



Vendor: Amazon

Exam Code: AIF-C01

Exam Name: AWS Certified AI Practitioner AIF-C01 Exam

Version: DEMO

QUESTION 1

A financial company stores patterns of fraudulent behavior in a database. The company uses this data to conduct investigations.

The company wants to use a graph-based ML solution to develop an AI tool that helps with these investigations.

Which AWS service will meet these requirements?

- A. Amazon OpenSearch Service
- B. Amazon Aurora
- C. Amazon Neptune
- D. Amazon MemoryDB

Answer: C

Explanation:

Amazon Neptune is a fully managed graph database service designed for graph-based data models and graph analytics, making it suitable for building ML solutions that analyze relationships and patterns such as fraudulent behavior networks for investigation purposes.

QUESTION 2

An AI Practitioner is using an LLM-as-a-judge in Amazon Bedrock to evaluate the quality of agent responses in a production environment. The AI practitioner wants to apply a built-in metric that assesses how thoroughly the agent responses address all parts of each prompt or question. Which metric will meet these requirements?

- A. Recall-Oriented Understudy for Gisting Evaluation (ROUGE)
- B. Completeness
- C. Following instructions
- D. Refusal

Answer: B

Explanation:

The completeness metric evaluates how fully an agent's response addresses all parts and requirements of a user's prompt, making it the correct built-in metric for assessing thorough coverage of each question or request.

QUESTION 3

A company is building a generative AI application with a foundation model (FM). The application needs to automatically generate marketing emails. The company wants the application's output text to be creative and short in length. Which configuration of inference parameters will meet these requirements?

- A. Decrease the temperature and the response length.
- B. Increase the temperature and the response length.
- C. Increase the temperature and decrease the response length.
- D. Decrease the temperature and increase the response length.

Answer: C

Explanation:

Increasing the temperature encourages more creative and diverse text generation, while decreasing the response length constrains the output to be short, which fits the requirement for

creative, concise marketing emails.

QUESTION 4

A company is using AI to build a toy recommendation website that suggests toys based on a customer's interests and age. The company notices that the AI tends to suggest stereotypically gendered toys. Which AWS service or feature should the company use to investigate the bias?

- A. Amazon Rekognition
- B. Amazon Q Developer
- C. Amazon Comprehend
- D. Amazon SageMaker Clarify

Answer: D

Explanation:

Amazon SageMaker Clarify is designed to detect, measure, and analyze bias in datasets and machine learning models, enabling investigation of biased outcomes such as stereotypical recommendations.

QUESTION 5

A company has a team of AI practitioners that builds and maintains AI applications in an AWS account. The company must keep records of the actions that each AI practitioner takes in the AWS account for audit purposes. Which AWS service will meet these requirements?

- A. AWS CloudTrail
- B. AWS Config
- C. AWS Audit Manager
- D. AWS Trusted Advisor

Answer: A

Explanation:

AWS CloudTrail records and logs all API calls and user actions in an AWS account, providing a complete audit trail of activities performed by each AI practitioner for compliance and audit purposes.

QUESTION 6

A company wants to use AI for budgeting. The company made one budget manually and one budget by using an AI model. The company compared the budgets to evaluate the performance of the AI model. The AI model budget produced incorrect numbers. Which option represents the AI model's problem?

- A. Hallucinations
- B. Safety
- C. Interpretability
- D. Cost

Answer: A

Explanation:

The AI model generated incorrect numerical outputs that are not grounded in reality or the input data, which is characteristic of hallucinations in AI systems.

QUESTION 7

A company trains image and text generation models on Amazon SageMaker AI. The company releases the models by using Amazon Bedrock. The company must retain a tamper-proof, queryable record of every API call from SageMaker AI, Amazon Bedrock, and AWS Identity and Access Management (IAM). Which AWS service will meet these requirements?

- A. AWS Trusted Advisor
- B. Amazon Macie
- C. AWS CloudTrail Lake
- D. Amazon Inspector

Answer: C

Explanation:

AWS CloudTrail Lake provides a tamper-proof, immutable, and queryable event data store that records API activity across services such as Amazon SageMaker AI, Amazon Bedrock, and IAM, enabling long-term auditing, investigation, and compliance reporting.

QUESTION 8

Which type of ML technique provides the MOST explainability?

- A. Linear regression
- B. Support vector machines
- C. Random cut forest (RCF)
- D. Neural network

Answer: A

Explanation:

Linear regression provides the highest explainability because its parameters directly show how each input feature influences the output, making the model's behavior transparent and easily interpretable.

QUESTION 9

A company is building a custom AI solution in Amazon SageMaker Studio to analyze financial transactions for fraudulent activity in real time. The company needs to ensure that the connectivity from SageMaker Studio to Amazon Bedrock traverses the company's VPC. Which solution meets these requirements?

- A. Configure AWS Identity and Access Management (IAM) roles and policies for SageMaker Studio to access Amazon Bedrock.
- B. Configure Amazon Macie to proxy requests from SageMaker Studio to Amazon Bedrock.
- C. Configure AWS PrivateLink endpoints for the Amazon Bedrock API endpoints in the VPC that SageMaker Studio is connected to.
- D. Configure a new VPC for the Amazon Bedrock usage. Register the VPCs as peers.

Answer: C

Explanation:

AWS PrivateLink enables private connectivity to AWS service APIs through VPC interface endpoints, ensuring SageMaker Studio traffic to the Amazon Bedrock API stays within the company's VPC and does not traverse the public internet.

QUESTION 10

An education company wants to build a private tutor application. The application will give users the ability to enter text or provide a picture of a question. The application will respond with a written answer and an explanation of the written answer. Which model type meets these requirements?

- A. Computer vision model
- B. Multimodal LLM
- C. Diffusion model
- D. Text-to-speech model

Answer: B

Explanation:

A multimodal large language model can process both text and image inputs and generate text outputs, enabling the application to accept written questions or images of questions and return a written answer with an explanation.

QUESTION 11

Which AI technique combines large language models (LLMs) with external knowledge bases to improve response accuracy?

- A. Reinforcement learning (RL)
- B. Natural language processing (NLP)
- C. Retrieval Augmented Generation (RAG)
- D. Transfer learning

Answer: C

Explanation:

Retrieval Augmented Generation integrates large language models with external knowledge sources by retrieving relevant information at inference time and grounding the model's responses in that data, improving accuracy and factual reliability.

QUESTION 12

A company has deployed an AI application in production on AWS. The application's responses have become less accurate over time.

The company needs a solution to send alerts when the application performance drifts.

Which AWS service or feature will meet this requirement?

- A. Amazon Augmented AI (Amazon A2I)
- B. Amazon SageMaker Model Monitor
- C. Amazon Rekognition
- D. AWS Trusted Advisor

Answer: B

Explanation:

Amazon SageMaker Model Monitor detects data drift and model performance degradation in production and can trigger alerts when accuracy or data quality deviates from expected baselines.

QUESTION 13

A company stores its AI datasets in Amazon S3 buckets. The company wants to share the S3 buckets with its business partners. The company needs to avoid accidentally sharing sensitive data. Which AWS service should the company use to discover sensitive data in the dataset?

- A. Amazon Kendra
- B. Amazon Macie
- C. Amazon Textract
- D. AWS Data Exchange

Answer: B

Explanation:

Amazon Macie automatically discovers, classifies, and protects sensitive data in Amazon S3 using machine learning, enabling the company to identify sensitive information before sharing datasets with business partners.

QUESTION 14

What is continues pre-training?

- A. The process of fine-tuning a pre-trained language model on labeled data for a specific task
- B. The process of providing unlabeled data to a pre-trained language model to improve the model's domain knowledge
- C. The process of training a language model from the beginning on a specific dataset
- D. The process of evaluating the performance of a pre-trained language model on a test set

Answer: B

Explanation:

Continued pre-training improves a pre-trained language model's domain knowledge by training it further on large volumes of unlabeled, domain-specific data, enabling better understanding and performance in that specific domain.

QUESTION 15

A real estate company is developing an ML model to predict house prices by using sales and marketing data. The company wants to use feature engineering to build a model that makes accurate predictions. Which approach will meet these requirements?

- A. Understand patterns by providing data visualization.
- B. Tune the model's hyperparameters.
- C. Create or select relevant features for model training.
- D. Collect data from multiple sources.

Answer: C

Explanation:

Feature engineering focuses on creating, transforming, and selecting the most relevant input variables for model training, which directly improves the model's ability to learn meaningful patterns and make accurate house price predictions.

QUESTION 16

Which statement describes a generative AI use case for multimodal models?

- A. Deploy multiple scalable and cost-effective versions of a model.
- B. Process large amounts of data to train multiple models.

- C. Write code in multiple programming languages.
- D. Process different data types, such as images, audio, and video.

Answer: D

Explanation:

Multimodal generative AI models are designed to process and reason across multiple data types such as images, audio, video, and text within a single model.

QUESTION 17

Which term is the speed at which a pre-trained foundation model (FM) processes requests and delivers output?

- A. Model size
- B. Inference latency
- C. Context window
- D. Fine-tuning

Answer: B

Explanation:

Inference latency refers to the time a foundation model takes to process an input request and produce an output, representing the speed of model response during inference.

QUESTION 18

A company is using a large collection of web data to produce a large language model (LLM). The company completes a random initialization of the model's weights. Next, the company fits the model to the data through a language objective modelling function. Which stage of the model training process does this scenario describe?

- A. Fine-tuning
- B. Pre-training
- C. Model selection
- D. Deployment

Answer: B

Explanation:

The model is being trained from randomly initialized weights on a large corpus using a general language modeling objective, which corresponds to the pre-training phase where foundational knowledge is learned.

QUESTION 19

Hotspot Question

A company is building an AI assistant application. The company must implement a core governance process for the application development project. The company must ensure that the application aligns with responsible AI practices.

Select and order the steps from the following list to correctly describe the implementation of a core governance process for this use case. Select each step one time.

- Step 1:
- | | |
|--|---|
| Select... | ▼ |
| Determine governance goals, risks, and policies. | |
| Put together a cross-functional AI governance group. | |
| Set up model monitoring mechanisms. | |
- Step 2:
- | | |
|--|---|
| Select... | ▼ |
| Determine governance goals, risks, and policies. | |
| Put together a cross-functional AI governance group. | |
| Set up model monitoring mechanisms. | |
- Step 3:
- | | |
|--|---|
| Select... | ▼ |
| Determine governance goals, risks, and policies. | |
| Put together a cross-functional AI governance group. | |
| Set up model monitoring mechanisms. | |

Answer:

Step 1:	Select... ▼ Determine governance goals, risks, and policies. Put together a cross-functional AI governance group. Set up model monitoring mechanisms.
Step 2:	Select... ▼ Determine governance goals, risks, and policies. Put together a cross-functional AI governance group. Set up model monitoring mechanisms.
Step 3:	Select... ▼ Determine governance goals, risks, and policies. Put together a cross-functional AI governance group. Set up model monitoring mechanisms.

QUESTION 20

A financial company is training a generative AI model to predict outcomes of loan applications. The training dataset is small. The dataset categorizes loan applicants as “younger-aged,” “middle-aged,” or “older-aged.” Most individuals in the dataset are characterized as “middle-aged.”

The company removes the age range feature from the training dataset.

Which model behavior will likely happen as a result of this change to the dataset?

- A. The model will inaccurately predict outcomes for younger and older age groups.
- B. The model will require less training data.
- C. The model will predict accurate outcomes for only younger age groups.
- D. The model will accurately predict outcomes for all ages.

Answer: A

Explanation:

Removing the age feature eliminates an important variable that differentiates applicant groups. Since the dataset is already imbalanced, the model will generalize based on dominant patterns and lose the ability to account for age-related differences, leading to poorer predictions especially for underrepresented younger and older groups.

Thank You for Trying Our Product

Lead2pass 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.lead2pass.com/all-products.html>



10% Discount Coupon Code: ASTR14