AIF-C01 ^{Q&As}

Amazon AWS Certified AI Practitioner

Pass Amazon AIF-C01 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.certbus.com/aif-c01.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon
Official Exam Center

⚙ **Instant Download** After Purchase

⚙ **100% Money Back** Guarantee

⚙ **365 Days** Free Update

⚙ **800,000+** Satisfied Customers

QUESTION 1

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.

Correct Answer: B

Domain adaptation fine-tuning involves training a foundation model (FM) further using a specific dataset that includes domain-specific terminology and content, such as scientific terms in research papers. This process allows the model to better understand and handle complex terminology, improving its performance on specialized tasks. Option B (Correct): "Use domain adaptation fine-tuning to adapt the FM to complex scientific terms":This is the correct answer because fine-tuning the model on domain-specific data helps it learn and adapt to the specific language and terms used in the research papers, resulting in better performance. Option A:"Use few-shot prompting to define how the FM can answer the questions" is incorrect because while few-shot prompting can help in certain scenarios, it is less effective than fine-tuning for handling complex domain-specific terms. Option C:"Change the FM inference parameters" is incorrect because adjusting inference parameters will not resolve the issue of the model\\'s lack of understanding of complex scientific terminology. Option D:"Clean the research paper data to remove complex scientific terms" is incorrect because removing the complex terms would result in the loss of important information and context, which is not a viable solution. AWS AI Practitioner References: Domain Adaptation in Amazon Bedrock:AWS recommends fine-tuning models with domain-specific data to improve their performance on specialized tasks involving unique terminology.

QUESTION 2

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

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.

Option

A (Correct): "On-Demand":This is the correct answer because on-demand pricing allows the company to use Amazon Bedrock without any long-term commitments and to manage costs according to their budget. Option B:"Model

customization" is a feature, not a pricing model. Option C:"Provisioned Throughput" involves reserving capacity ahead of time, which might not offer the desired flexibility and could lead to higher costs if the capacity is not fully used.

Option D:"Spot Instance" is a pricing model for EC2 instances and does not apply to Amazon Bedrock.

AWS AI Practitioner

References:

AWS Pricing Models for Flexibility:On-demand pricing is a key AWS model for services that require flexibility and no long-term commitment, ensuring cost- effectiveness for projects with variable usage patterns.

QUESTION 3

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

Correct Answer: A

Generative Pre-trained Transformers (GPT) are suitable for building an AI-based application that can generate SQL queries from natural language input provided by employees.

QUESTION 4

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

Batch transform in Amazon SageMaker is designed for offline processing of large datasets. It is ideal for scenarios where immediate predictions are not required, and the inference can be done on large datasets that are multiple gigabytes in

size. This method processes data in batches, making it suitable for analyzing archived data without the need for real-time access to predictions.

Option A (Correct): "Batch transform":This is the correct answer because batch transform is optimized for handling large datasets and is suitable when immediate access to predictions is not required.

Option B:"Real-time inference" is incorrect because it is used for low-latency, real- time prediction needs, which is not required in this case. Option C:"Serverless inference" is incorrect because it is designed for small-scale, intermittent

inference requests, not for large batch processing. Option D:"Asynchronous inference" is incorrect because it is used when immediate predictions are required, but with high throughput, whereas batch transform is more suitable for very large

datasets.

AWS AI Practitioner

References:

Batch Transform on AWS SageMaker:AWS recommends using batch transform for large datasets when real-time processing is not needed, ensuring cost- effectiveness and scalability.

QUESTION 5

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

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. Option B

(Correct): "Enable invocation logging in Amazon Bedrock":This is the correct answer as it directly enables the logging of all model invocations, ensuring transparency and traceability.

Option A:"Configure AWS CloudTrail" is incorrect because CloudTrail logs API calls but does not provide specific logging for model inputs and outputs. Option C:"Configure AWS Audit Manager" is incorrect as Audit Manager is used for

compliance reporting, not specific invocation logging for AI models. Option D:"Configure model invocation logging in Amazon EventBridge" is incorrect as EventBridge is for event-driven architectures, not specifically designed for logging AI

model inputs and outputs.

AWS AI Practitioner

References:

Amazon Bedrock Logging Capabilities:AWS emphasizes using built-in logging features in Bedrock to maintain data integrity and transparency in model operations.

QUESTION 6

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

Amazon S3 is the optimal choice for storing and uploading datasets used for machine learning model validation and training. It offers scalable, durable, and secure storage, making it ideal for holding datasets required by Amazon Bedrock for

validation purposes. Option A (Correct): "Amazon S3":This is the correct answer because Amazon S3 is widely used for storing large datasets that are accessed by machine learning models, including those in Amazon Bedrock.

Option B:"Amazon Elastic Block Store (Amazon EBS)" is incorrect because EBS is a block storage service for use with Amazon EC2, not for directly storing datasets for Amazon Bedrock.

Option C:"Amazon Elastic File System (Amazon EFS)" is incorrect as it is primarily used for file storage with shared access by multiple instances. Option D:"AWS Snowcone" is incorrect because it is a physical device for offline data transfer,

not suitable for directly providing data to Amazon Bedrock.

AWS AI Practitioner

References:

Storing and Managing Datasets on AWS for Machine Learning:AWS recommends using S3 for storing and managing datasets required for ML model training and validation.

QUESTION 7

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

Correct Answer: BC

To comply with multiple regulatory frameworks, the company must ensure data protection and threat detection. Data protection involves safeguarding sensitive customer information, while threat detection identifies and mitigates security

threats to the application. Option C (Correct): "Data protection":This is correct because data protection is critical for compliance with privacy and security regulations. Option B (Correct): "Threat detection":This is correct because detecting and

mitigating threats is essential to maintaining the security posture required for regulatory compliance.

Option A:"Auto scaling inference endpoints" is incorrect because auto-scaling does not directly relate to regulatory compliance.

Option D:"Cost optimization" is incorrect because it is focused on managing expenses, not compliance.

Option E:"Loosely coupled microservices" is incorrect because this architectural approach does not directly address compliance requirements.

AWS AI Practitioner

References:

AWS Compliance Capabilities:AWS offers services and tools, such as data protection and threat detection, to help companies meet regulatory requirements for security and privacy.

QUESTION 8

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

The scenario where a student copies content from generative AI to write essays represents the challenge ofplagiarismin responsible AI use.

QUESTION 9

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

When a model is trained on a dataset containing confidential or sensitive data, the model may inadvertently learn patterns from this data, which could then be reflected in its inference responses. To ensure that a model does not generate responses based on confidential data, the most effective approach is to remove the confidential data from the training dataset and then retrain the model. of Each Option: Option A (Correct):"Delete the custom model. Remove the confidential data from the training dataset. Retrain the custom model."This option is correct because it directly addresses the core issue: the model has been trained on confidential data. The only way to ensure that the model does not produce inferences based on this data is to remove the confidential information from the training dataset and then retrain the model from scratch. Simply deleting the model and retraining it ensures that no confidential data is learned or retained by the model. This approach follows the best practices recommended by AWS for handling sensitive data when using machine learning services like Amazon Bedrock. Option B:"Mask the confidential data in the inference responses by using dynamic data masking."This option is incorrect because dynamic data masking is typically used to mask or obfuscate sensitive data in a database. It does not address the core problem of the model being trained on confidential data. Masking data in inference responses does not prevent the model from using confidential data it learned during training. Option C:"Encrypt the confidential data in the inference responses by using Amazon SageMaker."This option is incorrect because encrypting the inference responses does not prevent the model from generating outputs based on confidential data. Encryption only secures the data at rest or in transit but does not affect the model\'s underlying knowledge or training process. Option D:"Encrypt the confidential data in the custom model by using AWS Key Management Service (AWS KMS)."This option is incorrect as well because encrypting the data within the model does not prevent the model from generating responses based on the confidential data it learned during training. AWS KMS can encrypt data, but it does not modify the learning that the model has already performed. AWS AI Practitioner References: Data Handling Best Practices in AWS Machine Learning:AWS advises practitioners to carefully handle training data, especially when it involves sensitive or confidential information. This includes preprocessing steps like data anonymization or removal of sensitive data before using it to train machine learning models. Amazon Bedrock and Model Training Security:Amazon Bedrock provides foundational models and customization capabilities, but any training involving sensitive data should follow best practices, such as removing or anonymizing confidential data to prevent unintended data leakage.

QUESTION 10

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)

Correct Answer: A

A confusion matrix is the correct metric for evaluating the performance of a classification model, such as the deep learning model built to classify types of materials in images.

QUESTION 11

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

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 12

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

Amazon Bedrock is a service that provides foundational models for building generative AI applications. When creating an application for children, it is crucial to ensure that the generated content is appropriate for the target audience.

"Guardrails" in Amazon Bedrock provide mechanisms to control the outputs and topics of generated content to align with desired safety standards and appropriateness levels. Option C (Correct): "Guardrails for Amazon Bedrock":This is the

correct answer because guardrails are specifically designed to help users enforce content moderation, filtering, and safety checks on the outputs generated by models in Amazon Bedrock. For a children\\'s application, guardrails ensure that all

content generated is suitable and appropriate for the intended audience. Option A:"Amazon Rekognition" is incorrect. Amazon Rekognition is an image and video analysis service that can detect inappropriate content in images or videos, but it

does not handle text or story generation. Option B:"Amazon Bedrock playgrounds" is incorrect because playgrounds are environments for experimenting and testing model outputs, but they do not inherently provide safeguards to ensure

content appropriateness for specific audiences, such as children.

Option D:"Agents for Amazon Bedrock" is incorrect. Agents in Amazon Bedrock facilitate building AI applications with more interactive capabilities, but they do not provide specific guardrails for ensuring content appropriateness for children.

AWS AI Practitioner References:

Guardrails in Amazon Bedrock:Designed to help implement controls that ensure generated content is safe and suitable for specific use cases or audiences, such as children, by moderating and filtering inappropriate or undesired content.

Building Safe AI Applications:AWS provides guidance on implementing ethical AI practices, including using guardrails to protect against generating inappropriate or biased content.

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

Correct Answer: A

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.

Option A (Correct): "AWS PrivateLink":This is the correct answer because it allows secure access to Amazon Bedrock and other AWS services from a VPC without internet access, ensuring compliance with regulatory standards. Option

B:"Amazon Macie" is incorrect because it is a security service for data classification and protection, not for managing private network traffic. Option C:"Amazon CloudFront" is incorrect because it is a content delivery network service and does

not provide private network connectivity. Option D:"Internet gateway" is incorrect as it enables internet access, which violates the VPC\\'s no-internet-traffic policy.

AWS AI Practitioner References:

AWS PrivateLink Documentation:AWS highlights PrivateLink as a solution for connecting VPCs to AWS services privately, which is essential for organizations with strict regulatory requirements.

QUESTION 14

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

An anomaly detection system is suitable for identifying unusual patterns or behaviors, such as suspicious IP addresses, which might indicate a potential threat.

QUESTION 15

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.

Correct Answer: A

Adjusting the prompt is the correct solution to align the LLM outputs with the company\\'s expectations for short, specific

language responses.

AIF-C01 VCE Dumps   AIF-C01 Practice Test   AIF-C01 Exam Questions