based solution to offer new applicants discounts based on specific business criteria. 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.

Correct Answer: A, C
Section:

QUESTION 17
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.

Correct Answer: B
Section:

QUESTION 18
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

Correct Answer: A
Section:

QUESTION 19
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

Correct Answer: B
Section:

QUESTION 20
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

Correct Answer: A
Section:

QUESTION 21
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.

Correct Answer: D
Section:

QUESTION 22
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

Correct Answer: A
Section:

QUESTION 23
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

Correct Answer: A
Section:

QUESTION 24
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.

Correct Answer: C
Section:

QUESTION 25
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

Correct Answer: A
Section:

QUESTION 26
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.

Correct Answer: D
Section:

QUESTION 27
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.

Correct Answer: A
Section:

QUESTION 28
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.

Correct Answer: B
Section:

QUESTION 29
A company needs to build its own large language model (LLM) based on only the company's private data. 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

Correct Answer: D
Section:

QUESTION 30
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.

Correct Answer: A
Section:

QUESTION 31
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
Section:

QUESTION 32
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.

Correct Answer: A
Section:

QUESTION 33
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

Correct Answer: A
Section:

QUESTION 34
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

Correct Answer: A
Section:

QUESTION 35
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

Correct Answer: B
Section:

QUESTION 36
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

Correct Answer: A
Section:

QUESTION 37
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

Correct Answer: C
Section:

QUESTION 38
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.

Correct Answer: C
Section:

QUESTION 39
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

Correct Answer: D
Section:

QUESTION 40
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

Correct Answer: C
Section:

QUESTION 41
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

Correct Answer: D
Section:

QUESTION 42
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

Correct Answer: C
Section:

QUESTION 43
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

Correct Answer: B
Section:

QUESTION 44
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)

Correct Answer: B
Section:

QUESTION 45
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

Correct Answer: B
Section:

QUESTION 46
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

Correct Answer: B
Section:

QUESTION 47
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.

Correct Answer: B
Section: