Vendor: Isaca

Exam Code: AAIA

Exam Name: Advanced in AI Audit (AAIA)

Version: DEMO

QUESTION 1
When converting data categories before training an AI model, which of the following scenarios represents the GREATEST risk?

A. One-hot encoding the data attribute car colors for the options red, blue, green, black, white
B. Creating dummy variables for the data attribute dog breed for the options labrador, terrier, beagle
C. One-hot encoding the data attribute customer rewards category for the options economy, business, first class
D. Creating dummy variables for the data attribute product flavor for the options vanilla, chocolate, strawberry, banana

Answer: C
Explanation:
The AAIATM Study Guide emphasizes that encoding categorical variables must preserve the semantic meaning and order of categories when relevant. The greatest risk occurs when ordinal data--such as customer rewards tiers--is treated as nominal through one-hot encoding, which removes the inherent order and may impair model learning.

"Improper encoding of ordinal variables as nominal can distort the model's understanding of relationships, leading to inaccurate predictions or biased outcomes."

Customer reward categories (economy < business < first class) have a natural order. One-hot encoding ignores this order, potentially degrading model accuracy. Other options represent nominal data and are appropriately encoded.

QUESTION 2
Which of the following BEST detects model drift or unexpected changes in AI model outputs?

A. Standardization of AI configurations
B. Anomaly monitoring
C. AI model documentation reviews
D. AI model retraining

Answer: B
Explanation:
Anomaly monitoring detects irregularities or deviations in input data or model outputs, which are key indicators of model drift. According to the AAIATM Study Guide, continuous anomaly detection is one of the most effective methods for identifying when a model is no longer functioning as expected due to changes in the data environment.

"Monitoring for output anomalies enables early identification of model drift. This proactive approach allows organizations to retrain or adjust models before significant performance degradation occurs."

QUESTION 3
A healthcare organization uses patient data to train an AI model for early disease detection. Which of the following practices provides the BEST assurance that personal data is secure and its integrity is maintained?

A. Encrypting stored data to reduce exposure and log access
B. Updating the AI model with new data and tracking changes
C. Implementing strict data access controls and conducting security tests

D. Anonymizing patient data and performing regular quality checks

Answer: D
Explanation:
In healthcare AI applications, protecting patient data is critical. The AAIATM Study Guide identifies anonymization as one of the most effective strategies to preserve privacy and maintain data integrity. When combined with quality checks, it ensures data accuracy and compliance with health data protection regulations (e.g., HIPAA, GDPR).

"Anonymizing sensitive data removes identifying attributes, significantly reducing risk if data is accessed or leaked. Ongoing data quality checks ensure the integrity and utility of the anonymized dataset."

QUESTION 4
Which of the following controls MOST effectively helps to ensure an AI model is resilient against external threats?

A. AI data set anonymization
B. Monitoring of AI model developers
C. Monitoring of AI access logs
D. AI model configuration testing

Answer: D
Explanation:
Ensuring AI model resilience against external threats involves validating that the model is configured to resist attacks, such as adversarial inputs, data poisoning, or misuse. The AAIATM Study Guide emphasizes configuration testing as a crucial control to simulate threat scenarios and assess robustness.

"Model configuration testing simulates real-world threat conditions to validate model resilience. This includes testing against adversarial attacks, input manipulation, and exposure of sensitive outputs."

While access monitoring (C) and anonymization (A) reduce risks, they don't actively validate model behavior under threat conditions. Therefore, D offers the most effective resilience measure.

QUESTION 5
Which of the following will provide the BEST evidence to support the alignment of an AI model with an organization's business objectives?

A. AI model vulnerability assessment
B. AI change management requests
C. AI model inventory
D. AI acceptable use policy

Answer: C
Explanation:
An AI model inventory documents the models in use, their purposes, and how they support specific business functions. According to the AAIATM Study Guide, maintaining a comprehensive AI model inventory allows auditors to trace model objectives, performance metrics, and use cases back to business goals.

"A well-maintained AI model inventory supports governance and alignment by offering a centralized view of model functions, business integration, and ownership. It ensures transparency and strategic coherence."

While policies and assessments are important, only the inventory directly shows which AI models exist and their connection to organizational objectives.

QUESTION 6
An IS auditor is looking to expedite reporting for an audit with complex issues. Which of the following would be the MOST effective way for the auditor to use generative AI?

A. Developing action items discussed in closing meetings for management action plans
B. Developing a draft of an executive summary based on detailed findings and audit scope
C. Revising audit conclusions with precise verbiage to describe the audit observations
D. Revising audit background and scope information based on new information from management

Answer: B
Explanation:
The AAIATM Study Guide highlights that generative AI tools can significantly enhance audit productivity, especially in content generation tasks. Drafting executive summaries--condensing complex findings into clear, concise formats--is a high-impact use case that benefits from generative AI's natural language generation capabilities.

"Generative AI can be leveraged by auditors to automate reporting sections, such as executive summaries, ensuring consistency and saving time without compromising content clarity."

QUESTION 7
When utilizing a machine learning (ML) model to predict whether a wind turbine electricity generator will fail, which model evaluation metric should be the PRIMARY focus?

A. Precision
B. Specificity
C. Accuracy
D. Recall

Answer: D
Explanation:
In predictive maintenance use cases--such as detecting turbine failure--the most critical concern is identifying as many actual failures as possible to prevent catastrophic events. The AAIATM Study Guide emphasizes that in such high-risk scenarios, Recall is the most appropriate metric because it measures the proportion of true positives correctly identified.

"Recall is critical in scenarios where missing a positive instance (e.g., a failure) is costly or dangerous. It ensures that most real issues are caught by the model, even at the expense of some false positives."

Precision measures correctness of positive predictions, specificity measures true negatives, and accuracy may be misleading if the data is imbalanced. Thus, D (Recall) is most appropriate.

QUESTION 8
Which of the following is the MOST effective way an IS auditor could use generative AI to plan an

audit of a new database storing transactional data?

A. Identifying separation of duties conflicts for database data changes
B. Developing architecture diagrams
C. Identifying technology-specific risk and considerations
D. Summarizing meeting transcripts from interviews with database administrators (DBAs)

Answer: C
Explanation:
Generative AI excels at synthesizing large datasets and technical documentation into understandable insights. The AAIATM Study Guide recommends leveraging generative AI to identify domain-specific risks and control considerations by analyzing complex environments and correlating them with industry risk patterns.

"AI can assist auditors during planning by generating tailored risk profiles for technologies under review, helping prioritize audit focus and scoping."

QUESTION 9
An organization uses an AI-powered tool to detect and respond to cybersecurity threats in real time. An IS auditor finds that the tool produces excessive false positives, increasing the workload of the security team. Which of the following techniques should the auditor recommend to BEST evaluate the tool's effectiveness in managing this issue?

A. Use a log analysis tool to examine the types and frequency of alerts generated.
B. Implement a benchmarking tool to compare the system's alerting capability with industry standards.
C. Conduct penetration testing to assess the system's ability to detect genuine threats.
D. Deploy a machine learning (ML) validation tool to increase the model's accuracy and performance.

Answer: D
Explanation:
The AAIATM Study Guide recommends using validation tools to fine-tune and evaluate ML models, particularly when high false positives undermine operational efficiency. ML validation can identify threshold adjustments, retraining needs, or feature misweighting contributing to excessive alerting.

"Model validation enables organizations to quantify performance, reduce false alarms, and recalibrate AI behavior to align with operational needs and threat landscapes."

QUESTION 10
Which of the following BEST ensures that an AI system complies with user data ownership rights under privacy regulations?

A. Applying data clustering techniques to anonymize data sets
B. Enforcing strict data retention policies to limit storage duration
C. Implementing a transparent data consent management process
D. Regularly conducting AI system performance testing for accuracy

Answer: C
Explanation:
A transparent data consent management process ensures users are informed about how their

data will be used, and enables them to exercise their rights to consent, access, rectify, or delete their data. This is a core requirement under regulations such as GDPR and CCPA.

"Consent management is fundamental to respecting data ownership rights. Organizations must clearly disclose data usage purposes, provide opt-in/out capabilities, and maintain audit trails of user interactions."

QUESTION 11
For a sales promotion, an AI system sorts customer attributes into several categories by analyzing transaction history. Verifying which of the following would BEST validate the effectiveness of this process?

A. Stress tests are regularly conducted to maintain consistent AI performance.
B. The applied methodology adequately reflects business objectives.
C. Sensitive attributes are converted to other data types prior to input.
D. Sampling of AI output is conducted to identify unusual decisions.

Answer: B
Explanation:
The effectiveness of an AI-driven business process--such as categorizing customers for promotional campaigns--depends on how well it supports defined business objectives. The AAIATM Study Guide recommends validating that AI methodology aligns with intended outcomes as part of performance auditing.

"Effectiveness is best measured by assessing whether the AI logic contributes meaningfully to business goals. Output alignment with organizational KPIs or campaign strategies provides clear evidence of functional success."

QUESTION 12
Which of the following is the PRIMARY objective of AI governance?

A. Implementing compliance and ethics controls for AI initiatives
B. Defining clear roles and responsibilities for AI development, use, and oversight
C. Ensuring controls over AI are designed well and operate effectively
D. Promoting a positive return on investment (ROI) from AI projects

Answer: B
Explanation:
The AAIATM Study Guide defines the primary objective of AI governance as establishing structure and accountability for AI initiatives. This includes clearly assigning responsibilities across development, deployment, risk management, and auditing roles to ensure that AI is used responsibly and transparently.

"AI governance establishes the policies, roles, and oversight structures that guide the ethical and secure deployment of AI. Clear accountability helps prevent unauthorized use and ensures strategic alignment."

QUESTION 13
An organization's system development process has been enhanced with AI. Which of the following features presents the GREATEST risk?

A. The AI allocates resources for new system development projects.
B. Non-technical users are validating AI results.
C. The AI personalizes applications for the user.
D. All codes are generated by AI without human oversight.

Answer: D
Explanation:
Allowing AI to autonomously generate code without human review introduces significant risks, including security vulnerabilities, logic errors, and noncompliance with organizational development standards. The AAIATM Study Guide strongly advocates for human-in-the-loop oversight, particularly in automated development contexts.

"AI-assisted development must include manual code reviews to ensure functionality, compliance, and security. Autonomous code generation without validation increases the risk of introducing undetected flaws."

QUESTION 14
An IS auditor is auditing a financial system in which a generative AI tool is used to identify trends in batches of 4,000 rows, while the generative AI tool has a limit of 3,000 tokens. Which of the following is the GREATEST concern?

A. The AI will process only a portion of the data set.
B. The AI will prioritize high-value entries.
C. The AI will reject the data set and not analyze the data.
D. The AI output will be biased toward the first 3,000 tokens.

Answer: D
Explanation:
Generative AI tools such as large language models often have token limitations that restrict how much input they can process in a single prompt. According to the AAIATM Study Guide, when input exceeds token limits, the model processes only the initial portion--leading to biased outputs based on early entries.

"Token limits in generative AI systems may lead to partial input processing, skewing model outputs toward the beginning of the data and introducing interpretive bias. Auditors must assess whether output reflects the full dataset."

QUESTION 15
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 AAIATM 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."

QUESTION 16
While evaluating a complex machine learning (ML) model used for regulatory compliance in a financial institution, which of the following should the IS auditor do to BEST ensure transparency?

A. Document sources and data processes.
B. Create dashboards to show outputs.
C. Provide periodic model audit reports.
D. Use tools that explain model decisions.

Answer: D
Explanation:
Transparency in AI, especially in regulated sectors like finance, is best achieved by using explainability tools that interpret the internal workings and outputs of complex ML models. The AAIATM Study Guide highlights model explainability as essential for both auditors and regulators to understand why certain decisions are made.

"Explainability tools allow auditors to trace predictions back to input features, helping verify fairness, logic, and compliance. These tools support transparency and trustworthiness in black-box models."

QUESTION 17
Which of the following is MOST important to consider when deciding whether to implement an AI solution?

A. The cost of AI implementation
B. The speed of AI implementation
C. The space required for AI hardware
D. The ethical implications of AI

Answer: D
Explanation:
While cost, speed, and infrastructure are practical considerations, the AAIATM Study Guide stresses that ethical implications must be the top priority. Implementing AI without evaluating potential harm, bias, discrimination, or privacy concerns can result in legal, social, and reputational damage.

"Ethical implications should be central to AI implementation decisions. Responsible AI practices involve evaluating societal impact, transparency, accountability, and fairness before deployment."

Therefore, D is the most essential consideration.

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**