



Vendor: Amazon

Exam Code: AIF-C01

Exam Name: AWS Certified AI Practitioner AIF-C01 Exam

Version: DEMO

QUESTION 1

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

Explanation:

Amazon SageMaker Ground Truth Plus is a managed data labeling service that includes human-in-the-loop (HITL) validation. This solution ensures high accuracy by involving human reviewers to validate the annotations and reduce the risk of incorrect annotations.

QUESTION 2

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

Explanation:

AWS CloudTrail is a service that enables governance, compliance, and operational and risk auditing of your AWS account. It tracks API calls and identifies unauthorized access attempts to AWS resources, including Amazon Bedrock.

QUESTION 3

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

Explanation:

Amazon Textract is a service that automatically extracts text and data from scanned documents, including PDFs. It is the best choice for converting resumes from PDF format to plain text for further processing.

QUESTION 4

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: A

Explanation:

Using Amazon Bedrock with large language models (LLMs) allows for efficient utilization of AI to answer queries based on context provided in product manuals. To achieve this cost-effectively, the company should avoid unnecessary use of resources.

QUESTION 5

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

Explanation:

Embeddings are numerical representations of objects (such as words, sentences, or documents) that capture the objects' semantic meanings in a form that AI and NLP models can easily understand. These representations help models improve their understanding of textual information by representing concepts in a continuous vector space.

QUESTION 6

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

Explanation:

Generative adversarial networks (GANs) are a type of deep learning model used for generating synthetic data based on existing datasets. GANs consist of two neural networks (a generator and

a discriminator) that work together to create realistic data.

QUESTION 7

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: A

Explanation:

Amazon Q Developer is a tool designed to assist developers in increasing productivity by generating code snippets, managing reference tracking, and handling open-source license tracking. These features help developers by automating parts of the software development process.

QUESTION 8

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

Explanation:

Amazon Bedrock offers an on-demand pricing model that provides flexibility without long-term commitments. This model allows companies to pay only for the resources they use, which is ideal for a limited budget and offers flexibility.

QUESTION 9

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

Explanation:

Amazon SageMaker Canvas is a visual, no-code machine learning interface that allows users to build machine learning models without having any coding experience or knowledge of machine learning algorithms. It enables users to analyze internal and external data, and make predictions using a guided interface.

QUESTION 10

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

Explanation:

Tokens in generative AI models are the smallest units that the model processes, typically representing words, subwords, or characters. They are essential for the model to understand and generate language, breaking down text into manageable parts for processing.

QUESTION 11

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

Explanation:

Amazon Bedrock provides an option to enable invocation logging to capture and store the input and output data of the models used. This is essential for monitoring and auditing purposes, particularly when handling customer data.

QUESTION 12

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

Answer: D

Explanation:

The Amazon EC2 Trn series (Trainium) instances are designed for high-performance, cost-effective machine learning training while being energy-efficient. AWS Trainium-powered instances are optimized for deep learning models and have been developed to minimize environmental impact by maximizing energy efficiency.

QUESTION 13

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

Explanation:

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.

QUESTION 161

Hotspot Question

A company is training its employees on how to structure prompts for foundation models.

Select the correct prompt engineering technique from the following list for each prompt template. Each prompt engineering technique should be selected one time.

"Classify the following text as either sports, politics, or entertainment: [input text]."

Select...
Select...
Chain-of-thought reasoning
Few-shot learning
Zero-shot learning

"A [image 1], [image 2], and [image 3] are examples of [target class]. Classify the following image as [target class]."

Select...
Select...
Chain-of-thought reasoning
Few-shot learning
Zero-shot learning

"[Question.] [Instructions to follow.] Think step by step and walk me through your thinking."

Select...
Select...
Chain-of-thought reasoning
Few-shot learning
Zero-shot learning

Answer:

"Classify the following text as either sports, politics, or entertainment: [input text]."

Select...
Select...
Chain-of-thought reasoning
Few-shot learning
Zero-shot learning

"A [image 1], [image 2], and [image 3] are examples of [target class]. Classify the following image as [target class]."

Select...
Select...
Chain-of-thought reasoning
Few-shot learning
Zero-shot learning

"[Question.] [Instructions to follow.] Think step by step and walk me through your thinking."

Select...
Select...
Chain-of-thought reasoning
Few-shot learning
Zero-shot learning

QUESTION 14

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

Explanation:

Inference is the correct answer because it is the AI process that occurs when a deployed model analyzes new data (such as an image) to make predictions or identify objects.

QUESTION 15

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

Explanation:

To manage the flow of data from Amazon S3 to SageMaker Studio notebooks securely, using a VPC with an S3 endpoint is the best solution.

Thank You for Trying Our Product

Braindump2go Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.braindump2go.com/all-products.html>



10% Discount Coupon Code: ASTR14