Version: 6.2

Question: 1

A financial institution plans to deploy an AI system to provide credit risk assessments for loan applications. Which of the following should be given the HIGHEST priority in the system's design to ensure ethical decision-making and prevent bias?

A. Regularly update the model with new customer data to improve prediction accuracy.

B. Integrate a mechanism for customers to appeal decisions directly within the system.

C. Train the system to provide advisory outputs with final decisions made by human experts.

D. Restrict the model's decision-making criteria to objective financial metrics only.

Answer: C

E xpl anati on:

In AI governance frameworks, credit scoring is treated as a high-risk application. For such systems, the highest-priority safeguard is human oversight to ensure fairness, accountability, and prevention of bias in automated decisions.

The AI Security Management™ (AAISM) domain of AI Governance and Program Management emphasizes that high-impact AI systems require explicit governance structures and human accountability. Human-in-the-loop design ensures that final decisions remain the responsibility of human experts rather than being fully automated. This is particularly critical in financial contexts, where biased outputs can affect individuals' access to credit and create compliance risks.

Official ISACA AI governance guidance specifies:

High-risk AI systems must comply with strict requirements, including human oversight, transparency, and fairness.

The purpose of human oversight is to reduce risks to fundamental rights by ensuring humans can intervene or override an automated decision.

Bias controls are strengthened by requiring human review processes that can analyze outputs and prevent unfair discrimination.

Why other options are not the highest priority:

A . Regular updates improve accuracy but do not guarantee fairness or ethical decision-making. Model drift can introduce new bias if not governed properly.

B . Appeals mechanisms are important for accountability, but they operate after harm has occurred. Governance frameworks emphasize prevention through human oversight in the decision loop.

D . Restricting criteria to “objective metrics” is insufficient, as even objective data can contain hidden proxies for protected attributes. Bias mitigation requires monitoring, testing, and human oversight, not only feature restriction.

AAISM Domain Alignment:

Domain 1 – AI Governance and Program Management: Ensures accountability, ethical oversight, and governance structures.

Domain 2 – AI Risk Management: Identifies and mitigates risks such as bias, discrimination, and lack of transparency.

Domain 3 – AI Technologies and Controls: Provides the technical enablers for implementing oversight mechanisms and bias detection tools.

Reference from AAISM and ISACA materials:

AAISM Exam Content Outline – Domain 1: AI Governance and Program Management (roles, responsibilities, oversight).

ISACA AI Governance Guidance (human oversight as mandatory in high-risk AI applications).

Bias and Fairness Controls in AI (human review and intervention as a primary safeguard).

Question: 2

A retail organization implements an AI-driven recommendation system that utilizes customer purchase history. Which of the following is the BEST way for the organization to ensure privacy and comply with regulatory standards?

A. Conducting quarterly retraining of the AI model to maintain the accuracy of recommendations

B. Maintaining a register of legal and regulatory requirements for privacy

C. Establishing a governance committee to oversee AI privacy practices

D. Storing customer data indefinitely to ensure the AI model has a complete history

Answer: B

E xpl anati on:

According to the AI Security Management™ (AAISM) study framework, compliance with privacy and regulatory standards must begin with a formalized process of identifying, documenting, and maintaining applicable obligations. The guidance explicitly notes that organizations should maintain a comprehensive register of legal and regulatory requirements to ensure accountability and alignment with privacy laws. This register serves as the foundation for all governance, risk, and control practices surrounding AI systems that handle personal data.

Maintaining such a register ensures that the recommendation system operates under the principles of privacy by design and privacy by default. It allows decision-makers and auditors to trace every AI data processing activity back to relevant compliance obligations, thereby demonstrating adherence to laws such as GDPR, CCPA, or other jurisdictional mandates.

Other measures listed in the options contribute to good practice but do not achieve the same direct compliance outcome. Retraining models improves technical accuracy but does not address legal obligations. Oversight committees are valuable but require the documented register as a baseline to oversee effectively. Indefinite storage of customer data contradicts regulatory requirements, particularly the principle of data minimization and storage limitation.

AAISM Domain Alignment:

This requirement falls under Domain 1 – AI Governance and Program Management, which emphasizes organizational accountability, policy creation, and maintaining compliance documentation as part of a structured governance program.

Reference from AAISM and ISACA materials:

AAISM Exam Content Outline – Domain 1: AI Governance and Program Management

AI Security Management Study Guide – Privacy and Regulatory Compliance Controls

ISACA AI Governance Guidance – Maintaining Registers of Applicable Legal Requirements

Question: 3

An organization is updating its vendor arrangements to facilitate the safe adoption of AI technologies. Which of the following would be the PRIMARY challenge in delivering this initiative?

A. Failure to adequately assess AI risk

B. Inability to sufficiently identify shadow AI within the organization

C. Unwillingness of large AI companies to accept updated terms

D. Insufficient legal team experience with AI

Answer: C

E xpl anati on:

In the AAISM™ guidance, vendor management for AI adoption highlights that large AI providers often resist contractual changes, particularly when customers seek to impose stricter security, transparency, or ethical obligations. The official study materials emphasize that while organizations must evaluate AI risk and build internal expertise, the primary challenge lies in negotiating acceptable contractual terms with dominant AI vendors who may not be willing to adjust their standardized agreements. This resistance limits the ability of organizations to enforce oversight, bias controls, and compliance requirements contractually.

Reference:

AAISM Exam Content Outline – AI Risk Management

AI Security Management Study Guide – Third-Party and Vendor Risk

Question: 4

After implementing a third-party generative AI tool, an organization learns about new regulations related to how organizations use AI. Which of the following would be the BEST justification for the organization to decide not to comply?

A. The AI tool is widely used within the industry

B. The AI tool is regularly audited

C. The risk is within the organization's risk appetite

D. The cost of noncompliance was not determined

Answer: C

E xpl anati on:

The AAISM framework clarifies that compliance decisions must always be tied to an organization's risk appetite and tolerance. When new regulations emerge, management may choose not to comply if the associated risk remains within the documented and approved risk appetite, provided that

accountability is established and governance structures support this decision. Other options such as widespread industry use, third-party audits, or lack of cost assessment do not justify noncompliance under the governance principles. The risk appetite framework is the only recognized justification under AI governance principles.

Reference:

AAISM Study Guide – AI Governance and Program Management

ISACA AI Risk Guidance – Risk Appetite and Compliance Decisions

Question: 5

Which of the following is the MOST serious consequence of an AI system correctly guessing the personal information of individuals and drawing conclusions based on that information?

A. The exposure of personal information may result in litigation

B. The publicly available output of the model may include false or defamatory statements about i ndi v i dual s

C. The output may reveal information about individuals or groups without their knowledge

D. The exposure of personal information may lead to a decline in public trust

Answer: C

E xpl anati on:

The AAISM curriculum states that the most serious privacy concern occurs when AI systems infer and disclose sensitive personal or group information without the knowledge or consent of the individuals. This constitutes a direct breach of privacy rights and data protection principles, including those enshrined in GDPR and other global regulations. While litigation, reputational damage, or loss of trust are significant consequences, the unauthorized revelation of personal information through inference is classified as the most severe, because it directly undermines individual autonomy and confidentiality.

Reference:

AAISM Exam Content Outline – AI Risk Management

AI Security Management Study Guide – Privacy and Confidentiality Risks