Category Breakdown

Category	Number of Questions
AI Governance and Risk Management	4
Ethical and Legal Considerations in AI	3
AI in Audit Processes	1
AI Fundamentals and Technologies	1
AI Operations and Performance	1
TOTAL	10

When using off-the-shelf AI models, which of the following is the MOST appropriate way for organizations to approach vendor management?

A. Ensure a minimum of three quotes have been obtained for market research and comparison.

B. Establish responsibility and clear terms for model updates and support.

C. Only use models from vendors with globally recognized accreditation.

D. Use the vendor only if the contract has been reviewed by the information security department.

Answer: B

Explanation

When organizations leverage off-the-shelf AI models, effective vendor management is critical to ensure operational reliability, compliance, and long-term support. The ISACA Advanced in AI Audit™ (AAIA™) Study Guide highlights that the "establishment of clear contractual terms regarding responsibilities for ongoing model updates, maintenance, support, and incident response is essential for managing third-party AI risks."

By clearly defining the roles and expectations for updates and support (option B), organizations reduce the risk of unaddressed vulnerabilities, outdated models, or unclear recourse in the event of an incident or system failure. This approach supports ongoing risk management and ensures that both parties understand their obligations throughout the model's lifecycle.

While market research, vendor accreditation, and contract review by information security are important due diligence steps, they do not directly address the need for clarity in ongoing vendor responsibilities, which is critical for effective governance and sustained operation of AI solutions.

Reference:ISACA Advanced in AI Audit™ (AAIA™) Study Guide, Section: "Vendor Management for AI Systems," Subsection: "Third-Party AI Risk and Contractual Obligations"

Practice Test Isaca - AAIA

An organization uses an AI image generation platform to create promotional materials. An IS auditor identifies that the platform includes copyrighted images in its training data. Which of the following is the auditor's BEST recommendation to address this issue?

A. Implement a manual review process to ensure no copyrighted images are used in generated outputs.

B. Use a platform that certifies the provenance and licensing of its training data.

C. Label all AI-generated images to disclaim the possibility of third-party content.

D. Suspend the use of the platform until the training data is sanitized.

Answer: B

Explanation

Ensuring that AI tools are trained on properly licensed and documented data sets is critical to avoiding copyright infringement and legal exposure. The AAIA™ Study Guide emphasizes using platforms with certified and traceable training data to meet ethical and legal standards.

“Organizations must verify the provenance and licensing of data used to train AI systems. Platforms that certify data sources reduce the risk of using protected intellectual property without consent.”

Manual review (A) is resource-intensive and may not detect embedded copyright violations. Labeling (C) is not sufficient for legal protection. Suspension (D) may be excessive without first attempting remediation.

Thus, B is the most strategic and effective recommendation.

Reference: ISACA Advanced in AI Audit™ (AAIA™) Study Guide, Section: “Ethical and Legal Considerations in AI,” Subsection: “Intellectual Property and Data Licensing in AI Systems”

Question #:3 - [AI Governance and Risk Management] Which of the following is MOST important to consider when auditing an organization's AI procedures?

A. Frequency of AI system updates to enhance security

B. Employee training on recognized AI best practices

C. Backup and recovery in the event of an AI data breach

D. AI data validation and filtration to prevent data poisoning

Answer: D

Explanation

The integrity of data fed into AI systems is a critical concern. The AAIA™ Study Guide emphasizes that validation and filtration processes are essential to mitigate the risk of data poisoning—an attack that can manipulate model behavior by injecting malicious inputs.

“Data poisoning represents a major vulnerability in AI pipelines. Effective controls include robust validation, filtration, and monitoring of training data sources. These preventive practices are essential to ensure model reliability and security.”

While options A, B, and C are important operational and training measures, only D addresses a technical risk that can directly compromise model outputs and trustworthiness.

Reference: ISACA Advanced in AI Audit™ (AAIA™) Study Guide, Section: “AI Governance and Risk Management,” Subsection: “AI Data Integrity and Attack Prevention”

Question #:4 - [Ethical and Legal Considerations in AI]

A retail organization uses an AI model to analyze customers' purchase history in order to offer personalized discounts. Which of the following practices represents the MOST ethical use of customer data?

A. Utilizing customer purchase data only after obtaining explicit consent and allowing customers to opt out

B. Retaining and analyzing all available customer data to ensure unbiased recommendations

C. Providing the public with access to review and audit the data set of collected customer information

D. Sharing customer purchase data with third-party vendors to improve advertising and communication

Answer: A

Explanation

The ethical use of customer data is rooted in respecting privacy, maintaining informed consent, and enabling data subjects to exercise control over their personal information. The AAIA™ Study Guide clearly outlines that obtaining explicit consent and providing opt-out capabilities align with principles of data protection and ethical AI.

“Ethical AI implementation includes transparency in data collection, clear consent mechanisms, and the right of users to opt out or control their personal data usage. Retail and consumer applications must ensure that personalized services do not override these data subject rights.”

Options B and D violate principles of minimal data use and consent, while C may create unnecessary privacy exposure. Therefore, A is the correct and most ethical practice.

Reference: ISACA Advanced in AI Audit™ (AAIA™) Study Guide, Section: “Ethical and Legal Considerations in AI,” Subsection: “Informed Consent and Customer Privacy”

Question #:5 - [AI in Audit Processes]

The GREATEST benefit of using AI auditing techniques over traditional methods is that AI auditing techniques can:

A. eliminate the need for human intervention.

B. ensure full compliance with regulations.

C. identify complex data patterns.

D. significantly reduce data bias.

Answer: C

Explanation

AI auditing techniques excel at identifying complex data patterns (option C), which is their primary advantage over manual or traditional audit approaches. The AAIA™ Study Guide states, “AI-based audit tools can process massive volumes of data at speed and depth, detecting anomalies, trends, or relationships that might be invisible to human auditors or unfeasible to uncover manually.”

AI does not fully eliminate the need for human involvement, nor does it guarantee compliance or the elimination of bias, but it can analyze intricate patterns in large, multidimensional data sets.

Reference:ISACA Advanced in AI Audit™ (AAIA™) Study Guide, Section: "Advantages of AI-Enabled Audit Approaches"

Question #:6 - [AI Fundamentals and Technologies]

An IS auditor notes the combined number of records utilized within the training, validation, and testing data sets exceeds the total number of records in the original data set. Which of the following is MOST important for the auditor to determine?

A. Whether the training, validation, and testing data sets were created in the correct order

B. Whether data leakage occurred from utilizing overlapping records in the data sets

C. Whether a sufficient number of records were utilized in the training data set

D. Whether the validation data set utilized the same number of records as the training data sets

Answer: B

Explanation

If the combined size of the training, validation, and testing sets exceeds the original data size, it suggests that records may have been reused across sets. This can lead to data leakage, where the model has access to test or validation information during training, resulting in overly optimistic performance metrics.

“Data leakage invalidates model evaluation because it introduces unintended data overlap. Auditors must ensure that the training, validation, and test sets are strictly partitioned.”

Options A, C, and D refer to process order or quantity, but only B addresses the root issue of compromised model integrity due to overlapping data.

Practice Test Isaca - AAIA

Reference: ISACA Advanced in AI Audit™ (AAIA™) Study Guide, Section: “AI Fundamentals and Technologies,” Subsection: “Data Partitioning and Leakage Risks”

Question #:7 - [AI Governance and Risk Management]

Which of the following is the GREATEST challenge facing IS auditors evaluating the explainability of generative AI models?

A. Differences of opinion regarding model types

B. Difficulties in preventing the input of biased data

C. Performance issues due to excessive computation

D. Algorithms changing as AI continues to learn

Answer: D

Explanation

The greatest challenge for IS auditors in evaluating the explainability of generative AI models is the changing nature of algorithms as AI continues to learn (option D). Generative AI models, especially those using advanced techniques like deep learning and reinforcement learning, often employ continuous or dynamic learning, which results in models that evolve over time. This adaptability can significantly hinder explainability because the logic, parameters, or decision pathways may shift with ongoing retraining or real- time learning.

The ISACA Advanced in AI Audit™ (AAIA™) Study Guide stresses that: “Continually learning AI systems present unique audit challenges, as their internal representations and reasoning can change after deployment, making it difficult to fully capture and explain the rationale for outputs at any given point.”

Other options, such as bias in input data or computational performance, are significant but do not pose as fundamental a challenge to explainability as a model whose internal workings can dynamically change.

Reference:ISACA Advanced in AI Audit™ (AAIA™) Study Guide, Section: "AI Explainability and Dynamic Models"

Question #:8 - [Ethical and Legal Considerations in AI]

An organization is adopting AI for its procurement and inventory teams, raising concern from stakeholders that they will lose their jobs due to AI. Which of the following is the BEST way for the IS auditor to assess whether the potential negative impacts were minimized?

A. Review human-centered design practices to determine how they were considered.

B. Review the AI roadmap for short-term and long-term milestones.

C. Review how the project management team collected feedback in engagement activities.

D. Review the current state assessment of how AI may impact the organization.

Answer: A

Explanation

Human-centered design (HCD) focuses on integrating stakeholder needs, ethical implications, and social impact into AI system development. The AAIA™ Study Guide emphasizes the role of HCD in minimizing negative workforce impacts and ensuring inclusive design.

“Auditors should review whether the AI system design process included human-centered practices— particularly for applications affecting jobs or critical human roles. HCD ensures responsible adoption.”

While feedback collection (C) and impact assessments (D) provide useful context, A directly addresses ethical impact mitigation at the design level.

Reference: ISACA Advanced in AI Audit™ (AAIA™) Study Guide, Section: “Ethical and Legal Considerations in AI,” Subsection: “Human-Centered AI and Workforce Impacts”

Question #:9 - [AI Governance and Risk Management]

Which of the following is the MOST important course of action for an organization prior to allowing end users to utilize an AI tool?

A. Develop an AI policy with guidelines on appropriate use.

B. Determine the impact to the disaster recovery plan (DRP).

C. Implement baseline performance metrics.

D. Ensure a cybersecurity insurance clause is in place to include the use of AI.

Answer: A

Explanation

An AI usage policy sets the foundation for safe, ethical, and effective AI deployment. According to the AAIA™ Study Guide, having an AI policy in place ensures that users understand acceptable behaviors, limitations, and responsibilities when interacting with AI tools.

“AI acceptable use policies promote governance by clearly outlining the dos and don’ts of AI interaction, preventing misuse and aligning user activity with organizational values and compliance standards.”

Other actions (B, C, D) are important in operations and risk management but should follow the establishment of governance protocols through a usage policy. Hence, A is the highest-priority prerequisite.

Reference: ISACA Advanced in AI Audit™ (AAIA™) Study Guide, Section: “AI Governance and Risk Management,” Subsection: “Policy Frameworks for End-User AI Interaction”

Question #:10 - [AI Operations and Performance]

Which of the following correctly summarizes the conclusions of the model card excerpt provided? Model Card – Electrical Grid Predictive Maintenance Model

Model Information:

Description: AI model designed to predict maintenance needs for electrical grid components, reduce unplanned downtime, and improve grid reliability.

Inputs: Real-time sensor data, historical maintenance records, and operational logs. Outputs: Maintenance needs predictions for 60 & 90 days.Evaluation:

Approach: Cross-validation and validation of accuracy, precision, and recall. Results: Accuracy 72%; Precision 60%; Recall 95%; F1 76%

A. The AI model correctly predicts maintenance needs 95% of the time.

B. The electrical grid uptime is expected to be 72% of the time.

C. Grid failure is predicted to occur after 90 days.

D. F1 indicates that the model identifies true maintenance needs 76% of the time.

Answer: D

Explanation

The F1 score is the harmonic mean of precision and recall, offering a balanced measure of model accuracy, especially when there is an imbalance in the classes (e.g., more “no-maintenance” than “needs-maintenance” outcomes). According to the AAIA™ Study Guide, the F1 score is used to evaluate how well the model identifies true positives while balancing the risk of false positives and false negatives.

“An F1 score summarizes the model's ability to correctly identify relevant events—in this case, true maintenance needs. A 76% F1 score means the model is relatively balanced and effective at catching maintenance requirements without generating too many false alerts.”

Thus, D is correct. Option A misrepresents recall as predictive accuracy. Option B misinterprets accuracy. Option C has no direct basis in the excerpt.

Reference: ISACA Advanced in AI Audit™ (AAIA™) Study Guide, Section: “AI Operations and Performance,” Subsection: “Understanding Evaluation Metrics and Model Cards”

About Easypathuni