TestkingIT

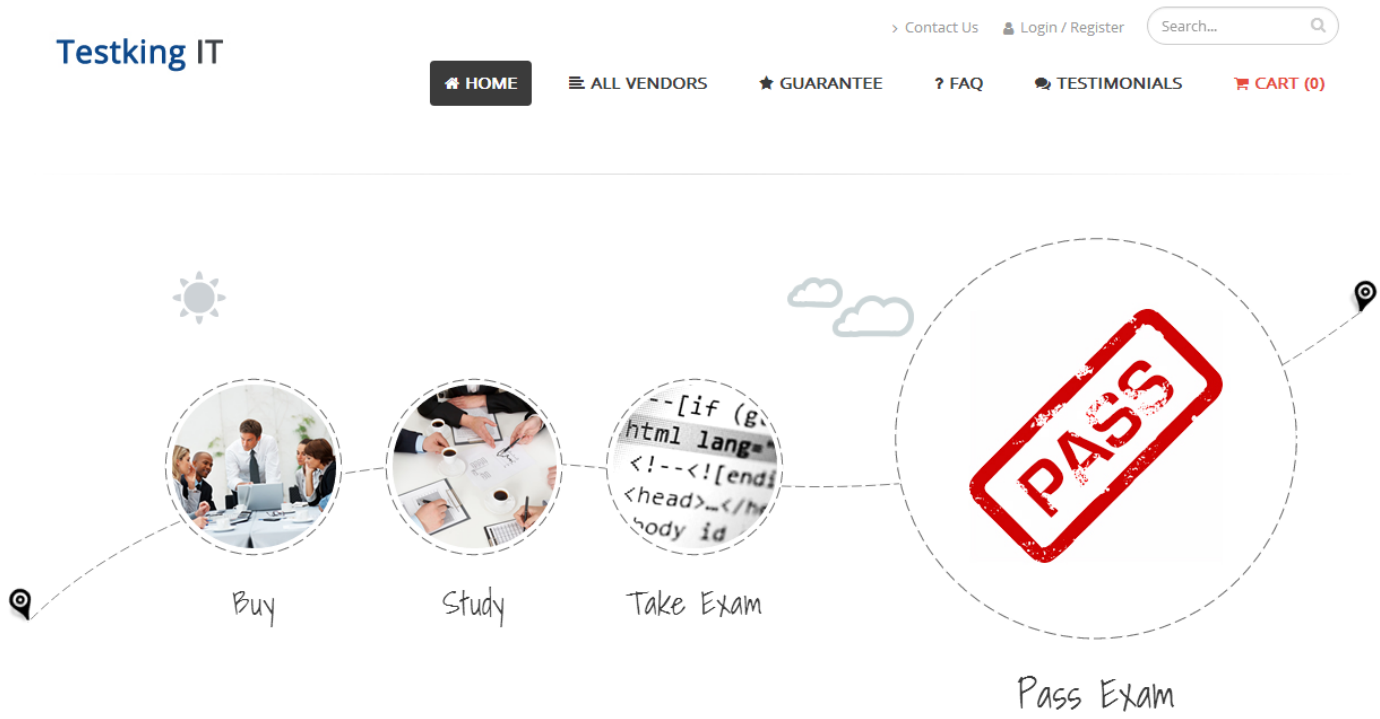

Try **Desktop Test Engine** before you buy

We're not the only ones **happy** about TestKingsIT Practice Material ...

48236+ customers in 100+ countries use TestKingsIT Test Engine. Meet our customers.

iMessenger VOREED GetCustom JET ORANGE iCompany Paradoxx

http://www.testkingit.com/

Latest practice material - Exam Cram - TestKingIT

Exam : **AAISM**

Title : ISACA Advanced in AI Security Management (AAISM) Exam

Vendor : ISACA

Version : DEMO

QUESTION NO: 1
Which of the following is the MOST important factor to consider when selecting industry frameworks to align organizational AI governance with business objectives?
A. Risk tolerance
B. Risk threshold
C. Risk register
D. Risk appetite
Answer: D
Explanation:
According to AAISM governance principles, the risk appetite of the organization is the most important factor in selecting appropriate frameworks for AI governance. Risk appetite defines the level of risk an organization is willing to accept in pursuit of its objectives, ensuring frameworks are aligned with strategic goals. Risk tolerance and thresholds are operational measures derived from appetite, and the risk register is a documentation tool. The foundational consideration for framework alignment is the organization's risk appetite.
References:
AAISM Exam Content Outline - AI Governance and Program Management (Risk Appetite in Governance Alignment) AI Security Management Study Guide - Framework Selection and Business Strategy

QUESTION NO: 2
A data scientist creating categories and training an algorithm on large data sets is performing which learning technique?
A. Supervised
B. Reinforcement
C. Unsupervised
D. Machine learning (ML)
Answer: A
Explanation:
AAISM identifies supervised learning as involving:
* labeled categories
* ground-truth datasets
* model training with known outcomes
This aligns exactly with categorizing data and training on labeled datasets.
Reinforcement (B) involves reward feedback loops. Unsupervised (C) uses unlabeled data.
"Machine learning" (D) is too broad and not a specific technique.
References: AAISM Study Guide - AI Learning Types; Supervised Learning Definition.

QUESTION NO: 3
An organization is reviewing an AI application to determine whether it is still needed. Engineers have been asked to analyze the number of incorrect predictions against the total number of predictions made. Which of the following is this an example of?
A. Control self-assessment (CSA)
B. Model validation

C. Key performance indicator (KPI)

D. Explainable decision-making

Answer: C

Explanation:

AAISM guidance identifies metrics like error rate versus total predictions as a key performance indicator (KPI) for evaluating AI model effectiveness. KPIs provide measurable values to assess performance against objectives. Model validation is broader and occurs prior to production use, testing the model against predefined standards. Control self-assessment relates to governance processes, not predictive accuracy.

Explainable decision-making refers to interpretability, not error-rate evaluation. Thus, analyzing incorrect predictions against total predictions is a performance measure, making it a KPI.

References:

AAISM Exam Content Outline - AI Governance and Program Management (Performance Metrics and KPIs) AI Security Management Study Guide - Accuracy and Error Metrics

QUESTION NO: 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

Explanation:

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.

References:

AAISM Study Guide - AI Governance and Program Management

ISACA AI Risk Guidance - Risk Appetite and Compliance Decisions

QUESTION NO: 5

Which of the following BEST describes an adversarial attack on an AI model?

A. Attacking underlying hardware

B. Providing inputs that mislead the model into incorrect predictions

C. Reverse-engineering the model using social engineering

D. Conducting denial-of-service attacks on AI APIs

Answer: B
Explanation:
AAISM defines adversarial attacks as manipulations of input data (text, image, audio, numeric values) designed to cause the model to produce incorrect or harmful predictions. Hardware attacks (A) are infrastructure threats. Social engineering (C) targets people, not models. DoS attacks (D) affect availability, not model decision pathways.
References: AAISM Study Guide - Adversarial Threats; Input Manipulation.

QUESTION NO: 6
Which of the following employee awareness topics would MOST likely be revised to account for AI-enabled cyber risk?
A. Clean desk policy
B. Social engineering
C. Malicious insider threats
D. Authentication controls
Answer: B
Explanation:
AAISM training guidance specifies that social engineering is the awareness topic most impacted by AI- enabled risks. With generative AI and deepfake technologies, attackers can create highly convincing phishing messages, synthetic voices, or fake executive requests, increasing the sophistication of social engineering attacks. Clean desk policies, insider threat awareness, and authentication procedures remain relevant but are not directly altered by AI advancements. The most likely revision to employee awareness programs in the AI era is therefore enhanced social engineering awareness.
References:
AAISM Exam Content Outline - AI Risk Management (Human Factors and Awareness) AI Security Management Study Guide - Social Engineering Risks with AI

QUESTION NO: 7
Which of the following is the GREATEST benefit of implementing an AI tool to safeguard sensitive data and prevent unauthorized access?
A. Timely analysis of endpoint activities
B. Timely initiation of incident response
C. Reduced number of false positives
D. Reduced need for data classification
Answer: C
Explanation:
The AAISM study materials highlight that AI-powered security tools provide the greatest benefit by reducing false positives in monitoring and access control systems. This improves efficiency, prevents alert fatigue, and enables security teams to focus on true threats. While timely analysis and incident response are benefits, they are not unique to AI-based tools and can be achieved with traditional methods. AI also does not remove the need for data classification, as classification underpins governance and compliance. The standout advantage is the improved accuracy and reduced false positives provided by AI.

References:
AAISM Study Guide - AI Technologies and Controls (Security Tools and Access Management) ISACA AI Security Management - Benefits of AI-Enabled Security

QUESTION NO: 8
To ensure AI tools do not jeopardize ethical principles, it is MOST important to validate that:
A. The organization has implemented a responsible development policy
B. Outputs of AI tools do not perpetuate adverse biases
C. Stakeholders have approved alignment with company values
D. AI tools are evaluated by the privacy department before implementation
Answer: B
Explanation:
AAISM highlights that the core ethical risk in AI is the perpetuation of bias that results in unfair or discriminatory outcomes. Therefore, the most important validation step is ensuring that outputs of AI systems are free from adverse biases. A responsible development policy, stakeholder approvals, and privacy reviews all contribute to governance, but they do not directly ensure ethical outcomes. Validation of output fairness is the critical safeguard for ensuring AI does not violate ethical principles.
References:
AAISM Study Guide - AI Risk Management (Bias and Ethics Validation)
ISACA AI Security Management - Ethical AI Practices

QUESTION NO: 9
Which of the following is the MOST important course of action prior to placing an in-house developed AI solution into production?
A. Perform a privacy, security, and compliance gap analysis
B. Deploy a prototype of the solution
C. Obtain senior management sign-off
D. Perform testing, evaluation, validation, and verification
Answer: D
Explanation:
AAISM lifecycle governance guidance specifies that before any AI solution is moved into production, it must undergo testing, evaluation, validation, and verification to ensure accuracy, resilience, security, and compliance with standards. These steps confirm that the solution performs as expected under varied conditions. Conducting gap analysis is part of compliance checks but comes earlier in design. Management sign-off provides approval but cannot substitute for assurance of technical reliability. Deploying prototypes is a testing method but not the final assurance step. The critical requirement is a complete cycle of testing, validation, and verification.
References:
AAISM Exam Content Outline - AI Risk Management (Lifecycle Testing and Validation) AI Security Management Study Guide - Production Readiness Checks

QUESTION NO: 10
An organization plans to use AI to analyze the shopping patterns of its customers to predict

interests and send targeted, customized marketing emails. Which of the following should be done FIRST?

A. Obtain customer consent

B. Train the marketing department

C. Update the terms of service

D. Verify customer email addresses

Answer: A

Explanation:

The first action, before any processing of personal data for AI-driven profiling and targeted communications, is to establish a lawful basis for processing. Under AAISM-aligned privacy governance, explicit and informed consent is prioritized for new or sensitive uses such as interest profiling and targeted marketing. Consent ensures purpose limitation, transparency, and user control prior to model ingestion and campaign activation.

Training teams, updating terms of service, or verifying contact details are important, but they do not provide legal authority to process data; therefore, they follow after consent is obtained.

References: AI Security Management™ (AAISM) Body of Knowledge - Privacy Governance and Lawful Basis; Purpose Limitation and Transparency; Consent Management in AI-enabled Marketing. AAISM Study Guide - Data Protection Controls for AI Profiling; Consent Capture and Record-Keeping.

QUESTION NO: 11

A security assessment revealed that attackers could access sensitive company data through chat interface injection. What is the BEST mitigation?

A. Conducting regular security audits

B. Manually reviewing AI model outputs

C. Implementing input validation and templates

D. Ensuring continuous monitoring and tagging

Answer: C

Explanation:

AAISM explains that prompt injection attacks are best mitigated by:

* strict input validation
* templated prompts
* controlled context windows
* guardrail enforcement

These prevent malicious instructions from overriding system prompts.

Audits (A) are periodic, not preventive. Manual review (B) is not scalable. Monitoring (D) detects issues but does not block injection.

References: AAISM Study Guide - Prompt Injection & Input Control Mechanisms.

QUESTION NO: 12

Which of the following is the MOST effective way to identify and address security risk in an AI model?

A. Assign staff to review AI model outputs for accuracy

B. Conduct threat modeling to identify vulnerabilities and possible attack methods

C. Encrypt the training data and model parameters to prevent unauthorized access

D. Add more data to the model to increase its accuracy and reduce errors

Answer: B

Explanation:

AI/ML threat modeling is the most effective structured method to both identify and address model security risks. It systematically surfaces attack classes (poisoning, evasion, membership inference, model extraction, inversion), maps system-specific attack surfaces (data pipelines, feature stores, training artifacts, inference APIs), and drives prioritized mitigations (ingestion validation, robust training, rate-limiting, watermarking, differential privacy, monitoring, red teaming). Output spot-checking (A) finds errors but not security vulnerabilities; encryption (C) protects confidentiality but does not reveal threats or mitigate inference-time attacks; adding data (D) may improve accuracy but does not target adversarial risk.

References: AI Security Management™ (AAISM) Body of Knowledge - AI Risk Identification & Threat Modeling; Attack Surface Analysis for ML; Risk Treatment Planning. AAISM Study Guide - Evasion /Poisoning/Extraction Controls; Mapping Risks to Controls; Validation and Assurance Activities.

QUESTION NO: 13

Which of the following is the BEST way to ensure an organization remains compliant with industry regulations when decommissioning an AI system used to record patient data?

A. Ensure backups are tested and access controls are recorded and audited to ensure compliance

B. Update governance policies based on lessons learned and ensure a feedback loop exists

C. Perform a post-destruction risk assessment to verify that there is no residual exposure of data

D. Ensure the certificate of destruction is received and archived in line with data retention policies

Answer: D

Explanation:

For regulated data such as patient information, AAISM requires provable data lifecycle closure at decommissioning. The authoritative evidence is a certificate of destruction (covering primary, replicas, backups, and caches) retained per the organization's records retention policy. While testing backups and auditing access (A), updating policies (B), and doing post-destruction risk assessment (C) are valuable practices, documented destruction attestation is the primary compliance proof point that the data was disposed of in accordance with regulatory and contractual obligations.

References: AI Security Management™ (AAISM) Body of Knowledge - Data Lifecycle Governance; Decommissioning & Secure Disposal; Records Retention and Evidence of Destruction.

QUESTION NO: 14

Which of the following AI data life cycle phases presents the GREATEST inherent risk?

A. Training

B. Maintenance
C. Monitoring
D. Preparation
Answer: D
Explanation:
The data Preparation phase-covering sourcing, collection, labeling, cleansing, and provenance-presents the greatest inherent risk because it is where privacy, consent, representativeness, bias, quality, lineage, and legality must be established. Decisions and defects here propagate into training and downstream use, amplifying ethical, regulatory, and accuracy risks.
While Training can introduce additional risks (e.g., poisoning, leakage), these are frequently mitigated by controls that depend on having trustworthy prepared data. Monitoring and Maintenance are later-life-cycle phases oriented toward detection and correction; they are critical but inherently rely on the foundation set during Preparation.
References: AI Security Management™ (AAISM) Body of Knowledge: "AI Data Lifecycle Risks-Sourcing, Consent, Provenance, and Bias," "Risk Treatment Priorities Across the Lifecycle"; AAISM Study Guide:
"Data Preparation Controls and Quality Gates," "Bias and Privacy Risk Controls at Ingestion."

QUESTION NO: 15
An aerospace manufacturing company that prioritizes accuracy and security has decided to use generative AI to enhance operations. Which of the following large language model (LLM) adoption plans BEST aligns with the company's risk appetite?
A. Developing a public LLM to automate critical functions
B. Purchasing an LLM dataset on the open market
C. Contracting LLM access from a reputable third-party provider
D. Developing a private LLM to automate non-critical functions
Answer: D
Explanation:
AAISM recommends aligning AI adoption with organizational risk appetite by limiting blast radius, protecting sensitive data, and staging adoption in lower-risk domains first. Building a private LLM for non- critical functions preserves data control, enables tighter governance (access control, logging, evaluation), and confines any model errors away from safety- or mission-critical operations. A public LLM for critical functions (A) is misaligned with a high-assurance posture; buying open-market datasets (B) raises provenance and licensing risk; third-party access (C) can be appropriate but still introduces vendor/visibility limits and data residency concerns that may not meet aerospace security needs.
References: AI Security Management™ (AAISM) Body of Knowledge - Risk Appetite Mapping to AI Use Cases; Criticality Segmentation; Data Control & Deployment Models. AAISM Study Guide - Phased Adoption for High-Assurance Environments; Private vs. Hosted LLM Trade -offs; Governance, Evaluation, and Containment Patterns.

QUESTION NO: 16
A global organization experienced multiple incidents of staff pasting confidential data into public chatbots.

Which action is MOST important to reduce short-term risk?

A. Deliver role-based, scenario-driven AI security training mapped to job functions

B. Require employees to complete an annual generic phishing and deepfake module

C. Publish an AI acceptable use policy and collect signatures

D. Block access to public LLMs at the network perimeter

Answer: A

Explanation:

AAISM states that the most effective short-term mitigation for unintentional data leakage into public AI tools is targeted, role-based AI security awareness, focused on:

* what data cannot be entered
* consequences of leakage
* real-world scenarios employees face

An acceptable-use policy (C) is necessary but insufficient alone. Blocking LLMs (D) may reduce access but does not change user behavior and may cause shadow AI usage. Generic phishing training (B) does not address AI misuse risks.

References: AAISM Study Guide - AI Security Awareness; Human-Centric Data Leakage Prevention.

QUESTION NO: 17

An organization uses an AI tool to scan social media for product reviews. Fraudulent social media accounts begin posting negative reviews attacking the organization's product. Which type of AI attack is MOST likely to have occurred?

A. Model inversion

B. Deepfake

C. Availability attack

D. Data poisoning

Answer: C

Explanation:

The AAISM materials classify availability attacks as attempts to disrupt or degrade the functioning of an AI system so that its outputs become unreliable or unusable. In this scenario, the fraudulent social media accounts are deliberately overwhelming the AI tool with misleading negative reviews, undermining its ability to deliver accurate sentiment analysis. This aligns directly with the concept of an availability attack. Model inversion relates to reconstructing training data from outputs, deepfakes involve synthetic content generation, and data poisoning corrupts the training set rather than manipulating inputs at runtime. Therefore, the fraudulent review campaign is most accurately identified as an availability attack.

References:

AAISM Study Guide - AI Risk Management (Adversarial Threats and Availability Risks)

ISACA AI Security Management - Attack Classifications

QUESTION NO: 18

An organization is implementing AI agent development across engineering teams. What should AI-specific training focus on?

A. Prompt injection, agent memory control, insecure tool execution
B. Dataset bias, explainability, fairness
C. Output moderation, hallucination handling, policy alignment
D. API abuse, data leakage, third-party plug-in risk
Answer: A
Explanation:
AAISM states that AI agent security training should focus on the unique risks of agentic systems, which include:
* prompt injection
* memory control and context hijacking
* unsafe tool execution (agents triggering unauthorized actions)
These risks are specific to autonomous or semi-autonomous AI agents.
Bias, fairness (B) and output moderation (C) are important but not the most critical for agent security. API abuse and plug-in risk (D) matter but are secondary.
References: AAISM Study Guide - Agentic AI Security; Prompt Injection and Tool Execution Risks.

QUESTION NO: 19
Within an incident handling process, which of the following would BEST help restore end user trust with an AI system?
A. The AI model prioritizes incidents based on business impact
B. AI is being used to monitor incident detection and alerts
C. The AI model's outputs are validated by team members
D. Remediation of the AI system based on lessons learned
Answer: C
Explanation:
Restoring end user trust during incident handling requires visible, immediate assurance that system outcomes are safe and appropriate. AAISM prescribes human oversight and approval gates for high-risk AI decisions, with human validation of outputs before use as a primary control to maintain trust while technical remediation is underway. Prioritization (A) and monitoring (B) aid operations but do not directly rebuild user confidence in outcomes. Post-incident improvements (D) are essential for long-term assurance but do not provide the immediate trust restoration that supervised, human-validated outputs deliver.
References: AI Security Management™ (AAISM) Body of Knowledge - Incident Handling & Communications; Human Oversight and Approval Gates; Trust Restoration During AI Incidents.

QUESTION NO: 20
In the context of generative AI, which of the following would be the MOST likely goal of penetration testing during a red-teaming exercise?
A. Generate outputs that are unexpected using adversarial inputs
B. Stress test the model's decision-making process
C. Degrade the model's performance for existing use cases
D. Replace the model's outputs with entirely random content

Answer: A
Explanation:
AAISM's risk management content describes red-teaming in generative AI as focused on deliberately crafting adversarial prompts to test whether the model produces unexpected or undesired outputs that violate safety, integrity, or compliance standards. The goal is not to stress system performance or randomly disrupt outputs, but rather to uncover vulnerabilities in how the model responds to manipulative inputs. This allows organizations to improve resilience against prompt injection, jailbreaking, or harmful content generation. The correct answer is therefore generate outputs that are unexpected using adversarial inputs.
References:
AAISM Exam Content Outline - AI Risk Management (Red-Team Testing and Adversarial Exercises) AI Security Management Study Guide - Penetration Testing in Generative AI Contexts

QUESTION NO: 21
The PRIMARY purpose of adopting and implementing AI architecture as part of an organizational AI program is to:
A. ensure the development of powerful, efficient, and scalable AI systems
B. deploy fast and cost-efficient AI systems for rapidly changing environments
C. align the system components of AI with the business goals of the organization
D. provide a basis for identification of threats and vulnerabilities
Answer: C
Explanation:
An AI architecture, within program governance, exists to align AI system components and lifecycle processes with business goals and policy constraints. Architecture provides the organizing structure linking strategy, capabilities, processes, data, models, controls, and assurance so that AI outcomes are traceable to business value, risk appetite, and compliance expectations. Efficiency, speed, and threat analysis are important architectural qualities, but they are not the primary purpose; the primary purpose is strategic and governance alignment so that technical choices and controls consistently realize organizational objectives.
References:* AI Security Management™ (AAISM) Body of Knowledge: AI Program Architecture - alignment of capabilities, processes, and controls to business objectives* AI Security Management™ Study Guide: Architecture-driven governance, traceability from business goals to technical and control design

QUESTION NO: 22
Which of the following is the MOST likely cause of model drift?
A. Data poisoning
B. Perfect knowledge
C. Membership inference
D. Model stealing
Answer: A
Explanation:

Model drift occurs when the statistical properties of input data and/or the relationship between features and outcomes change over time, causing degraded model performance. The AAISM guidance classifies data- centric causes (distribution shift, concept drift, and contamination) as the primary drivers and highlights that malicious contamination of training or incremental learning data (data poisoning) is a direct, high-likelihood driver of observable drift in production because it changes the effective data-generating process the model learns from. In contrast:
* Perfect knowledge is an attacker capability descriptor, not a drift cause.
* Membership inference targets privacy of the training set and does not inherently shift data distributions.
* Model stealing targets IP/confidentiality; it does not change the victim model's data distribution or decision boundary in situ.
References:* AI Security Management™ (AAISM) Body of Knowledge: Model Risk & Drift; Data Integrity Risks; Adversarial ML-Poisoning vs. Evasion* AAISM Study Guide: Production Monitoring & Drift Management; Risk Scenarios-Data Poisoning Impacts and Controls* AAISM Mapping to Standards:
Lifecycle Risk Treatment-Robustness to Data Contamination; Continuous Monitoring and Feedback

QUESTION NO: 23
As organizations increasingly rely on vendors to develop AI systems, which of the following is the MOST effective way to monitor vendors and ensure compliance with ethical and security standards?
A. Conducting regular audits of vendor processes and adherence to AI development guidelines
B. Requiring vendors to monitor their adherence to ethics and security standards
C. Mandating that vendors share source code and AI documentation with the contracting party
D. Allowing vendors to self-attest ethical AI compliance and implement benchmark monitorin g
Answer: A
Explanation:
AAISM vendor governance guidance identifies regular audits of vendor processes as the most effective method of ensuring compliance with ethical and security standards. Independent audits provide verifiable assurance that vendors are meeting agreed-upon requirements. Self-attestation, internal monitoring, or documentation sharing provide some transparency but do not guarantee compliance. The best practice, particularly for high-risk AI deployments, is independent and recurring audits of vendor processes.
References:
AAISM Exam Content Outline - AI Risk Management (Vendor Oversight)
AI Security Management Study Guide - Vendor Audit and Compliance Assurance

QUESTION NO: 24
How can an organization BEST protect itself from payment diversions caused by deepfake attacks impersonating management?

A. Require mandatory deepfake detection training for all employees
B. Mandate that payments be sent only once per week
C. Issue a security policy on deepfakes
D. Implement resilient payment approval processes
Answer: D
Explanation:
AAISM's risk management framework stresses that the most effective defense against deepfake-enabled fraud, such as payment diversion, is resilient payment approval processes. This includes multi-step verification, segregation of duties, and independent confirmations for high-value transactions. Employee training, policies, or limiting payment frequency may reduce exposure, but they cannot guarantee prevention.
Only process-based controls enforce structural safeguards that prevent fraudulent instructions from being executed, even if a deepfake impersonation attempt is successful.
References:
AAISM Exam Content Outline - AI Risk Management (Fraud and Deepfake Risk) AI Security Management Study Guide - Transactional Resilience and Controls

QUESTION NO: 25
After deployment, an AI model's output begins to drift outside of the expected range. Which of the following is the development team's BEST course of action?
A. Take the AI model offline
B. Adjust the hyperparameters of the AI model
C. Create an emergency change request to correct the issue
D. Return to an earlier phase in the AI life cycle
Answer: D
Explanation:
AAISM emphasizes that when model drift occurs, the best response is not a quick fix but rather to revisit an earlier phase of the AI life cycle to address data quality, retraining, or evaluation processes. Simply taking the model offline halts functionality without resolution, while adjusting hyperparameters or issuing emergency changes treats the symptom rather than the root cause. Proper governance requires returning to the design or training phases to re-establish stability, accuracy, and compliance of the model. Thus, the correct approach is to return to an earlier AI lifecycle phase.
References:
AAISM Exam Content Outline - AI Risk Management (Model Drift and Lifecycle Responses) AI Security Management Study Guide - Continuous Improvement in AI Lifecycle

QUESTION NO: 26
When using AI as part of incident response, which of the following BEST ensures the automation aligns with regulatory and governance obligations?
A. Use deep learning models to autonomously classify all incidents
B. Train the AI incident response platform to mirror legacy response workflows and log containment
C. Apply anomaly detection models to filter incoming threats and automate containment

D. Implement a tiered automation strategy where severity ratings inform the need for human oversight

Answer: D

Explanation:

AAISM prescribes risk-based, human-in-the-loop orchestration for safety-critical or regulated actions. A tiered automation strategy that gates autonomy by incident severity, data sensitivity, and regulatory requirements ensures accountability, auditability, and proportionality, satisfying governance obligations. Full autonomy (A) risks non-compliance; simply mirroring legacy workflows (B) may not meet current obligations; broad auto-containment (C) lacks necessary oversight controls.

References: AI Security Management™ (AAISM) Body of Knowledge - Governance of AI -Driven Security Automation; Human Oversight and Escalation; Risk-Based Orchestration. AAISM Study Guide - Incident Response with AI: Controls, Approvals, and Auditability.