

ISACA.AAISM.v2026-09-07.q200

Exam Code:	AAISM
Exam Name:	ISACA Advanced in AI Security Management (AAISM) Exam
Certification Provider:	ISACA
Free Question Number:	200
Version:	v2026-09-07
# of views:	112
# of Questions views:	2154
https://www.freecram.net/torrent/ISACA.AAISM.v2026-09-07.q200.html	

NEW QUESTION: 1

Which of the following is the BEST way for an organization to mitigate the use of deepfakes for social engineering?

- A. Develop biometric authentication controls.
- B. Implement a robust business continuity plan (BCP).
- C. Increase investment in the security operations center (SOC).
- D. Require an additional method of identity verification.

Answer: (SHOW ANSWER)

Deepfakes can convincingly mimic voices or appearances, making single-factor verification unreliable. Requiring an additional independent method of identity verification ensures that requests are validated through multiple channels, reducing the effectiveness of social engineering attempts.

NEW QUESTION: 2

A large corporation has received an influx of sophisticated credential-phishing emails and wants to leverage an AI solution to detect and quarantine these messages before they reach employees. Which of the following blue-team AI features is BEST suited to this task?

- A. Large language model (LLM)
- B. Natural language processing (NLP)
- C. Natural language generation (NLG)
- D. Retrieval-augmented generation (RAG)

Answer: (SHOW ANSWER)

For pre-delivery phishing detection and classification, the most appropriate capability is NLP--tokenization, feature extraction, semantic similarity, and supervised classifiers (e.g., transformer-based classifiers) tuned on phishing corpora and indicators. NLP models score messages and drive automated quarantine policies.

NEW QUESTION: 3

A critical AI system shows biased outcomes. What is the BEST course of action?

- A. Activate the kill switch
- B. Conduct audits of data and model
- C. Perform root cause analysis to identify mitigation
- D. Retrain the model with a new diverse dataset

Answer: ([SHOW ANSWER](#))

AAISM emphasizes root cause analysis as the primary step when AI behavior deviates unexpectedly.

RCA determines whether issues stem from:

- data drift
- model drift
- poisoned data
- missing features
- logic errors
- configuration changes

NEW QUESTION: 4

A newly hired programmer suspects that the organization's AI solution is inferring users' sensitive information and using it to advise future decisions. Which of the following is the programmer's BEST course of action?

- A. Conduct a code review
- B. Alert the CIO to the risk
- C. Suggest fine-tuning the AI solution
- D. Inform the governance panel

Answer: ([SHOW ANSWER](#))

AAISM directs personnel to use established AI governance channels for suspected privacy, ethics, or compliance risks. The governance panel (risk, privacy, legal/compliance, security, product/data science) is chartered to triage, record, investigate, and direct remediation for potential inference of sensitive attributes and resulting decision impacts. Direct technical action (A or C) bypasses due process and accountability; escalating directly to a single executive (B) lacks the structured, cross-functional oversight required for regulated and ethical AI risk handling.

NEW QUESTION: 5

During which phase of the AI life cycle are models assessed for security and determined to be free of malicious manipulations?

- A. Verification
- B. Validation
- C. Evaluation
- D. Testing

Answer: ([SHOW ANSWER](#))

The verification phase focuses on ensuring the model has been built correctly and is free from tampering or malicious manipulation by checking integrity, security controls, and adherence to design specifications before deployment.

NEW QUESTION: 6

Which of the following is the MOST effective defense against cyberattacks that alter input data to avoid detection by the model?

- A. Conducting periodic monitoring activities on the model's decisions
- B. Enhancing model robustness through adversarial training
- C. Implementing restricted access to the model's internal parameters
- D. Applying differential privacy controls on training datasets

Answer: ([SHOW ANSWER](#))

Evasion attacks manipulate inputs to induce misclassification while leaving the model unchanged. AAISM prescribes adversarial robustness controls, with adversarial training as a primary measure:

incorporate adversarially perturbed examples into training/validation to harden decision boundaries and improve resilience across threat models (e.g., Lp-bounded perturbations).

NEW QUESTION: 7

An organization deploying an LLM is concerned input manipulations could compromise security. What is the MOST effective way to determine an acceptable risk threshold?

- A. Deploy real-time logging and monitoring
- B. Restrict all inputs containing special characters
- C. Assess the business impact of known threats
- D. Implement a static threshold limiting LLM outputs

Answer: ([SHOW ANSWER](#))

AAISM instructs that acceptable risk thresholds must be determined using business impact analysis. This aligns with the broader enterprise risk management principle of defining tolerances based on:

- potential harm
- regulatory exposure
- financial impact
- operational disruption

NEW QUESTION: 8

Which of the following is the MOST effective use of model confidence scores?

- A. Evaluating uncertainty in AI model outputs to classify results requiring action
- B. Triggering secondary verification of resilience to external tampering
- C. Assigning probabilities to predictive patterns in training data
- D. Managing user expectations regarding model precision and accuracy

Answer: ([SHOW ANSWER](#))

Model confidence scores are most effectively used to measure the AI model's level of certainty in its outputs. This helps identify uncertain or high-risk results that may require human review, additional validation, or alternative handling before action is taken.

NEW QUESTION: 9

Which of the following is the MOST effective way to identify deepfakes generated by AI?

- A. Examining irregularities in metadata or data traces
- B. Detecting irregular pixel patterns or distorted frames
- C. Analyzing biometric inconsistencies
- D. Comparing the model card to the output generated by AI

Answer: ([SHOW ANSWER](#))

Analyzing biometric inconsistencies is the most effective method because deepfakes often fail to perfectly replicate natural human physiological and behavioral characteristics such as facial movements, blinking patterns, voice characteristics, or lip synchronization. These inconsistencies provide strong indicators of AI-generated manipulation.

NEW QUESTION: 10

Which of the following BEST addresses risk associated with hallucinations in AI systems?

- A. Recursive chunking
- B. Automated output validation
- C. Content enrichment
- D. Human oversight

Answer: ([SHOW ANSWER](#))

AAISM prescribes human-in-the-loop (HITL) controls as the primary safeguard for high-impact generative AI use cases to mitigate hallucination risk. Human oversight ensures critical outputs are reviewed, corrected, and approved before use, with accountability, escalation, and documented decision trails. Automated validators and enrichment help reduce errors but are secondary; recursive chunking is a prompting tactic, not a governance control.

NEW QUESTION: 11

Which of the following would BEST help an organization align its AI initiatives with business objectives?

- A. Complying with applicable AI-related regulations
- B. Ensuring ethical use of AI technologies in projects
- C. Establishing an AI governance committee
- D. Protecting enterprise information used by AI projects

Answer: ([SHOW ANSWER](#))

An AI governance committee provides cross-functional oversight to align AI strategy, investment, and risk appetite with business goals. It sets policies, prioritizes portfolios, ensures accountability, and integrates compliance, ethics, and security into decision-making. While compliance, ethics,

and information protection are essential, governance is the primary mechanism that systematically connects AI initiatives to enterprise objectives.

NEW QUESTION: 12

Which defense is MOST effective against cyberattacks that alter input data to avoid detection?

- A. Enhancing model robustness through adversarial training
- B. Restricting access to internal model parameters
- C. Conducting periodic monitoring of decisions
- D. Applying differential privacy to training data

Answer: ([SHOW ANSWER](#))

AAISM lists adversarial training as the strongest method to harden models against input manipulation attacks. By exposing models to adversarial examples during training, the system learns to resist evasion techniques.

NEW QUESTION: 13

Which of the following would BEST ensure safe implementation of AI systems?

- A. Selection of AI projects that align with business objectives.
- B. Compliance with the AI risk management framework and applicable regulations.
- C. Establishment of an AI program steering committee to govern AI implementation.
- D. Phased implementation with human-in-the-loop validation before full rollout.

Answer: ([SHOW ANSWER](#))

A phased implementation combined with human-in-the-loop validation allows the organization to monitor AI behavior, identify issues early, validate outputs, and maintain human oversight before full deployment. This approach most effectively ensures safe and controlled implementation of AI systems.

NEW QUESTION: 14

An organization is implementing a large language model (LLM) to answer questions for both employees and customers. The CIO wants the LLM to provide employees with additional internal data that should not be shared with customers. Which of the following is the MOST secure approach?

- A. Retrieval-augmented generation (RAG) with access control lists (ACLs)
- B. Model fine tuning with increased use of function calls
- C. Prompt sanitization
- D. Allow list employee access

Answer: ([SHOW ANSWER](#))

Retrieval-augmented generation (RAG) combined with access control lists (ACLs) ensures that the LLM retrieves and exposes only the data each user is authorized to access. This approach provides granular, role-based protection of internal information while supporting different data visibility requirements for employees and customers.

NEW QUESTION: 15

Which of the following methods provides the MOST effective protection against model inversion attacks?

- A. Using adversarial training
- B. Reducing the model's complexity
- C. Implementing regularization output
- D. Increasing the number of training iterations

Answer: C ([LEAVE A REPLY](#))

AAISM classifies model inversion as a privacy leakage threat where adversaries infer sensitive attributes or training records from model outputs. The recommended technical risk treatments emphasize reducing overfitting and information leakage via regularization and output-side constraints. Regularization (e.g., stronger penalties, output smoothing, confidence calibration, temperature limiting, and related techniques) reduces the model's tendency to memorize training data and curtails exploitable signal in outputs.

NEW QUESTION: 16

A CISO has been tasked with providing key performance indicators (KPIs) on the organization's newly launched AI chatbot. Which of the following are the BEST metrics for the CISO to recommend?

- A. Explainability and F1 score
- B. Customer effort score and user retention rate
- C. Response time and throughput
- D. Error rate and bias detection

Answer: ([SHOW ANSWER](#)**)**

For executive security and governance reporting, AAISM prioritizes risk-and harm-oriented KPIs that reflect safety, reliability, and responsible behavior of AI systems. Error rate (safety/quality signal) and bias detection (fairness/compliance signal) provide leading indicators of model risk, potential user harm, and regulatory exposure--key interests for a CISO.

Valid AAISM Dumps shared by EduDump.com for Helping Passing AAISM Exam!

EduDump.com now offer the **newest AAISM exam dumps**, the EduDump.com AAISM exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AAISM dumps with Test Engine here:

<https://www.edudump.com/exams/ISACA/AAISM/premium/> (**432** Q&As Dumps, **35%OFF**

Special Discount Code: **freecram**)

NEW QUESTION: 17

Which of the following is the MOST important consideration when deciding how to compose an AI red team?

- A. Resource availability
- B. AI use cases
- C. Time-to-market constraints
- D. Compliance requirements

Answer: ([SHOW ANSWER](#))

AAISM materials specify that the composition of an AI red team must be tailored to the organization's AI use cases. The purpose of red-teaming is to simulate realistic adversarial conditions aligned with the actual applications of AI. For example, testing a generative model requires different expertise than testing a fraud detection system. While resource availability, compliance requirements, and time-to-market pressures are practical considerations, they are secondary to aligning team expertise with use case scenarios. The most important factor is therefore the AI use cases themselves.

NEW QUESTION: 18

Which of the following BEST supports the evaluation of AI system outputs?

- A. Manual review
- B. Scenario analysis
- C. Model card assessment
- D. Input controls

Answer: ([SHOW ANSWER](#))

Manual review allows direct assessment of AI outputs for accuracy, relevance, and potential issues in real-world contexts, making it the most effective method for evaluating the quality and reliability of outputs.

NEW QUESTION: 19

A large pharmaceutical company using a new AI solution to develop treatment regimens is concerned about potential hallucinations with the introduction of real-world data. Which of the following is MOST likely to reduce this risk?

- A. Penetration testing
- B. Human-in-the-loop
- C. AI impact analysis
- D. Data asset validation

Answer: ([SHOW ANSWER](#))

AAISM materials identify human-in-the-loop governance as the most effective safeguard against risks such as hallucinations in AI systems used in high-stakes domains like healthcare. By ensuring that human experts validate outputs before they influence patient treatment decisions, organizations preserve accountability, safety, and accuracy. Penetration testing is a cybersecurity measure, not relevant to hallucination risk. AI impact analysis helps evaluate systemic effects but does not directly prevent faulty outputs. Data validation improves input quality but cannot fully prevent generative hallucinations. The key safeguard is human-in-the-loop oversight.

NEW QUESTION: 20

What BEST protects trade secrets related to AI technologies during their life cycle?

- A. Enforcing trademark rights
- B. Restricting access to sensitive data
- C. Patenting AI algorithms and data
- D. Watermarking AI output

Answer: ([SHOW ANSWER](#))

AAISM emphasizes access control and data security as the strongest mechanisms to protect trade secrets, including:

- proprietary algorithms
- training datasets
- model weights
- internal design documentation

NEW QUESTION: 21

Which of the following is the BEST way to ensure role clarity and staff effectiveness when implementing AI-assisted security monitoring tools?

- A. Defer implementation until the security team can be expanded with data scientists.
- B. Update the security program to include cross-functional AI-specific responsibilities.
- C. Transition responsibilities for AI tools to external consultants for improved scalability.
- D. Increase training budgets for business staff to obtain vendor-neutral AI certifications.

Answer: B ([LEAVE A REPLY](#))

AAISM directs that AI adoption in security operations be governed through explicit operating models and RACI-mapped responsibilities spanning security operations, data science/ML, platform engineering, privacy, and compliance. Role clarity comes from updating the security program to codify AI-specific responsibilities (model monitoring, incident handling for AI failures, data governance, change control for models, bias/fairness reviews, supplier assurance) rather than deferring implementation, outsourcing core accountability, or relying on generic certifications. This ensures measurable accountability, reduces hand-off risk, and aligns day-to-day SOC practices with AI control objectives.

NEW QUESTION: 22

An organization recently introduced a generative AI chatbot that can interact with users and answer their queries. Which of the following would BEST mitigate hallucination risk identified by the risk team?

- A. Performing model testing and validation
- B. Training the foundational model on large data sets
- C. Ensuring model developers have been trained in AI risk
- D. Fine-tuning the foundational model

Answer: ([SHOW ANSWER](#))

AAISM highlights fine-tuning foundational models as one of the most effective strategies for reducing hallucination risk. By tailoring the model with domain-specific, curated, and verified datasets, organizations can reduce the frequency of irrelevant or fabricated outputs. Testing and validation help evaluate risks but do not directly minimize hallucinations. Training on larger datasets may improve generalization but does not guarantee accuracy. Developer training in AI risk supports governance but is not a technical control against hallucinations. The best mitigation is fine-tuning to align the chatbot with trusted, context-specific knowledge.

NEW QUESTION: 23

Which of the following AI data management techniques involves creating validation and test data?

- A. Training
- B. Annotating
- C. Splitting
- D. Learning

Answer: ([SHOW ANSWER](#))

Data splitting partitions a labeled dataset into training, validation, and test subsets to enable unbiased model tuning and evaluation. Training (A) consumes the training split; annotating (B) adds labels; learning (D) is a general term for model optimization, not a data management step.

NEW QUESTION: 24

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 ([LEAVE A REPLY](#))

In the AAISMTM 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.

NEW QUESTION: 25

An organization plans to implement a new AI system. Which of the following is the MOST important factor in determining the level of risk monitoring activities required?

- A. The organization's risk appetite
- B. The organization's number of AI system users
- C. The organization's risk tolerance

D. The organization's compensating controls

Answer: ([SHOW ANSWER](#))

AAISM risk management guidance clarifies that the organization's risk tolerance is the most important factor in determining how much monitoring is needed. Risk tolerance specifies the amount of risk the organization is willing to accept and defines the threshold for triggering monitoring or mitigation activities. Risk appetite is broader and strategic, while tolerance sets the operational limits. The number of users may influence scale, and compensating controls may affect resilience, but neither dictates monitoring intensity as directly as risk tolerance.

NEW QUESTION: 26

A healthcare provider plans to use an external vendor's AI system to analyze patient medical images. Management is concerned about protecting patient data and managing the risk of the new technology. Which of the following should be done FIRST before signing the vendor contract?

- A.** Include guaranteed 24/7 system support and technical performance requirements in the contract terms.
- B.** Focus only on the vendors with a good reputation in healthcare and their industry certifications.
- C.** Ask patients to sign consent forms acknowledging that AI will be used in their care.
- D.** Review the vendor's data handling, regulatory compliance, and supply chain dependencies.

Answer: ([SHOW ANSWER](#))

Reviewing the vendor's data handling practices, compliance with healthcare regulations, and supply chain dependencies ensures that risks related to patient data protection and third-party exposure are understood and addressed before entering into a contractual agreement.

NEW QUESTION: 27

To process customer feedback for business decisions, an organization purchases an AI-powered sentiment analysis service implemented as a black box API. Which of the following represents the GREATEST security risk to the business?

- A.** The organization's API key for the service could be stolen and used by an unauthorized party.
- B.** The API endpoint for the service could be unavailable due to a distributed denial of service (DDoS) attack.
- C.** The service may secretly exfiltrate customer feedback data for its own purposes.
- D.** The service could be subjected to data poisoning attacks that causes sentiment analysis to be manipulated.

Answer: ([SHOW ANSWER](#))

Using a black box API means the organization has limited visibility into how data is handled. If the service exfiltrates customer feedback data, it directly compromises confidentiality and trust, posing the most significant security risk to the business.

NEW QUESTION: 28

Which of the following BEST ensures the integrity of data sets used to train AI models?

- A. Collection and retention of only necessary data sets
- B. Tracking and verification of data sets via cryptographic controls
- C. Appropriate storage of data sets according to documented classification processes
- D. Clear documentation of data sources, types used, and processing steps

Answer: ([SHOW ANSWER](#))

AAISM defines cryptographic tracking and verification as the best control for ensuring the integrity of training data. By applying hashing and verification methods, organizations can confirm that datasets remain unaltered and authentic throughout collection, storage, and processing.

Collecting only necessary data, proper storage, or clear documentation all support governance and compliance, but they do not guarantee that the data has not been tampered with. Integrity is specifically ensured by cryptographic verification techniques.

NEW QUESTION: 29

When integrating AI for innovation, which of the following can BEST help an organization manage security risk?

- A. Re-evaluating the risk appetite
- B. Seeking third-party advice
- C. Evaluating compliance requirements
- D. Adopting a phased approach

Answer: ([SHOW ANSWER](#))

AAISM emphasizes that when introducing innovative AI systems, organizations reduce security and compliance risk by following a phased adoption approach. This allows incremental deployment, controlled testing, and gradual scaling while monitoring risks in real time. Re-evaluating risk appetite and evaluating compliance are important governance steps but do not directly mitigate risks during implementation. Seeking third-party advice can add expertise but does not provide the structured control that phased integration offers. The most effective risk management approach for AI innovation is to adopt a phased rollout strategy.

NEW QUESTION: 30

Which of the following controls would BEST help to prevent data poisoning in AI models?

- A. Increasing the size of the training data set
- B. Implementing a strict data validation mechanism
- C. Establishing continuous monitoring
- D. Regularly updating the foundational model

Answer: ([SHOW ANSWER](#))

The most direct preventative control against data poisoning is robust data validation/ingestion gating: provenance checks, schema and constraint validation, anomaly/outlier screening, label consistency tests, and whitelist/blacklist source controls before data reaches training pipelines.

NEW QUESTION: 31

A large language model (LLM) has been manipulated to provide advice that serves an attacker's objectives. Which of the following attack types does this situation represent?

- A. Privilege escalation
- B. Data poisoning
- C. Model inversion
- D. Evasion attack

Answer: ([SHOW ANSWER](#))

AAISM categorizes the manipulation of an LLM at inference time, where crafted inputs cause outputs to serve attacker objectives, as an evasion attack. Evasion attacks exploit weaknesses in the model's decision-making boundaries by altering queries to produce compromised or misleading outputs. Privilege escalation refers to unauthorized access rights, data poisoning targets the training phase, and model inversion reconstructs training data. In this case, manipulation of outputs to align with an attacker's goals reflects an evasion attack.

Valid AAISM Dumps shared by EduDump.com for Helping Passing AAISM Exam!
EduDump.com now offer the **newest AAISM exam dumps**, the EduDump.com AAISM exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AAISM dumps with Test Engine here:
<https://www.edudump.com/exams/ISACA/AAISM/premium/> (**432** Q&As Dumps, **35%OFF**
Special Discount Code: freecram)

NEW QUESTION: 32

Secure aggregation enhances federated learning security by:

- A. Encrypting individual model updates so only the server can access them
- B. Applying differential privacy to training data
- C. Ensuring client contributions remain confidential even if the server is compromised
- D. Processing client updates in isolation

Answer: ([SHOW ANSWER](#))

AAISM explains that secure aggregation ensures the server only sees aggregated model updates--not individual client contributions--so privacy is preserved even if the server is breached.

NEW QUESTION: 33

Which of the following AI controls would BEST protect against deepfake content?

- A. AI prompt validation
- B. Data provenance
- C. Data ownership
- D. AI model testing

Answer: ([SHOW ANSWER](#))

Data provenance establishes the origin and authenticity of content, enabling verification of whether media has been altered or fabricated, which is essential for detecting and protecting against deepfakes.

NEW QUESTION: 34

Which of the following AI incidents would be BEST contained via a kill switch?

- A. The agent attempts self-replication.
- B. The query takes too long to complete.
- C. A model version update fails.
- D. Graphics processing unit utilization decreases.

Answer: A ([LEAVE A REPLY](#))

A kill switch is designed to immediately stop or disable an AI system when it exhibits dangerous, uncontrolled, or unauthorized behavior. An agent attempting self-replication represents a critical safety and security threat that requires rapid containment through a kill switch mechanism.

NEW QUESTION: 35

Which of the following is the BEST approach to establish continuous monitoring of AI systems after rolling out AI services in a global organization?

- A. Assess independent control reports and implement vendor dashboard monitoring.
- B. Emphasize pre-deployment testing and model validation reports, deferring runtime monitoring until the services start producing results.
- C. Implement predefined metrics, automate logging, and escalate alerts to the security operations center (SOC).
- D. Quickly schedule manual reviews of model accuracy on a monthly basis.

Answer: C ([LEAVE A REPLY](#))

Defining monitoring metrics, automating logging, and integrating alerting with the SOC enables real-time visibility and response to issues. This provides continuous oversight of AI system behavior and risks across the organization after deployment.

NEW QUESTION: 36

Which of the following is the MOST effective and secure way to prevent model attribute attacks against an organization's centralized AI system training data that includes sensitive information about individuals?

- A. Monitoring the output of AI learning models continuously
- B. Using federated learning to train AI algorithms
- C. Implementing a cross-validation AI learning model
- D. Removing private data from AI training datasets

Answer: ([SHOW ANSWER](#)**)**

Federated learning keeps sensitive data decentralized by training models locally and sharing only model updates rather than raw data. This significantly reduces exposure of individual-level information and helps prevent attribute inference attacks against centralized training datasets.

NEW QUESTION: 37

An organization is designing an AI-based credit risk assessment system integrating sensitive financial data. Which option BEST supports security-by-design?

- A. Integrating differential privacy mechanisms into model training
- B. Applying threat modeling specific to AI components before deployment
- C. Segmenting AI services across containers
- D. Restricting access to AI models using IP allow lists

Answer: ([SHOW ANSWER](#))

AAISM identifies AI-specific threat modeling as an essential early-stage control in security-by-design, particularly for high-risk systems like credit scoring. It systematically identifies:

- data poisoning
- bias vulnerabilities
- model evasion
- model extraction
- misuse scenarios

NEW QUESTION: 38

Which of the following is an information security manager's BEST recommendation when the new release of a proprietary AI model is generating inconsistent output?

- A. Report the issue to the AI governance committee.
- B. Check the model scorecard to confirm data sources.
- C. Retrain the model with more datasets.
- D. Revert to a prior version of the model.

Answer: ([SHOW ANSWER](#))

Reverting to a prior stable version is the best immediate recommendation because inconsistent output from a new model release indicates potential instability or unintended behavior. Restoring the previously validated version minimizes operational and security risk while the issue is investigated and corrected.

NEW QUESTION: 39

Which of the following is the BEST reason to involve a diverse team when evaluating the data used within an AI system?

- A. It reduces the likelihood of a risk being overlooked.
- B. It ensures that scenarios are reviewed thoroughly.
- C. It reduces the perception of bias by the media and users.
- D. It ensures the output is representative of target end users.

Answer: ([SHOW ANSWER](#))

Involving a diverse team provides broader perspectives, experiences, and expertise during evaluation of AI data. This helps identify hidden biases, ethical concerns, security issues, and other risks that might otherwise be overlooked by a less diverse review group.

NEW QUESTION: 40

Which of the following BEST describes the role of transparency in AI?

- A. Talking through a decision tree to better understand how the algorithm made each of its choices
- B. Publishing AI mechanisms, data sources, and decision-making processes while making them openly available
- C. Explaining the AI system in an understandable and logical way so reasons for decisions can be given
- D. Persuading someone that the AI tool in use is beneficial and operates as expected

Answer: ([SHOW ANSWER](#))

Transparency in AI is a governance principle requiring that systems be explainable to stakeholders in ways that are understandable and meaningful, enabling clear articulation of how decisions were reached and why. Within an AI program, transparency supports accountability, auditability, and trust by ensuring that reasons for decisions can be communicated and scrutinized. Option C reflects this definition by focusing on intelligible, logical explanations of system behavior and decision rationale.

NEW QUESTION: 41

An organization implementing an LLM application sees unexpected cost increases due to excessive computational resource usage. Which vulnerability is MOST likely in need of mitigation?

- A. Excessive agency
- B. Sensitive information disclosure
- C. Unbounded consumption
- D. System prompt leakage

Answer: ([SHOW ANSWER](#))

AAISM categorizes unbounded consumption (also known as "resource exhaustion" or "infinite queries") as an AI-specific vulnerability where attackers (or faulty prompts) trigger excessive computation, leading to high costs and degraded service.

NEW QUESTION: 42

Which of the following is MOST important to consider when obtaining senior management and stakeholder support for a new AI governance program?

- A. Historical data related to security breaches and remediation efforts
- B. Alignment with existing security and privacy policies
- C. Risk visualizations that incorporate AI-specific threats and vulnerabilities
- D. Impact on diverse enterprise objectives and processes

Answer: ([SHOW ANSWER](#))

Senior management and stakeholders are most influenced by how an AI governance program affects organizational objectives, operations, risk management, compliance, and business

processes. Demonstrating enterprise-wide impact helps secure strategic alignment, sponsorship, and long-term support for the governance initiative.

NEW QUESTION: 43

Which approach should an organization prioritize to effectively verify the security of its AI models?

- A.** Automating vulnerability identification
- B.** Developing a testing strategy including AI-specific threat modeling and adversarial attack simulations
- C.** Testing team competencies in IT threat mitigation
- D.** Using standard penetration testing methods

Answer: ([SHOW ANSWER](#))

The AAISM standard explicitly states that traditional penetration tests alone are insufficient for AI systems. Effective AI security testing requires:

- AI-specific threat modeling (e.g., data poisoning, prompt injection, model theft)
- Adversarial attack simulations (white-box, black-box, gradient-based attacks)
- Evaluation of robustness and manipulation resistance

NEW QUESTION: 44

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: ([SHOW ANSWER](#))

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.

NEW QUESTION: 45

A programmer suspects an AI system is inferring sensitive user information. What is the BEST action?

- A.** Inform the governance panel
- B.** Suggest fine-tuning
- C.** Conduct a code review
- D.** Alert the CIO

Answer: (SHOW ANSWER)

AAISM directs that potential privacy, ethical, or compliance risks must be escalated to the AI Governance Panel, the body responsible for oversight, risk approval, and corrective action.

NEW QUESTION: 46

An automotive manufacturer uses AI-enabled sensors on machinery to monitor variables such as vibration, temperature, and pressure. Which of the following BEST demonstrates how this approach contributes to operational resilience?

- A. Scheduling repairs for critical equipment based on real-time condition monitoring
- B. Performing regular maintenance based on manufacturer recommendations
- C. Conducting monthly manual reviews of maintenance schedules
- D. Automating equipment repairs without any human intervention

Answer: A (LEAVE A REPLY)

AAISM highlights that AI-enabled predictive maintenance improves operational resilience by using real-time sensor monitoring to schedule repairs based on actual conditions rather than fixed schedules. This prevents unexpected breakdowns, reduces downtime, and ensures continuity of operations. Regular maintenance based on recommendations is static and may not reflect real conditions. Manual reviews are slow and inefficient. Full automation of repairs without human oversight is not realistic or safe in critical manufacturing. The approach that best demonstrates resilience is real-time condition-based repair scheduling.

Valid AAISM Dumps shared by EduDump.com for Helping Passing AAISM Exam!

EduDump.com now offer the **newest AAISM exam dumps**, the EduDump.com AAISM exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AAISM dumps with Test Engine here:

<https://www.edudump.com/exams/ISACA/AAISM/premium/> (**432** Q&As Dumps, **35%OFF**

Special Discount Code: freeexam)

NEW QUESTION: 47

Which of the following presents the GREATEST challenge to gaining an understanding of how a deep learning system operates?

- A. Ineffectiveness of traditional data flow mapping approaches
- B. Insufficient output volumes and diversity to draw definitive conclusions
- C. Complex non-linear relationships among features and parameters
- D. Steep learning curve for interpretation of system testing results

Answer: C (LEAVE A REPLY)

Deep learning systems rely on highly complex non-linear relationships among large numbers of features, hidden layers, and parameters. These interactions make it difficult to understand, interpret, and explain how the model reaches specific decisions or outputs.

NEW QUESTION: 48

A model producing contradictory outputs based on highly similar inputs MOST likely indicates the presence of:

- A. Poisoning attacks
- B. Evasion attacks
- C. Membership inference
- D. Model exfiltration

Answer: ([SHOW ANSWER](#))

The AAISM study framework describes evasion attacks as attempts to manipulate or probe a trained model during inference by using crafted inputs that appear normal but cause the system to generate inconsistent or erroneous outputs. Contradictory results from nearly identical queries are a typical symptom of evasion, as the attacker is probing decision boundaries to find weaknesses. Poisoning attacks occur during training, not inference, while membership inference relates to exposing whether data was part of the training set, and model exfiltration involves extracting proprietary parameters or architecture. The clearest indication of contradictory outputs from similar queries therefore aligns directly with the definition of evasion attacks in AAISM materials.

NEW QUESTION: 49

During training, the PRIMARY security benefit of using synthetic data derived from real records is:

- A. reduced inference latency.
- B. faster convergence.
- C. lowered re-identification risk.
- D. higher processing efficiencies.

Answer: ([SHOW ANSWER](#))

Synthetic data derived from real records helps reduce the risk that individuals can be directly identified from training datasets. This lowers re-identification risk while still enabling AI models to learn useful patterns from representative data.

NEW QUESTION: 50

Which of the following is the PRIMARY objective of performing a privacy impact assessment (PIA) for an AI system?

- A. Validating data provenance for all training datasets
- B. Documenting AI model training methodologies and parameter settings
- C. Identifying unexpected processing of personal information
- D. Adapting AI acceptable use policies to jurisdiction-specific privacy regulations

Answer: ([SHOW ANSWER](#))

The primary objective of a privacy impact assessment is to identify and evaluate risks related to the processing of personal information, including unexpected, excessive, or inappropriate data

use. This helps ensure that privacy risks are recognized and mitigated before or during AI system deployment.

NEW QUESTION: 51

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: ([SHOW ANSWER](#))

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.

NEW QUESTION: 52

Which of the following would be of GREATEST concern to an information security manager when a small, startup organization is proposing the use of a pre-trained open-source AI solution?

- A. The organization has limited visibility to how data is being used.
- B. There may be a lack of support if the project is abandoned.
- C. The model is not compatible with enterprise architecture (EA).
- D. The open-source model includes a hidden backdoor.

Answer: D ([LEAVE A REPLY](#))

A hidden backdoor represents a direct and severe security risk, potentially enabling unauthorized access, data exfiltration, or system compromise. This type of threat can undermine the entire security posture of the organization and is therefore the most critical concern compared to operational or compatibility issues.

NEW QUESTION: 53

An organization is implementing an AI-based credit assessment engine using internal and third-party customer data. Which of the following BEST aligns with data management controls for the AI life cycle?

- A. Documented procedures for data sourcing, lineage tracking, and quality validation
- B. Use of hashed identifiers to anonymize datasets used for model validation and internal analytics
- C. Encrypted isolation and dynamic access controls on training data pipelines
- D. Limitation of model training to structured data from vetted sources to minimize ingestion risk

Answer: ([SHOW ANSWER](#))

AAISM emphasizes that data governance over the full AI life cycle is foundational. The official content describes effective AI data management as including documented procedures for: (1) how data is sourced, (2) how lineage is tracked from origin to model, and (3) how data quality is validated and monitored. This ensures transparency, accountability, and auditability, which are especially critical in regulated areas like credit assessments.

NEW QUESTION: 54

Which of the following should be the PRIMARY objective of implementing differential privacy techniques in AI models leveraging fraud detection systems?

- A. Enhancing the accuracy of predictions to desired levels
- B. Increasing model training speed for an efficient launch
- C. Protecting individual data contributions while allowing statistical analysis
- D. Reducing computational resources required for the model training phase

Answer: ([SHOW ANSWER](#))

Differential privacy aims to protect the privacy of any single individual's data contribution while still enabling useful aggregate learning and statistical analysis. Noise mechanisms are calibrated so that results remain informative for modeling (e.g., fraud patterns) without revealing whether any particular person's data was included or enabling inference about them. Accuracy, speed, and compute efficiency can be secondary considerations, but the primary objective is privacy protection with utility preserved.

NEW QUESTION: 55

The CEO of a large multi-sector organization has learned that, in the absence of formal guidance, business units have created shadow IT by integrating AI into their daily operations. In response, an AI governance group has been created to address the issue. Which of the following is MOST important for the group to collect from each business unit?

- A. Inventory of models in use
- B. Business continuity plan (BCP)
- C. API key management practices
- D. Software bill of materials

Answer: ([SHOW ANSWER](#))

Establishing a comprehensive inventory of AI models in use provides immediate visibility into shadow IT, enabling the governance group to understand scope, assess risks, and prioritize controls. This foundational step is necessary before addressing operational practices or implementing governance measures.

NEW QUESTION: 56

A large financial services organization is integrating a third-party AI solution into its critical fraud detection system. Which of the following is the BEST way for the organization to reduce risk associated with AI vendor and supply chain dependencies?

- A. Conducting annual vulnerability assessments of the fraud detection system after integration

- B. Focusing on performance testing to ensure the solution meets operational requirements
- C. Establishing contractual agreements requiring vendors to provide evidence of secure development practices
- D. Implementing isolated virtual environments to validate the integration of the fraud detection system with the solution

Answer: ([SHOW ANSWER](#))

AAISM emphasizes supplier assurance through contractual obligations as the foundational control for AI supply chain risk. Contracts should require verifiable evidence of secure development practices (e.g., secure SDLC, model and data provenance documentation, SBOM/MBOM where applicable, vulnerability disclosure, patch SLAs, audit rights, incident notification, and regulatory compliance assertions). This creates enforceable, continuous assurance beyond point-in-time tests.

NEW QUESTION: 57

Which of the following is the GREATEST benefit of performing AI security risk assessments?

- A. Appropriate privacy risk controls are implemented for AI models
- B. The appropriate level of funding is secured for AI security risk
- C. The risk register is updated with the latest AI risk
- D. Risk prioritization decisions are made for AI security

Answer: ([SHOW ANSWER](#))

AAISM emphasizes that the core outcome of AI risk assessments is prioritization: mapping threat likelihood and business impact to determine which risks to treat first, at what strength, and with which controls. Implementing privacy controls (A), funding alignment (B), and updating registers (C) are important outputs, but the greatest benefit is making defensible, prioritized decisions that align with risk appetite and optimize control selection and resource allocation.

NEW QUESTION: 58

An organization's HR chatbot was discovered generating fake employee details using a third-party API. Which would be the information security team's BEST recommendation?

- A. Retrain the chatbot to validate input.
- B. Limit the chatbot's communication with the third-party API.
- C. Review the chatbot's test reports.
- D. Request the third-party API logs for review.

Answer: ([SHOW ANSWER](#))

The issue stems from the chatbot relying on an external API that is producing unreliable outputs. Limiting or restricting communication with the third-party API reduces exposure to untrusted responses and prevents further generation of incorrect or fabricated employee information, directly mitigating the risk.

NEW QUESTION: 59

Which of the following would BEST validate automated business processes enabled by the introduction of AI technology?

- A. Complying with regulations to ensure public trust
- B. Ensuring human-in-the-loop oversight
- C. Aligning with core principles for disclosure of information
- D. Conducting periodic reviews of business process integrity

Answer: ([SHOW ANSWER](#))

Human-in-the-loop oversight is the best way to validate automated business processes enabled by AI because it ensures that human judgment, review, and intervention remain part of critical decision-making processes. This helps verify accuracy, appropriateness, and ethical operation of AI-driven activities.

NEW QUESTION: 60

A hospital adopts an AI solution from an external vendor to help diagnose rare diseases. Which of the following BEST demonstrates the verification of security requirements for this technology?

- A. Require contractual clauses stipulating periodic information security reviews.
- B. Preview contractual clauses of mandatory security AI software update.
- C. Perform regular in-depth vulnerability scans on the AI system.
- D. Integrate requirements for sensitive data aligned with healthcare regulations.

Answer: ([SHOW ANSWER](#))

Verification of security requirements for an AI used in diagnosis must focus on how sensitive health data (PHI) is protected in line with regulations such as HIPAA and similar healthcare laws. Integrating explicit requirements for handling, storing, transmitting, and accessing sensitive data in accordance with healthcare regulations is a direct verification that the technology's security controls meet the hospital's legal and security obligations.

NEW QUESTION: 61

Which of the following BEST describes how supervised learning models help reduce false positives in cybersecurity threat detection?

- A. They analyze patterns in data to group legitimate activity from actual threats
- B. They use real-time feature engineering to automatically adjust decision boundaries
- C. They learn from historical labeled data
- D. They dynamically generate new labeled data sets

Answer: ([SHOW ANSWER](#))

According to AAISM technical content, supervised learning models reduce false positives by learning from historical labeled data that distinguishes between legitimate activity and actual threats. This training enables the model to recognize patterns and improve its discrimination ability over time. Grouping patterns (A) describes clustering, an unsupervised method. Real-time feature engineering (B) and generating new labeled data (D) are advanced techniques but not the fundamental supervised learning approach. The essence of supervised learning is leveraging labeled data to minimize misclassification, including false positives.

Valid AAISM Dumps shared by EduDump.com for Helping Passing AAISM Exam!
EduDump.com now offer the **newest AAISM exam dumps**, the EduDump.com AAISM exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AAISM dumps with Test Engine here:
<https://www.edudump.com/exams/ISACA/AAISM/premium/> (**432 Q&As Dumps, 35%OFF**
Special Discount Code: freecram)

NEW QUESTION: 62

A health organization wants to use an AI tool for diagnostic decisions. Which of the following would BEST help the organization comply with privacy regulations?

- A. De-identifying training data
- B. Sanitizing training data
- C. Conducting a privacy impact assessment (PIA)
- D. Pseudonymizing AI model output data

Answer: (SHOW ANSWER)

Conducting a privacy impact assessment (PIA) is the best way to help ensure compliance with privacy regulations because it systematically identifies, evaluates, and mitigates privacy risks associated with the collection, processing, and use of sensitive health data within the AI diagnostic system.

NEW QUESTION: 63

Which of the following is the GREATEST risk for an organization that has adopted AI-powered chatbots for its customer service platform?

- A. Model distillation
- B. Prompt injection
- C. Model inversion
- D. Data poisoning

Answer: (SHOW ANSWER)

Prompt injection directly targets chatbot interactions by manipulating inputs to override instructions or extract sensitive information, making it a primary and immediate risk in customer-facing AI systems.

NEW QUESTION: 64

An organization is planning to commission a third-party AI system to make decisions using sensitive data. Which of the following metrics is MOST important for the organization to consider?

- A. Model response time
- B. Service availability
- C. Accessibility rating

D. Accuracy thresholds

Answer: ([SHOW ANSWER](#))

When AI systems make consequential decisions over sensitive data, AAISM requires explicit performance thresholds tied to decision quality--i.e., accuracy (and related error/false-rate limits) aligned to business risk appetite and regulatory expectations. Availability and latency are important service metrics, but decision integrity and error bounds are primary risk drivers in sensitive contexts. Establishing, monitoring, and enforcing minimum accuracy thresholds (with subgroup performance checks) is essential to reduce harm, ensure fairness/compliance, and support auditability.

NEW QUESTION: 65

Which of the following is the MOST effective way to reduce the risk of reverse engineering of an AI model?

- A.** Watermarking the model
- B.** Utilizing proprietary APIs
- C.** Encapsulating parameters
- D.** Limit the number of queries

Answer: ([SHOW ANSWER](#))

Limiting the number of queries reduces the amount of information an attacker can collect from the model's inputs and outputs. Reverse engineering attacks often rely on extensive querying to approximate the model's behavior, architecture, or parameters, so restricting queries is the most effective mitigation.

NEW QUESTION: 66

Which of the following involves documenting and monitoring the complete journey of data as it flows through an AI system?

- A.** Lineage
- B.** Transformation
- C.** Origin
- D.** Processing

Answer: ([SHOW ANSWER](#))

Data lineage records and monitors the end-to-end journey of data--sources, movements, transformations, storage locations, uses, and dependencies--providing traceability, auditability, and accountability across the AI lifecycle. "Origin" is a single point (provenance), "transformation" is one step within the flow, and "processing" is a general activity rather than a governance record of the entire path.

NEW QUESTION: 67

An organization is adopting multiple AI solutions for customer support and management is concerned about potential privacy and confidentiality violations due to the implementation of the new solution. Which of the following is the BEST risk mitigation strategy in this situation?

- A. Rely on best practice manuals from AI vendors to guide how employees use the systems.
- B. Implement key risk indicators (KRIs) to monitor AI usage and update the risk register.
- C. Create responsible AI use policies that address ethics, acceptable use, and transparency.
- D. Use consumer-level AI guidelines with rules, terms, and recommendations for end users.

Answer: ([SHOW ANSWER](#))

Establishing responsible AI use policies defines clear expectations for handling data, ethical use, and transparency, directly addressing privacy and confidentiality risks across multiple AI solutions and guiding consistent behavior organization-wide.

NEW QUESTION: 68

Within which stage of the AI development life cycle should effective feature engineering be conducted?

- A. Testing
- B. Define
- C. Design
- D. Development

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 69

Embedding unique identifiers into AI models would BEST help with:

- A. Preventing unauthorized access
- B. Tracking ownership
- C. Eliminating AI system biases
- D. Detecting adversarial attacks

Answer: ([SHOW ANSWER](#))

The AAISM framework explains that embedding unique identifiers--such as digital watermarks or model fingerprints--enables organizations to trace and verify model provenance. This technique is used for tracking ownership and intellectual property rights over models, particularly when sharing, licensing, or distributing AI systems. While identifiers may support certain security functions, their primary control objective is ownership verification, not preventing access, bias removal, or adversarial detection. The correct alignment with AAISM controls is tracking ownership.

NEW QUESTION: 70

Which of the following BEST describes the role of model cards in AI solutions?

- A. They are primarily used to visualize the performance of AI models
- B. They are used to automatically fine-tune AI models by adjusting hyperparameters based on user feedback
- C. They provide a standardized way to document the training data and AI model use cases
- D. They help developers create synthetic data and train AI models

Answer: ([SHOW ANSWER](#))

AAISM positions model cards as standardized documentation artifacts that record intended use and out-of-scope use, training/evaluation data characteristics, performance metrics across groups, limitations/risks, and governance controls/owners. Their purpose is transparency and assurance, not automated tuning or synthetic data generation. Visualization (A) may appear within a card, but the core role is structured documentation for governance, risk, and compliance.

NEW QUESTION: 71

A PRIMARY objective of responsibly providing AI services is to:

- A. Enable AI models to operate autonomously
- B. Ensure the confidentiality and integrity of data processed by AI models
- C. Build trust for decisions and predictions made by AI models
- D. Improve the ability of AI models to learn from new data

Answer: ([SHOW ANSWER](#))

AAISM emphasizes that the primary objective of responsible AI is to establish and maintain trust in AI-driven decisions and predictions. Trust is achieved through transparency, accountability, fairness, and governance. While confidentiality and integrity are critical technical objectives, they are not the overarching purpose of responsible AI service provision. Autonomy and learning ability are features of AI, but without trust, adoption and compliance falter. The correct answer is that responsible AI services must focus on building trust in AI outcomes.

NEW QUESTION: 72

In a new supply chain management system, AI models used by participating parties are interactively connected to generate advice in support of management decision making. Which of the following is the GREATEST challenge related to this architecture?

- A. Establishing clear lines of responsibility for AI model outputs
- B. Identifying hallucinations returned by AI models
- C. Determining the aggregate risk of the system
- D. Explaining the overall benefit of the system to stakeholders

Answer: ([SHOW ANSWER](#))

When multiple AI models interact in a connected supply chain system, risks can propagate and compound across models. Assessing the overall, aggregate risk - including errors, biases, or failures from individual models - is challenging but critical for ensuring reliable and safe decision-making.

NEW QUESTION: 73

An organization uses an AI model to optimize traffic signals and predict congestion. Due to misconfiguration and lack of validation, the system over-prioritizes low-traffic areas, causing gridlock in major cities. Which of the following should be done FIRST to address this issue?

- A. Retrain the model in real time using new traffic data.
- B. Roll back the AI system to a previous version.
- C. Continue running the AI model and deploy a secondary system.

D. Disable the AI system and issue a public statement.

Answer: B ([LEAVE A REPLY](#))

Rolling back to a previously stable version immediately restores known correct behavior and mitigates ongoing impact. This provides a controlled state before further validation or retraining is performed.

NEW QUESTION: 74

Which BEST describes the role of model cards in AI solutions?

- A.** They visualize AI model performance
- B.** They document training data and AI model use cases
- C.** They help developers create synthetic data
- D.** They automatically fine-tune AI models

Answer: B ([LEAVE A REPLY](#))

AAISM explains that model cards provide structured documentation about AI models, including:

- intended use cases
- training data characteristics
- ethical considerations
- known limitations
- risk factors
- performance benchmarks

NEW QUESTION: 75

Which of the following BEST helps an organization protect confidential data while using a pre-trained open-source AI solution?

- A.** Monitor the solution's performance for reliability.
- B.** Verify and sanitize both inputs and outputs.
- C.** Periodically test the solution for vulnerabilities.
- D.** Continuously monitor usage to detect fraud and abuse.

Answer: ([SHOW ANSWER](#)**)**

Verifying and sanitizing inputs prevents sensitive data from being unintentionally exposed to the model, while sanitizing outputs helps ensure that confidential information is not leaked in responses, directly protecting data during use.

NEW QUESTION: 76

An organization wants to determine whether their recent AI implementation and organizational objectives are in alignment. Which tool would BEST support this?

- A.** AI incident playbook
- B.** Business impact analysis (BIA)
- C.** AI risk assessment
- D.** Key performance indicators (KPIs)

Answer: ([SHOW ANSWER](#)**)**

KPIs provide measurable metrics that directly reflect how well the AI implementation is achieving the organization's objectives, making them the most effective tool for assessing alignment between the system and business goals.

Valid AAISM Dumps shared by EduDump.com for Helping Passing AAISM Exam!
EduDump.com now offer the **newest AAISM exam dumps**, the EduDump.com AAISM exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AAISM dumps with Test Engine here:

<https://www.edudump.com/exams/ISACA/AAISM/premium/> (**432** Q&As Dumps, **35%OFF**

Special Discount Code: **freecram**)

NEW QUESTION: 77

An information security manager has implemented key performance indicators (KPIs) and key risk indicators (KRIs) to oversee the security posture of an AI-based fraud detection system.

Which of the following would be of GREATEST concern?

- A. Indicators also assess network traffic unrelated to AI functions.
- B. Indicators do not include comparisons against industry competitors.
- C. Indicators fail to capture adversarial techniques targeting AI components.
- D. Indicators primarily focus on quantitative fraud and threat data.

Answer: (SHOW ANSWER)

If KPIs and KRIs do not account for adversarial techniques targeting AI components, critical AI-specific threats such as evasion, poisoning, or model manipulation may go undetected, leaving significant gaps in the security posture.

NEW QUESTION: 78

An organization lacks visibility into the decision-making processes of a newly deployed AI model, and there are concerns regarding accuracy, bias, and stakeholder trust. Which of the following is the BEST way to improve model accuracy?

- A. Implement a governance system that defines procedures for model testing and validation.
- B. Establish monitoring controls to track AI decision-making and escalate anomalies for further analysis.
- C. Increase the size of training datasets to encompass as many various scenarios as possible.
- D. Require explainability features and review of AI outputs as part of model governance.

Answer: (SHOW ANSWER)

Expanding the training dataset with diverse and representative scenarios improves the model's ability to generalize and reduces errors, directly enhancing accuracy.

NEW QUESTION: 79

Which of the following datasets is used to tune hyperparameters?

- A. Validation
- B. Test
- C. Configuration
- D. Training

Answer: ([SHOW ANSWER](#))

Per AAISM's ML lifecycle controls, hyperparameter tuning is performed on the validation set, reserving the test set strictly for final, unbiased performance estimation. The training set is used to fit parameters; the validation set guides model selection and hyperparameter optimization; the test set is untouched until the end to prevent leakage and optimistic bias. "Configuration" is not a dataset type in the lifecycle split.

NEW QUESTION: 80

Which of the following is the MOST important consideration for an organization that has decided to adopt AI to leverage its competitive advantage?

- A. Develop a comprehensive strategic roadmap for AI integration
- B. Develop a comprehensive risk management process to address AI-related issues
- C. Develop internal training programs on AI governance, risk, and compliance (GRC)
- D. Develop a business case for the procurement of AI monitoring tools

Answer: ([SHOW ANSWER](#))

AAISM's governance guidance emphasizes that adopting AI for competitive advantage must begin with a comprehensive strategic roadmap for integration. This roadmap aligns AI adoption with business objectives, sets priorities, defines milestones, and ensures coordination across functions. Risk management, training, and tool procurement are essential, but they are tactical steps that follow once the strategic direction is defined. Without a roadmap, adoption becomes fragmented and risks misalignment with business strategy. The most important consideration at the adoption stage is therefore creating a strategic integration roadmap.

NEW QUESTION: 81

Which of the following assessments MOST effectively addresses concerns related to ethical and trustworthy AI use?

- A. Fundamental rights impact assessment
- B. Information security risk assessment
- C. Business impact analysis (BIA)
- D. Privacy impact assessment (PIA)

Answer: ([SHOW ANSWER](#))

A fundamental rights impact assessment evaluates how AI systems may affect human rights, fairness, discrimination, autonomy, and ethical principles. This makes it the most effective assessment for addressing concerns related to ethical and trustworthy AI use.

NEW QUESTION: 82

Which of the following AI system vulnerabilities is MOST easily exploited by adversaries?

- A. Inaccurate generalizations from new data by the AI model
- B. Weak controls for access to the AI model
- C. Lack of protection against denial of service (DoS) attacks
- D. Inability to detect input modifications causing inappropriate AI outputs

Answer: ([SHOW ANSWER](#))

AAISM study materials stress that weak access controls are the most easily exploited vulnerability in AI systems. Without strong access restrictions, adversaries can directly query, extract, manipulate, or overload models, leading to data leakage or compromised outputs. While inaccurate generalizations, DoS vulnerabilities, or susceptibility to input manipulation are serious, they typically require more effort or specific conditions. Weak access control provides the most direct and immediate entry point for attackers. As such, it is identified as the most easily exploited vulnerability.

NEW QUESTION: 83

An internet service provider (ISP) uses a chatbot based on a public model to help resolve connectivity issues. Customers report that the chatbot's instructions lack details about the ISP's proprietary equipment. Which of the following is the BEST way to improve the user experience?

- A. Retrieval-augmented generation (RAG)
- B. Input filtering
- C. Output sanitization
- D. Use of additional public models

Answer: ([SHOW ANSWER](#))

Retrieval-augmented generation integrates the chatbot with the organization's proprietary knowledge sources, enabling it to provide accurate and detailed responses specific to the ISP's equipment while maintaining the strengths of the base model.

NEW QUESTION: 84

An organization is implementing AI agent development across multiple engineering teams. Which of the following is the MOST important focus of AI-specific security training for developers?

- A. Prompt injection, agent memory control, and insecure tool execution
- B. Dataset bias, explainability, and fairness in model decisions
- C. Output moderation, hallucination handling, and policy alignment
- D. API abuse, data leakage, and third-party plug-in risk

Answer: ([SHOW ANSWER](#))

For developer-facing, near-term hardening of AI agents, AAISM prioritizes secure agent design and runtime controls against prompt injection, unsafe memory/tool use, and tool-execution compromise. These are primary exploitation paths for agents that read external content, persist memory, and call tools with elevated privileges. Training must center on: guarding tool invocation, constraining memory scope, sanitizing/validating inputs, and isolating high-risk actions.

NEW QUESTION: 85

A hospital is conducting an AI impact assessment for a proposed diagnostic imaging tool. To ensure conformity with regulatory and ethical standards, which risk requires the MOST domain-specific assessment beyond what is typically found in a standard IT risk assessment?

- A.** System downtime may increase due to hardware failure in the hosting environment, preventing medical professionals from accessing the tool.
- B.** Software vulnerabilities in the application's web server may allow an attacker to gain unauthorized access to the system.
- C.** The model may produce systematically biased diagnoses for specific demographics, potentially leading to widespread clinical harm.
- D.** Noncompliance with data residency laws may occur if patient data is processed in a cloud region outside the required jurisdiction.

Answer: ([SHOW ANSWER](#))

Bias in diagnostic outputs directly affects clinical decision-making and patient outcomes, requiring specialized medical, ethical, and regulatory expertise beyond standard IT risk assessments. This type of risk involves domain-specific evaluation of fairness, safety, and potential harm to patients.

NEW QUESTION: 86

The PRIMARY reason to conduct a privacy impact assessment (PIA) on an AI system is to:

- A.** Identify applicable regulations
- B.** Determine whether personal data is poisoned
- C.** Build customer confidence
- D.** Analyze how personal data is handled

Answer: ([SHOW ANSWER](#))

According to AAISM privacy governance guidance, the primary reason for conducting a PIA is to analyze how personal data is collected, processed, shared, and retained by an AI system. This analysis ensures compliance with privacy laws, mitigates risks to individuals, and informs necessary safeguards. Identifying regulations is part of compliance but is secondary to analyzing actual data handling. Building customer confidence is an outcome, not the main purpose. Checking for poisoned data relates to data quality, not privacy assessment. The fundamental purpose of a PIA is therefore to analyze the handling of personal data.

NEW QUESTION: 87

Which of the following is the PRIMARY benefit of incorporating an AI component into an organization's security program?

- A.** Providing focus on vulnerabilities and attacks specific to AI while applying cybersecurity best practices
- B.** Maintaining clear separation between AI security and existing information security governance initiatives
- C.** Specifying tailored AI training and learning pathways for security team members
- D.** Ensuring consistent AI integration across all organizational business units

Answer: ([SHOW ANSWER](#))

Incorporating an AI component ensures that security efforts address AI-specific threats and vulnerabilities while aligning with established cybersecurity practices, strengthening the overall security posture.

NEW QUESTION: 88

When evaluating a third-party AI service provider, which master services agreement (MSA) provision is MOST critical for managing security risk?

- A. Guaranteeing unlimited model retraining requests
- B. Sharing real-time log information
- C. Prohibiting the use of customer data for model training
- D. Restricting query volume thresholds

Answer: (SHOW ANSWER)

AAISM emphasizes strong contractual restrictions on how vendors use customer data, especially prohibiting vendors from using customer inputs to train or fine-tune shared models.

This protects against:

- data leakage
- intellectual property exposure
- regulatory violations
- shadow training of external models

NEW QUESTION: 89

During the deployment of a generative AI platform, a risk assessment highlighted threats such as data leakage and prompt manipulation. Which of the following is the BEST way to ensure appropriate control selection?

- A. Rely primarily on vendor-provided security features and seek third-party certifications
- B. Map identified AI threats to enterprise control catalogs and integrate AI-specific safeguards where gaps exist
- C. Apply AI-specific controls from external frameworks without customization and initiate monitoring to expedite compliance
- D. Postpone control selection until deployment and address risk through enhanced monitoring

Answer: (SHOW ANSWER)

AAISM requires that control selection be threat-led and context-specific, aligning AI threats to the organization's existing enterprise control catalogs (security, privacy, resilience) and augmenting them with AI-specific safeguards where coverage is insufficient. This ensures consistency with the risk appetite, removes duplication, and closes AI-unique gaps (e.g., prompt injection, data leakage from context windows, model misuse). Generic reliance on vendors or uncustomized external frameworks does not ensure fit-for-purpose coverage, and deferring control selection to post-deployment contradicts proactive risk treatment.

NEW QUESTION: 90

Which of the following is the BEST technique for addressing the likelihood of a model evasion attack?

- A. Integrating multiple detection mechanisms
- B. Implementing hashing techniques
- C. Developing a catalog of AI-specific attack scenarios
- D. Establishing continuous monitoring

Answer: ([SHOW ANSWER](#))

Integrating multiple detection mechanisms creates layered defenses, making it significantly harder for adversarial inputs to bypass all controls. This diversity increases resilience against evasion attempts by detecting anomalies through different methods.

NEW QUESTION: 91

During red-team testing of an AI system used to make lending decisions, which of the following techniques BEST simulates a data poisoning attack?

- A. Inputting encrypted data into the model
- B. Adding noise to output predictions
- C. Stealing model weights from a deployed API
- D. Corrupting training data sets to manipulate outcomes

Answer: D ([LEAVE A REPLY](#))

AAISM defines data poisoning as the intentional manipulation of training data so that the learned model behaves incorrectly (e.g., skewed lending approvals/denials) while appearing valid. The correct simulation in red-team exercises is to corrupt or seed the training set with adversarial examples or mislabeled records to induce biased or erroneous decision boundaries.

Valid AAISM Dumps shared by EduDump.com for Helping Passing AAISM Exam!
EduDump.com now offer the **newest AAISM exam dumps**, the EduDump.com AAISM exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AAISM dumps with Test Engine here:

<https://www.edudump.com/exams/ISACA/AAISM/premium/> (**432** Q&As Dumps, **35%OFF**)

Special Discount Code: [freecram](#))

NEW QUESTION: 92

A malicious actor is able to retrieve the original data that was used to train an AI model through supplying inputs and observing the corresponding outputs of the model. Which of the following BEST describes this type of attack?

- A. Model inversion
- B. Partial knowledge attack
- C. Membership inference
- D. Model distillation

Answer: ([SHOW ANSWER](#))

A model inversion attack occurs when an attacker reconstructs or retrieves sensitive training data by analyzing the outputs generated from carefully crafted inputs to the AI model. The attack exploits information leakage from the model to infer original training data.

NEW QUESTION: 93

An organization concerned about the ethical and responsible use of a newly developed AI product should consider implementing:

- A. Model cards
- B. Vendor monitoring
- C. An accountability model
- D. Security by design

Answer: C ([LEAVE A REPLY](#))

The AAISM framework highlights that organizations adopting AI must ensure accountability structures are in place to govern ethical and responsible use. An accountability model assigns clear responsibility for decisions, outputs, and risks related to AI systems. While model cards provide transparency about a model's design and performance, they are primarily documentation tools. Vendor monitoring focuses on third-party oversight, not internal accountability. Security by design improves resilience but does not by itself address ethical use. The governance approach that most directly supports responsible and ethical AI deployment is an accountability model.

NEW QUESTION: 94

Which of the following is the GREATEST concern when decommissioning an agentic AI system that autonomously generated transient identities and dynamically created code during its operation?

- A. Diminished clarity in the decision-making audit trail
- B. Persistence of artifacts that could allow unauthorized activities after system shutdown
- C. Increased cost due to indefinite retention of operational logs
- D. Loss of system performance data needed for retrospective evaluation

Answer: ([SHOW ANSWER](#)**)**

Agentic AI systems that generate transient identities and dynamic code may leave behind residual artifacts, credentials, scripts, or processes after shutdown. These remnants could be exploited for unauthorized access or malicious activity, making persistence of artifacts the greatest security concern during decommissioning.

NEW QUESTION: 95

An organization develops and implements an AI-based plug-in for users that summarizes their individual emails. Which of the following is the GREATEST risk associated with this application?

- A. Lack of application vulnerability scanning
- B. Data format incompatibility
- C. Insufficient rate limiting for APIs
- D. Inadequate controls over parameters

Answer: D ([LEAVE A REPLY](#))

According to AAISM risk management guidance, the greatest risk in AI applications handling personal communication data is inadequate parameter controls, which may allow unintended access, manipulation, or leakage of sensitive information. Plug-ins that interact with emails must enforce strict parameter validation and security restrictions to prevent unauthorized or manipulated inputs. While vulnerability scanning, format incompatibility, and API rate limiting are valid concerns, they are secondary. The primary risk is a lack of strong parameter controls that could expose sensitive content.

NEW QUESTION: 96

A backdoor into an AI model was identified through an investigation into a specific set of fraudulent transactions. Which of the following is the BEST course of action?

- A. Inspect the model source code.
- B. Filter and sanitize input.
- C. Review access to the model weights.
- D. Revert to a prior version of the model.

Answer: ([SHOW ANSWER](#))

A backdoor indicates the model itself has been compromised. Reverting to a previously trusted version removes the malicious modification and restores integrity, providing the quickest and most effective way to eliminate the backdoor before further remediation steps are taken.

NEW QUESTION: 97

The BEST approach for aligning AI security requirements with the enterprise security program is to:

- A. embed security by design principles throughout the AI life cycle.
- B. conduct periodic external audits to validate control alignment.
- C. provide organization-wide training on ethical and responsible AI use.
- D. treat AI systems as standard IT assets and apply established security practices.

Answer: ([SHOW ANSWER](#))

Embedding security by design throughout the AI life cycle ensures that security requirements are integrated into planning, development, deployment, operation, and retirement activities from the beginning. This creates consistent alignment between AI initiatives and the enterprise security program while proactively addressing AI-specific risks.

NEW QUESTION: 98

Which of the following will BEST reduce data bias in machine learning (ML) algorithms?

- A. Adopting a more simplified model
- B. Utilizing unstructured data sets
- C. Diversifying the model training data
- D. Securing the model training data

Answer: ([SHOW ANSWER](#))

AAISM guidance clearly states that the most effective way to mitigate data bias is through diverse training data that fairly represents all relevant populations, scenarios, and contexts. Simplified models may reduce complexity but do not remove bias. Unstructured data sets may introduce new errors without addressing fairness. Securing training data protects confidentiality and integrity but does not resolve representational imbalance. Therefore, the best practice for reducing bias in ML is diversification of training datasets.

NEW QUESTION: 99

Which of the following is the PRIMARY advantage of applying data minimization and federated learning methods within AI security management?

- A. Facilitating the adaptation of data governance frameworks to AI use cases
- B. Mitigating privacy concerns while fostering responsible use of AI
- C. Streamlining AI security infrastructure configurations by unifying data repositories
- D. Ensuring personal data is only used by AI solutions when there is a legitimate business need

Answer: ([SHOW ANSWER](#))

Data minimization and federated learning reduce exposure of sensitive information by limiting the amount of centralized personal data used during AI training and processing. This helps protect privacy while still enabling effective and responsible AI model development and operation.

NEW QUESTION: 100

When an attacker uses synthetic data to reverse engineer an organization's AI model, it is an example of which of the following types of attack?

- A. Distillation
- B. Inversion
- C. Prompt
- D. Poisoning

Answer: ([SHOW ANSWER](#))

Explanation:

Model distillation attacks involve using synthetic or carefully crafted data to query an AI model repeatedly and reconstruct its underlying parameters or behavior. This allows attackers to approximate the original model without direct access, effectively reverse engineering it.

NEW QUESTION: 101

Which area of intellectual property law presents the GREATEST challenge in determining copyright protection for AI-generated content?

- A. Enforcing trademark rights associated with AI systems
- B. Determining the rightful ownership of AI-generated creations
- C. Protecting trade secrets in AI technologies
- D. Establishing licensing frameworks for AI-generated works

Answer: ([SHOW ANSWER](#))

AAISM governance content highlights that the greatest intellectual property challenge in the context of AI-generated works is determining rightful ownership. Traditional copyright law requires human authorship, but AI-generated creations blur authorship and ownership boundaries, raising legal uncertainty about who can claim rights. Trademark enforcement, trade secret protection, and licensing frameworks are established areas of IP law but do not present the same fundamental challenge as ownership attribution. For AI-generated content, the central legal dilemma is ownership of the creation.

NEW QUESTION: 102

Which of the following is the MOST effective way to mitigate the effects of outdated data in a large language model (LLM)?

- A. Replacing the model with a more complex model in the solution
- B. Adding a new model trained with new information to the solution
- C. Deploying retrieval-augmented generation (RAG) data to the model
- D. Retraining the model periodically with an updated training dataset

Answer: ([SHOW ANSWER](#))

Retrieval-augmented generation (RAG) enables the LLM to access current external information at query time instead of relying solely on static training data. This is the most effective way to mitigate outdated knowledge because it provides dynamically updated and context-relevant information without requiring full model retraining.

NEW QUESTION: 103

The MOST important consideration when establishing key risk indicators (KRIs) to monitor AI model drift in an operational machine learning (ML) system is adopting KRIs that:

- A. measure utilization rates and trends for AI computational resources.
- B. identify variations in model accuracy and outputs.
- C. monitor latency and AI infrastructure uptime.
- D. track user satisfaction levels with AI-generated outputs.

Answer: ([SHOW ANSWER](#))

AI model drift occurs when a model's performance or behavior changes over time due to evolving data patterns or environments. KRIs that identify variations in model accuracy and outputs are the most important because they directly indicate whether the model is deviating from expected performance and reliability.

NEW QUESTION: 104

Which of the following would MOST effectively obtain ongoing support from stakeholders to align AI initiatives with business objectives?

- A. Conducting periodic organization-wide AI staff training
- B. Addressing and optimizing AI-related risk
- C. Developing and monitoring the AI strategic roadmap
- D. Quantifying and communicating the value of AI solutions

Answer: (SHOW ANSWER)

Sustained stakeholder sponsorship hinges on demonstrated, quantified business value communicated in terms they own (KPIs, ROI, cost-to-serve, risk-adjusted outcomes). AAISM frames stakeholder alignment as a value-assurance loop: define value hypotheses, measure realized value, and continuously communicate results to sponsors.

NEW QUESTION: 105

Which of the following is the BEST way to reduce the risk of misuse of an AI agent that has access to critical data and systems?

- A. Validate agent compliance with output restrictions
- B. Allow users to configure the agent for productivity
- C. Prohibit users from manipulating agent behavior
- D. Limit human review of AI decisions

Answer: (SHOW ANSWER)

AAISM frameworks highlight output-based controls, including output filtering, restriction validation, and policy-aligned guardrails as primary defenses for AI agents with high privileges. Ensuring the agent does not output unauthorized instructions or sensitive data directly mitigates misuse.

NEW QUESTION: 106

Which of the following is the MOST significant risk associated with the use of AI for autonomous vehicle navigation?

- A. Inefficient vehicle routing
- B. Disclosure of proprietary data
- C. False object detection
- D. Denial of service (DoS) attacks

Answer: C (LEAVE A REPLY)

False object detection poses the greatest risk because incorrect identification or interpretation of objects can directly lead to unsafe driving decisions, accidents, injuries, or loss of life. In autonomous vehicle navigation, accurate real-time perception is critical for safety.

Valid AAISM Dumps shared by EduDump.com for Helping Passing AAISM Exam!

EduDump.com now offer the **newest AAISM exam dumps**, the EduDump.com AAISM exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AAISM dumps with Test Engine here:

<https://www.edudump.com/exams/ISACA/AAISM/premium/> (**432** Q&As Dumps, **35%OFF**

Special Discount Code: freecram)

NEW QUESTION: 107

A field service organization wants to adopt an AI solution from an external vendor. Management feels pressure to implement AI quickly to take advantage of the latest technology. Which of the following is the MOST important course of action before engaging with a vendor?

- A. Choose the vendor with the largest market share and follow the vendor's advice on how to use AI in the field service industry.
- B. Define the organization's business needs, regulatory requirements, and risk tolerance to select the most appropriate solution.
- C. Pick a vendor that bundles AI with existing field service support tools and deploy the solution quickly to show results.
- D. Test several AI tools in different departments and decide which one delivers the most value.

Answer: ([SHOW ANSWER](#))

Clearly defining business needs, regulatory requirements, and risk tolerance ensures the organization selects an AI solution that aligns with its objectives and risk posture. This prevents misalignment, reduces compliance risks, and provides a structured basis for evaluating and engaging with vendors.

NEW QUESTION: 108

An AI application development team has been given access to user information and now must format it to be readable by the AI model. During which phase of the data life cycle would this MOST likely occur?

- A. Data minimization
- B. Data preparation
- C. Data collection
- D. Data normalization

Answer: ([SHOW ANSWER](#))

According to AAISM's data life-cycle model, data preparation is the phase where raw data is transformed into a model-ready format. The materials describe this phase as including "cleaning, encoding, formatting, feature engineering, and other transformations required for model consumption." This directly matches the scenario where a team formats user information to be readable by an AI model.

NEW QUESTION: 109

A financial organization is concerned about the risk of prompt injection attacks on its customer service chatbot. Which of the following controls BEST addresses this concern?

- A. Human-in-the-loop
- B. Input validation
- C. Increasing model parameters
- D. Continuous monitoring

Answer: ([SHOW ANSWER](#))

AAISM describes prompt injection as an attack where adversaries craft inputs that manipulate model behavior or override system instructions. The recommended control pattern is to

implement robust input validation and constraint mechanisms that sanitize and structure user inputs before they are processed by the model. The guidance includes techniques such as template-based prompts, restricted instruction sets, and validation rules to filter malicious or out-of-scope content.

NEW QUESTION: 110

Which of the following would BEST help to prevent the compromise of a facial recognition AI system through the use of alterations in facial appearance?

- A. Enhancing training data to increase variance
- B. Monitoring the system for misuse cases
- C. Fine-tuning the AI model to decrease hallucinations
- D. Implementing a secondary AI system to confirm images

Answer: ([SHOW ANSWER](#))

AAISM materials note that adversaries may attempt to bypass facial recognition by disguising or altering appearance. The most effective mitigation is to enhance training data with a wide range of variances in facial features, lighting, and disguises so the system can robustly detect authentic users despite adversarial attempts. Monitoring and secondary confirmation are supportive controls but are reactive. Fine-tuning to reduce hallucinations is irrelevant in this context, as hallucinations apply more to generative AI. The best preventive measure is strengthening the model with diverse, variance-rich training data.

NEW QUESTION: 111

The PRIMARY benefit of implementing moderation controls in generative AI applications is that it can:

- A. Increase the model's ability to generate diverse and creative content
- B. Optimize the model's response time
- C. Ensure the generated content adheres to privacy regulations
- D. Filter out harmful or inappropriate content

Answer: D ([LEAVE A REPLY](#))

AAISM materials identify the primary benefit of moderation controls in generative AI systems as their ability to filter out harmful, offensive, or inappropriate content before it is delivered to users. This safeguards organizational reputation, compliance, and user trust. While moderation may indirectly support compliance with privacy requirements, its main function is ensuring that outputs align with ethical and safety standards. Moderation does not enhance creativity or response speed. Its primary value is in controlling the quality of generated outputs by blocking harmful content.

NEW QUESTION: 112

Which of the following technologies can be used to manage deepfake risk?

- A. Systematic data tagging
- B. Multi-factor authentication (MFA)

- C. Blockchain
- D. Adaptive authentication

Answer: ([SHOW ANSWER](#))

The AAISM study material highlights blockchain as a control mechanism for managing deepfake risk because it provides immutable verification of digital media provenance. By anchoring original data signatures on a blockchain, organizations can verify authenticity and detect tampered or synthetic content. Data tagging helps organize but does not guarantee authenticity. MFA and adaptive authentication strengthen identity security but do not address content manipulation risks. Blockchain's immutability and traceability make it the recognized technology for mitigating deepfake challenges.

NEW QUESTION: 113

Which of the following control strategies BEST mitigates the risk of user harm from a large language model (LLM) generating offensive or biased content?

- A. Static prompt templates
- B. Output filtering and moderation
- C. Minimization of logging
- D. Rate limiting and timeout enforcement

Answer: ([SHOW ANSWER](#))

Output filtering and moderation directly address the risk of harmful, offensive, or biased content by reviewing and controlling generated responses before they reach users. This helps enforce safety, ethical, and compliance standards in LLM outputs.

NEW QUESTION: 114

An auditor is conducting a third-party review of an AI model used in construction project scheduling. Which documentation would BEST support transparency and regulatory compliance?

- A. A model card
- B. A CI/CD pipeline log
- C. An organization-wide data dictionary
- D. A performance tuning log

Answer: ([SHOW ANSWER](#))

A model card provides structured documentation about the AI model's purpose, data sources, performance, limitations, and ethical considerations, supporting transparency and demonstrating compliance with regulatory and governance expectations.

NEW QUESTION: 115

Which of the following is the GREATEST benefit of employing human-in-the-loop when deploying AI models?

- A. Enhanced accuracy
- B. Enhanced safety
- C. Improved automation

D. Improved security

Answer: ([SHOW ANSWER](#))

Human-in-the-loop provides oversight and intervention in AI decision-making, reducing the risk of harmful or unintended outcomes and ensuring safe operation, especially in critical or high-impact scenarios.

NEW QUESTION: 116

Which of the following actions would BEST help an organization be prepared for a deepfake attack?

- A. Improving compliance with regional AI laws and regulations**
- B. Distributing AI-specific policies and procedures to employees**
- C. Training employees on AI-specific incident response**
- D. Implementing code obfuscation to hide code from attackers**

Answer: C ([LEAVE A REPLY](#))

Training employees on AI-specific incident response equips them to quickly recognize and appropriately respond to deepfake incidents, reducing the impact and improving organizational readiness for this type of threat.

NEW QUESTION: 117

Which of the following BEST ensures AI components are validated during disaster recovery testing?

- A. Running simulated data-loss scenarios by deleting test feature-store records**
- B. Disconnecting model training clusters to test retraining workflows**
- C. Simulating DoS attacks on AI APIs**
- D. Monitoring model performance during failover and recovery**

Answer: ([SHOW ANSWER](#))

AAISM states that AI disaster recovery testing must validate that models behave correctly during failover. The only option that tests actual operational continuity of AI components is: monitoring model performance during failover

This validates stability, functionality, and resilience under disaster conditions.

NEW QUESTION: 118

During the deployment of a generative AI platform, a risk assessment identified threats related to data leakage and prompt manipulation. Which of the following is the BEST way to ensure appropriate control selection?

- A. Rely on vendor-provided security features and seek third-party certifications.**
- B. Map identified AI threats to enterprise control catalogs and integrate AI-specific safeguards.**
- C. Apply AI-specific controls from external frameworks and initiate monitoring to expedite compliance.**
- D. Select controls after deployment and address risk through enhanced monitoring.**

Answer: ([SHOW ANSWER](#))

Mapping identified AI threats to enterprise control catalogs and integrating AI-specific safeguards ensures that controls are selected based on the organization's existing risk management framework while also addressing unique AI-related risks such as prompt manipulation and data leakage. This provides comprehensive, risk-aligned, and governance-consistent protection.

NEW QUESTION: 119

Which of the following BEST reduces the risk of exposing sensitive data through the output of large language models (LLMs) in applications?

- A. Encrypting data in transit and at rest
- B. Conducting adversarial testing
- C. Implementing data sanitization techniques
- D. Enforcing least privilege access

Answer: (SHOW ANSWER)

AAISM materials make clear that the best safeguard against sensitive information being leaked through the outputs of LLMs is data sanitization. This involves filtering, redacting, or masking sensitive content before the model can use it, thereby preventing unintended disclosure in outputs. Encryption protects confidentiality in storage and transmission but does not stop output leaks. Adversarial testing helps identify vulnerabilities but does not prevent exposure by itself. Least privilege access restricts who can interact with the model but does not sanitize the content of its outputs. The control most directly tied to preventing leakage is implementing data sanitization techniques.

NEW QUESTION: 120

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: (SHOW ANSWER)

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.

NEW QUESTION: 121

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 (LEAVE A REPLY)

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.

Valid AAISM Dumps shared by EduDump.com for Helping Passing AAISM Exam!
EduDump.com now offer the **newest AAISM exam dumps**, the EduDump.com AAISM exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AAISM dumps with Test Engine here:

<https://www.edudump.com/exams/ISACA/AAISM/premium/> (**432 Q&As Dumps, 35%OFF**

Special Discount Code: freecram)

NEW QUESTION: 122

Which of the following actions BEST enables the evaluation of bias during an AI impact assessment?

- A. Assessing the AI system's training data to ensure it represents all relevant end-user groups
- B. Comparing the AI system's output against historical data benchmarks
- C. Analyzing the AI system's reaction time under peak workload conditions
- D. Measuring the AI system's performance processing speed under predefined varying workloads

Answer: (SHOW ANSWER)

In AI Security Management, bias evaluation is primarily a data problem before it is a model or performance problem. The official AI Security Management content explains that impact assessments must specifically analyze training data representativeness against the demographics and characteristics of all relevant users and stakeholders. If the data used to train an AI system underrepresents or omits particular groups, the resulting model will systematically disadvantage them, regardless of how fast or "accurate" it is on average.

NEW QUESTION: 123

An aerospace manufacturer prioritizing accuracy and security wants to use generative AI. Which LLM adoption plan BEST aligns with its risk appetite?

- A. Developing a private LLM to automate non-critical functions
- B. Contracting LLM access from a reputable third-party provider

- C. Developing a public LLM to automate critical functions
- D. Purchasing an LLM dataset on the open market

Answer: ([SHOW ANSWER](#))

AAISM notes that high-security industries (e.g., aerospace) should prefer private, controlled environments with restricted data exposure. Developing a private LLM for non-critical workloads minimizes operational and security risk while enabling innovation.

NEW QUESTION: 124

Which of the following MOST effectively addresses bias in generative AI models?

- A. Data minimization
- B. Data augmentation
- C. Adversarial training
- D. Fairness constraints

Answer: ([SHOW ANSWER](#))

AAISM identifies fairness constraints as a direct mechanism to mitigate and control model bias by embedding fairness requirements into optimization objectives during training.

NEW QUESTION: 125

Which of the following is the MOST effective action an organization can take to address data security risk when using generative AI features in an application?

- A. Rely on the AI provider's independent third-party audit reports for assurance
- B. Establish policies and awareness training for acceptable use of AI
- C. Require opt-out provisions for data usage in service agreements
- D. Establish guidelines and best practices with third parties for intellectual property ownership

Answer: ([SHOW ANSWER](#))

AAISM directs organizations to manage third-party AI risks through contractual and technical controls that explicitly govern data use, retention, training/fine-tuning, isolation, and deletion. The most effective data-security action when consuming generative AI features is to require enforceable opt-out provisions that prohibit the provider from using the organization's data for training or secondary purposes and that mandate retention limits and secure deletion.

NEW QUESTION: 126

An AI fraud detection system in a bank uses transaction data that includes customer names and account numbers. Which of the following BEST helps the bank comply with privacy regulations and data classification policies?

- A. Performing periodic audits of the AI model for compliance
- B. Masking or tokenizing customer-related data in use
- C. Security training to avoid leakage of personal data
- D. Conducting awareness training for AI developers

Answer: ([SHOW ANSWER](#))

Masking or tokenizing customer names and account numbers directly protects sensitive and regulated data while allowing the AI system to continue processing transactions. This is the most effective control for supporting compliance with privacy regulations and data classification requirements.

NEW QUESTION: 127

An attacker modifies portions of training data over time within an organization's developing AI system. If this manipulation goes undetected, which of the following is the MOST likely risk?

- A. Biased decision-making within the model
- B. Regulatory noncompliance
- C. Lack of detailed system transparency
- D. Intellectual property infringement

Answer: ([SHOW ANSWER](#))

Gradual manipulation of training data is a form of data poisoning, which directly affects how the model learns patterns. This leads to skewed or biased outputs, making biased decision-making the most likely and immediate risk.

NEW QUESTION: 128

When using third-party pre-trained AI models for customer due diligence in a highly regulated financial industry, which of the following is the PRIMARY security concern?

- A. Lower energy efficiency leading to higher operational costs
- B. Excessive model size requiring additional computing resources
- C. Presence of unvalidated dependencies and typosquatting
- D. Inability to deploy model fully on-premises and vendor lock-in

Answer: ([SHOW ANSWER](#))

Third-party pre-trained models often rely on external libraries and components that may include unvalidated or malicious dependencies. Risks such as typosquatting can introduce compromised packages, creating a direct supply chain security threat that can lead to system compromise or data exposure.

NEW QUESTION: 129

A financial organization is deploying an AI-driven customer service platform that uses customer financial data and personalized responses. To ensure the platform remains compliant with data protection regulations, which of the following controls would be MOST important to implement?

- A. Encrypting all data used by the platform before processing
- B. Providing clear, granular, and revocable consent options to customers
- C. Logging all interactions and storing records to service level agreements (SLAs)
- D. Implementing role-based access control (RBAC) and monitoring activities

Answer: ([SHOW ANSWER](#))

Clear, granular, and revocable customer consent is essential for compliance with data protection regulations when processing personal financial data. It ensures customers understand and

control how their information is used by the AI platform, supporting lawful processing, transparency, and user rights management.

NEW QUESTION: 130

An attacker triggers several AI agents to execute CPU-heavy tasks simultaneously, exhausting compute quotas. Which of the following agentic AI threats is this an example of?

- A. Unbounded consumption
- B. Reflection-loop trap
- C. Intent-breaking manipulation
- D. Identity spoofing

Answer: ([SHOW ANSWER](#))

Unbounded consumption occurs when AI agents are manipulated into excessively consuming computational resources, such as CPU, memory, or API quotas. Triggering multiple CPU-intensive tasks simultaneously to exhaust compute capacity is a direct example of this threat.

NEW QUESTION: 131

Which of the following is the MOST effective action an organization can take to address data security risk when using generative AI features in an application?

- A. Require opt-out provisions for data usage in service agreements
- B. Rely on the AI provider's independent third-party audit reports for assurance
- C. Establish policies and awareness training for acceptable use of AI
- D. Establish guidelines and best practices with third parties for intellectual property ownership

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 132

What is the PRIMARY purpose of a dedicated AI management system policy?

- A. Minimizing environmental impact
- B. Optimizing AI model accuracy
- C. Complying with external regulations
- D. Providing a framework to set AI objectives

Answer: ([SHOW ANSWER](#))

AAISM states that an AI management system policy provides organizational structure by:

- defining AI objectives
- aligning governance
- outlining accountability
- defining roles, responsibilities, and guiding principles

NEW QUESTION: 133

An organization's business model relies on an AI application that handles sensitive customer information. From a governance perspective, which of the following is the MOST important consideration?

- A. Leadership accountability for AI initiatives
- B. Robust awareness training for responsible AI use
- C. Sufficient funding for AI security controls
- D. Diverse AI model training data

Answer: ([SHOW ANSWER](#))

Leadership accountability is the most important governance consideration because governance requires clear ownership, oversight, and responsibility for how AI systems manage sensitive customer information, comply with regulations, and align with organizational risk tolerance and ethical standards.

NEW QUESTION: 134

Which of the following AI-driven systems should have the MOST stringent recovery time objective (RTO)?

- A. Health support system
- B. Credit risk modeling system
- C. Car navigation system
- D. Industrial control system

Answer: ([SHOW ANSWER](#))

AAISM risk guidance notes that the most stringent recovery objectives apply to industrial control systems, as downtime can directly disrupt critical infrastructure, manufacturing, or safety operations. Health support systems also require high availability, but industrial control often underpins safety-critical and real-time environments where delays can result in catastrophic outcomes. Credit risk models and navigation systems are important but less critical in terms of immediate physical and operational impact. Thus, industrial control systems require the tightest RTO.

NEW QUESTION: 135

Which of the following is the MOST effective way to prevent a model inversion attack?

- A. Monitor model output for anomalies
- B. Utilize data pseudonymization
- C. Implement differential privacy during model training
- D. Ensure data minimization

Answer: ([SHOW ANSWER](#))

AAISM identifies differential privacy as the primary mitigation technique against model inversion attacks, which attempt to reconstruct sensitive training data by probing model outputs.

NEW QUESTION: 136

When robust input controls are not practical on a large language model (LLM) to prevent prompt injection attacks from external threats, which of the following would be the BEST compensating control to address the risk?

- A. Review and annotate the AI system's outputs

- B. Implement identity and access management (IAM)
- C. Conduct human reviews of the AI system's inputs
- D. Fine-tune the system to validate the AI system's inputs

Answer: ([SHOW ANSWER](#))

When preventive input hardening isn't feasible for LLMs, AAISM prescribes compensating detective and corrective controls--notably human review and annotation of outputs prior to downstream action--to reduce harm from prompt injection. Output-side review gates prevent untrusted instructions from propagating, enable rapid suppression/feedback loops, and provide labeled examples for subsequent model hardening.

Valid AAISM Dumps shared by EduDump.com for Helping Passing AAISM Exam!
EduDump.com now offer the **newest AAISM exam dumps**, the EduDump.com AAISM exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AAISM dumps with Test Engine here:
<https://www.edudump.com/exams/ISACA/AAISM/premium/> (**432** Q&As Dumps, **35%OFF**)
Special Discount Code: freecram)

NEW QUESTION: 137

The MOST effective AI-driven capability to ensure real-time business continuity is:

- A. Automating system failover to backup servers
- B. Keeping logs of incidents for future analysis
- C. Predicting potential disruptions without taking preventive action
- D. Sending alerts while waiting for manual intervention

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 138

Which of the following is the GREATEST threat to an organization where shadow generative AI is prevalent?

- A. Accidental leakage of proprietary data
- B. Software licensing violations
- C. Employee misuse of the AI system
- D. Limited technical knowledge of AI operations

Answer: ([SHOW ANSWER](#))

Shadow generative AI often operates outside governance and security controls, increasing the likelihood that sensitive or proprietary information is entered into uncontrolled systems, leading to unintended exposure or data leakage.

NEW QUESTION: 139

An organization's CIO provided the AI steering committee with a list of AI technologies in use and tasked them with categorizing the technologies by risk. Which of the following should the committee do FIRST?

- A. Begin grouping similar AI products and solutions together
- B. Identify vulnerabilities related to the technologies in use
- C. Ensure the AI technologies are included in the asset inventory
- D. Assess risk levels based on risk appetite and regulatory requirements

Answer: C ([LEAVE A REPLY](#))

AAISM governance practices state that before categorizing technologies by risk, the first step is to ensure that all AI systems are documented in the organizational asset inventory. A complete inventory provides the foundation for subsequent risk analysis, accountability, and governance. Grouping solutions, identifying vulnerabilities, and assessing risk levels come afterward, once inventory accuracy is established. Without confirming that the technologies are recorded in the inventory, risk categorization may miss critical assets.

NEW QUESTION: 140

Which risk is BEST mitigated by automatically clearing in-memory state at the end of each agent application session?

- A. Exfiltrated residual data
- B. Model bias drift
- C. Transport Layer Security (TLS) downgrade attacks
- D. Excessive token billing

Answer: ([SHOW ANSWER](#)**)**

Automatically clearing in-memory state after each session reduces the risk that sensitive or residual session data remains accessible to unauthorized users or processes. This helps prevent exposure or exfiltration of confidential information retained in memory after the session ends.

NEW QUESTION: 141

An organization plans to use an external vendor's AI service for customer retention analytics. Management is worried about third-party dependency and the risk of unauthorized access to customer data. Which of the following is the BEST approach to reduce this risk?

- A. Add data protection requirements to the vendor contract and regularly monitor compliance.
- B. Strengthen internal access controls and encrypt customer data before sharing with the vendor.
- C. Choose a vendor with a strong market reputation and integrate the AI service into enterprise architecture (EA).
- D. Conduct an initial risk assessment of the vendor and follow up with AI impact assessments.

Answer: A ([LEAVE A REPLY](#))

Including data protection requirements in the vendor contract establishes enforceable obligations for safeguarding customer data, while ongoing monitoring ensures the vendor continues to meet those obligations. This directly addresses both third-party dependency and unauthorized access risks in a sustainable and accountable manner.

NEW QUESTION: 142

An AI system sponsored by an organization has expanded by connecting with an increased number of third-party systems. Which of the following would be an information security manager's BEST recommendation in this situation?

- A. Monitor server utilization to reach the optimal operational state.
- B. Define clear lines of accountability for control ownership.
- C. Conduct due diligence on new third-party contracts.
- D. Introduce cost distribution methods that are fair to all third parties.

Answer: ([SHOW ANSWER](#))

Expanding connections to third-party systems increases the attack surface and dependency risks. Conducting due diligence ensures that each third party meets security, privacy, and compliance requirements before integration, directly addressing the most critical risk introduced by this expansion.

NEW QUESTION: 143

An attacker crafts inputs to a large language model (LLM) to exploit output integrity controls. Which of the following types of attacks is this an example of?

- A. Prompt injection
- B. Jailbreaking
- C. Remote code execution
- D. Evasion

Answer: ([SHOW ANSWER](#))

According to the AAISM framework, prompt injection is the act of deliberately crafting malicious or manipulative inputs to override, bypass, or exploit the model's intended controls. In this case, the attacker is targeting the integrity of the model's outputs by exploiting weaknesses in how it interprets and processes prompts. Jailbreaking is a subtype of prompt injection specifically designed to override safety restrictions, while evasion attacks target classification boundaries in other ML contexts, and remote code execution refers to system-level exploitation outside of the AI inference context. The most accurate classification of this attack is prompt injection.

NEW QUESTION: 144

An audit reveals that a third-party AI vendor of a financial organization does not provide documentation on model training and testing. Which of the following is the MOST critical risk for the organization?

- A. Lack of model explainability and potential regulatory noncompliance
- B. Inability to ensure data confidentiality and security during training
- C. Slower customer onboarding process and system administration
- D. Higher infrastructure and fine-tuning costs due to unknown model complexity

Answer: ([SHOW ANSWER](#))

Without documentation on how the model was trained and tested, the organization cannot demonstrate transparency, explainability, or proper validation. This creates significant regulatory compliance risk, especially in financial services where decisions must be explainable and auditable.

NEW QUESTION: 145

Which of the following BEST enables an organization to minimize risk from shadow AI?

- A. Implementing a robust framework for AI performance evaluation
- B. Implementing detective controls and training employees on the use of AI
- C. Prohibiting the use of shadow AI and establishing penalties for noncompliance
- D. Prohibiting the use of any type of AI in the organization

Answer: ([SHOW ANSWER](#))

Detective controls provide visibility into unauthorized AI usage, while employee training promotes awareness and responsible behavior. Together, they address both detection and prevention, making them the most effective approach to minimizing risks associated with shadow AI.

NEW QUESTION: 146

AI developers often find deep learning systems difficult to explain PRIMARILY because:

- A. Knowledge dynamically changes without logs
- B. Neural network architectures include statistical methods not fully understood
- C. Algorithms rely on probability theories
- D. Training data is spread across public domains

Answer: ([SHOW ANSWER](#))

AAISM notes that deep learning systems lack transparency due to complex neural architectures, where internal representations are statistical, nonlinear, and not directly interpretable.

NEW QUESTION: 147

An AI security manager is tasked with establishing key risk indicators (KRIs) to effectively monitor the security posture of a newly deployed generative AI system. Which of the following should be the PRIMARY focus of these KRIs?

- A. Compliance with existing AI governance standards and guidelines
- B. Identification of vulnerabilities and control gaps unique to AI systems
- C. Anonymization and de-identification protocols for AI training datasets
- D. Frequency of external AI security audits and reviews

Answer: ([SHOW ANSWER](#))

KRIs should focus on identifying and tracking risks that directly affect the system's security posture. For generative AI, this means monitoring vulnerabilities and control gaps specific to AI components, such as prompt injection, data leakage, or model abuse, which are critical to managing ongoing risk.

NEW QUESTION: 148

Implementing which of the following would MOST effectively address bias in generative AI models?

- A. Data augmentation
- B. Data minimization
- C. Adversarial training
- D. Fairness constraints

Answer: ([SHOW ANSWER](#))

AAISM identifies fairness constraints (e.g., constrained optimization, debiasing objectives, conditional generation controls, and post-processing calibrations) as the most direct, measurable method to mitigate disparate outcomes in generative systems. While data augmentation can help with coverage, and adversarial training improves robustness, fairness constraints explicitly target distributional fairness and outcome equity in generated content, aligning with governance and compliance goals.

NEW QUESTION: 149

Which of the following is the MOST critical success factor for an AI implementation project?

- A. Developing and using model cards
- B. Ensuring AI risk is captured in the risk register
- C. Mapping data throughout the life cycle
- D. Obtaining senior management buy-in

Answer: D ([LEAVE A REPLY](#))

AAISM identifies executive sponsorship and senior management buy-in as the foremost success factor for AI initiatives. It secures resources, resolves cross-functional conflicts, sets risk appetite, and enforces adherence to governance and controls. Model cards (A), risk registers (B), and lifecycle data mapping (C) are vital practices within the program, but without top-level commitment, adoption, funding, and accountability often fail.

NEW QUESTION: 150

Which of the following is the GREATEST contributor to end user security risk associated with generative AI?

- A. Insufficient awareness training
- B. Lack of data classification
- C. Focus on unintentional errors
- D. Lack of insider threat assessments

Answer: ([SHOW ANSWER](#))

End users are a primary interface with generative AI, and without sufficient awareness training they may unknowingly misuse the technology, expose sensitive data, or fall victim to AI-enabled threats, making this the largest contributor to security risk.

NEW QUESTION: 151

AI developers often find it difficult to explain the processes inside deep learning systems PRIMARILY because:

- A. Training data input for learning is spread throughout the public domain and continues to change
- B. Generated knowledge dynamically changes in memory without being tracked by change history logs
- C. Applied algorithms are based on probability theories to improve system performance
- D. Neural network architectures can include statistical methods that are not fully understood

Answer: D (LEAVE A REPLY)

Deep learning models learn high-dimensional, non-linear representations through layered parameterization that resists simple causal narratives. The internal mechanisms (e.g., distributed feature representations and complex statistical transformations) are difficult to map to human-interpretable rules, making explanation challenging. This is the primary reason for explainability difficulty in deep learning.

Valid AAISM Dumps shared by EduDump.com for Helping Passing AAISM Exam!
EduDump.com now offer the **newest AAISM exam dumps**, the EduDump.com AAISM exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AAISM dumps with Test Engine here:
<https://www.edudump.com/exams/ISACA/AAISM/premium/> (**432 Q&As Dumps, 35%OFF**
Special Discount Code: freecram)

NEW QUESTION: 152

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

- A. Monitoring
- B. Maintenance
- C. Preparation
- D. Training

Answer: (SHOW ANSWER)

AAISM identifies training as the phase with the highest inherent risk because this is where:

- data poisoning can occur
- sensitive data may be exposed
- bias can be introduced
- model inversion risks originate
- security and privacy vulnerabilities are embedded

NEW QUESTION: 153

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: ([SHOW ANSWER](#))

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.

NEW QUESTION: 154

Which of the following is the GREATEST concern when a vendor enables generative AI features for an organization's critical system?

- A. Access to the model
- B. Proposed regulatory enhancements
- C. Security monitoring and alerting
- D. Bias and ethical practices

Answer: ([SHOW ANSWER](#))

When enabling genAI capabilities in a critical system, AAISM prioritizes controlling access to the model and its interfaces (prompt surfaces, context windows, tools/functions, and connected data) because exposure expands the attack surface for prompt injection, data exfiltration, jailbreaks, and misuse.

NEW QUESTION: 155

Which of the following is the MOST important course of action to ensure an AI system produces ethical results?

- A. Ensuring the use of proprietary algorithms
- B. Embedding responsible practices throughout the life cycle
- C. Training the model with public data
- D. Aligning business priorities with ethical outcomes

Answer: ([SHOW ANSWER](#))

Embedding responsible practices throughout the AI life cycle ensures that ethical considerations are continuously addressed during design, development, deployment, operation, and retirement. This approach systematically promotes fairness, accountability, transparency, privacy, and risk management across the entire AI system lifecycle.

NEW QUESTION: 156

The GREATEST security risk associated with the use of pre-trained models from public repositories is that the model may:

- A. contain hidden triggers activated by specific inputs.
- B. have lower accuracy than a custom-built model.
- C. require significant maintenance resources.
- D. not be compatible with commercial use.

Answer: ([SHOW ANSWER](#))

Pre-trained models from public repositories may contain hidden backdoors or malicious triggers intentionally inserted during training. These triggers can activate under specific conditions and cause unauthorized or harmful behavior, making this the greatest security risk.

NEW QUESTION: 157

The PRIMARY goal of data poisoning attacks is to:

- A. compromise the confidentiality of output data from the model
- B. compromise the confidentiality of model input data
- C. manipulate the behavior of the model during development
- D. undermine the integrity of the AI system's outputs

Answer: ([SHOW ANSWER](#))

AAISM defines data poisoning as the insertion of malicious or corrupted data into training (or fine-tuning) pipelines to degrade or bias model behavior, thereby compromising output integrity in production.

NEW QUESTION: 158

A CISO learns that a video showing an unverified safety issue with the organization's key product has gone viral. The individual in the video often appears to be blurry, and the video contains random lower-quality sections throughout. Which type of threat is this MOST likely an example of?

- A. Hallucinations
- B. Model drift
- C. Data poisoning
- D. Deepfake

Answer: ([SHOW ANSWER](#))

AAISM defines deepfakes as manipulated media where individuals appear in synthetic or altered video/audio. Indicators include:

- blurred or inconsistent facial rendering
- mismatched frames
- low-quality or distorted transitions

These characteristics match the scenario provided.

NEW QUESTION: 159

Which of the following is BEST to review in order to detect issues related to data used throughout an AI system's life cycle?

- A. Data quality requirements

- B. User access logs
- C. Data retention policies
- D. Data flow diagrams

Answer: ([SHOW ANSWER](#))

Data flow diagrams provide a comprehensive view of how data moves through the system across all stages of the AI life cycle, enabling identification of issues related to data handling, transformation, and usage.

NEW QUESTION: 160

An organization has discovered that employees have started regularly utilizing open-source generative AI without formal guidance. Which of the following should be the CISO's GREATEST concern?

- A. Lack of monitoring
- B. Policy violations
- C. Data leakage
- D. Model hallucinations

Answer: ([SHOW ANSWER](#))

The greatest immediate risk from unsanctioned use of public or open-source generative AI tools is data leakage--employees may paste confidential or regulated information into third-party systems, resulting in loss of confidentiality, regulatory exposure, and loss of intellectual property. AAISM emphasizes that when AI use occurs outside approved channels, the top control priority is preventing exfiltration of sensitive data via prompts, attachments, and context sharing. Monitoring and policy are necessary enablers, but leakage is the highest-impact failure mode in the short term; hallucinations primarily affect accuracy, not confidentiality.

NEW QUESTION: 161

Which of the following should be the PRIMARY consideration for an organization concerned about liabilities associated with unforeseen behavior from agentic AI systems?

- A. Model dependencies
- B. Approved base models
- C. Accountability model
- D. Acceptable risk level

Answer: ([SHOW ANSWER](#))

AAISM governance guidance stresses that when dealing with agentic AI systems capable of autonomous decision-making, the primary consideration is accountability. Without clear accountability structures, unforeseen or harmful outcomes may result in unmitigated liability for the organization. While dependencies, base models, and defined risk levels are important, they do not directly address who is responsible when systems act unpredictably. The key governance safeguard is the implementation of an accountability model that ensures liability and oversight are properly assigned.

NEW QUESTION: 162

Which of the following BEST enables an organization to strengthen information security controls around the use of generative AI applications?

- A. Ensuring controls exceed industry benchmarks
- B. Monitoring AI outputs against policy
- C. Validating AI model training data
- D. Implementing a kill switch

Answer: ([SHOW ANSWER](#))

For generative AI, the primary enterprise security exposure is data and content exfiltration or policy violations at output, including leakage of sensitive data, toxic content, or regulatory non-compliance. AAISM prescribes policy-aligned output monitoring (e.g., DLP checks, PII/PHI detection, toxicity/safety filters, watermark/attribution checks) integrated into inference gateways to enforce organizational policies and evidence compliance.

NEW QUESTION: 163

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: ([SHOW ANSWER](#))

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.

NEW QUESTION: 164

An organization using an AI model for financial forecasting identifies inaccuracies caused by missing data. Which of the following is the MOST effective data cleaning technique to improve model performance?

- A. Increasing the frequency of model retraining with the existing data set
- B. Applying statistical methods to address missing data and reduce bias
- C. Deleting outlier data points to prevent unusual values impacting the model
- D. Tuning model hyperparameters to increase performance and accuracy

Answer: ([SHOW ANSWER](#))

The AAISM study content emphasizes that data quality management is a central pillar of AI risk reduction. Missing data introduces bias and undermines predictive accuracy if not addressed systematically. The most effective remediation is to apply statistical imputation and related methods to fill in or adjust for missing values in a way that minimizes bias and preserves data

integrity. Retraining on flawed data does not solve the underlying issue. Deleting outliers may harm model robustness, and hyperparameter tuning optimizes model mechanics but cannot resolve missing information. Therefore, the proper corrective technique for missing data is the application of statistical methods to reduce bias.

NEW QUESTION: 165

Which of the following processes involves tagging and labeling raw datasets to create training data sets?

- A. Categorization
- B. Transformation
- C. Annotation
- D. Classification

Answer: ([SHOW ANSWER](#))

Annotation is the process of tagging, labeling, or marking raw data so it can be used as structured training data for machine learning and AI models. This enables the model to learn patterns, categories, or relationships from the labeled examples.

NEW QUESTION: 166

A school district contracts a third-party provider for AI-based curriculum recommendations. Which of the following is the BEST way to ensure the vendor uses AI responsibly?

- A. Confirming the AI solution supports single sign-on (SSO)
- B. Verifying the vendor has updated terms of service
- C. Requiring the vendor to provide the model card
- D. Ensuring the vendor offers 24/7 technical support

Answer: ([SHOW ANSWER](#))

AAISM emphasizes transparency artifacts from vendors to enable due diligence and assurance. A model card documents intended use, data sources, limitations, performance across subgroups, known risks, and evaluation procedures--information necessary to assess safety, fairness, and compliance for sensitive contexts like education. SSO and support are useful operational features; generic ToS updates are insufficient without model-specific disclosures.

Valid AAISM Dumps shared by EduDump.com for Helping Passing AAISM Exam!

EduDump.com now offer the **newest AAISM exam dumps**, the EduDump.com AAISM exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AAISM dumps with Test Engine here:

<https://www.edudump.com/exams/ISACA/AAISM/premium/> (**432** Q&As Dumps, **35%OFF**

Special Discount Code: **freecram**)

NEW QUESTION: 167

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: ([SHOW ANSWER](#))

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.

NEW QUESTION: 168

A healthcare provider uses a third-party large language model (LLM) to summarize patient records. A risk analysis indicates a high likelihood of potential data leakage in excess of the organization's risk tolerance. Which of the following is the BEST course of action taken by the vendor to reduce risk to an acceptable level?

- A. Implement blue team testing and validate through joint testing.
- B. Conduct red team exercises to simulate injection attempts and model inversion.
- C. Apply input filtering and restrict sensitive data exposure.
- D. Migrate the LLM workload to a private instance and restrict access.

Answer: ([SHOW ANSWER](#))

Migrating the LLM workload to a private instance and restricting access most effectively reduces the risk of sensitive patient data leakage. This approach provides stronger control over data handling, processing environments, access management, and exposure pathways, making it the most appropriate risk reduction measure for healthcare data.

NEW QUESTION: 169

A global retail organization is procuring a third-party sentiment analysis API to process customer feedback. The vendor refuses to provide details and documentation on the training data, model provenance, and security testing. Which of the following is the MOST significant risk of proceeding with this procurement?

- A. Inability to ensure the model's outputs are fair, unbiased, and legally compliant
- B. Decreased performance accuracy compared to an in-house built model

- C. Overreliance on a specific vendor's proprietary interface and technology stack
- D. Increased total cost of ownership (TCO) due to API call fees

Answer: ([SHOW ANSWER](#))

Without transparency into training data, model provenance, and security testing, the organization cannot adequately assess whether the model produces fair, unbiased, secure, and legally compliant outcomes. This creates significant ethical, regulatory, and governance risks when processing customer-related data and decisions.

NEW QUESTION: 170

Which of the following is the MOST important course of action when implementing continuous monitoring and reporting for AI-based systems?

- A. Establish an automated alert system for threshold breaches in risk metrics
- B. Develop standardized risk reporting templates for different stakeholder groups
- C. Implement real-time monitoring of key risk indicators (KRIs) for AI systems
- D. Implement a risk dashboard for visualizing and tracking AI-related risk over time

Answer: ([SHOW ANSWER](#))

The AAISM governance framework specifies that the foundation of continuous monitoring is real-time tracking of key risk indicators. This ensures immediate detection of deviations, model drift, and operational anomalies. Automated alerts, dashboards, and reporting templates all support monitoring, but they rely on the presence of accurate, real-time KRI measurement as their source. Without live monitoring, the other controls are reactive rather than proactive. The most important course of action in establishing effective continuous monitoring is therefore real-time KRI tracking.

NEW QUESTION: 171

Which of the following is MOST important for effective AI risk management?

- A. Utilization of best practice AI risk management frameworks
- B. Internal stakeholder participation in AI risk management processes
- C. Risk measurement during an early stage of the AI system life cycle
- D. Creation of separate risk management processes for AI-specific risk

Answer: ([SHOW ANSWER](#))

AAISM positions early and continuous risk assessment as a critical success factor. The guidance states that AI risk should be "identified, analyzed, and measured starting from the design and concept phases, before significant investment and deployment." This ensures that high-impact risks (e.g., bias, privacy violations, safety issues) can be mitigated or designed-out before they become embedded in production systems.

NEW QUESTION: 172

What is the GREATEST benefit of performing AI security risk assessments?

- A. Updating the risk register
- B. Implementing privacy controls
- C. Enabling risk prioritization

D. Securing appropriate funding

Answer: ([SHOW ANSWER](#))

AAISM emphasizes that the primary value of AI security risk assessments is prioritizing risks based on likelihood, impact, and business relevance.

NEW QUESTION: 173

An organization plans to apply an AI system to its business, but developers find it difficult to predict system results due to lack of visibility to the inner workings of the AI model. Which of the following is the GREATEST challenge associated with this situation?

- A.** Gaining the trust of end users through explainability and transparency
- B.** Assigning a risk owner who is responsible for system uptime and performance
- C.** Determining average turnaround time for AI transaction completion
- D.** Continuing operations to meet expected AI security requirements

Answer: ([SHOW ANSWER](#))

AAISM materials identify explainability and transparency as the greatest challenge when models operate as "black boxes" where inner logic is opaque. Inability to interpret how results are produced undermines the trust of business users, customers, regulators, and auditors. Explainability is emphasized as a critical governance requirement, because without it, ethical validation, accountability, and regulatory compliance are at risk. Assigning risk owners or measuring transaction times are operational concerns, but they do not address the core trust deficit caused by lack of visibility. The greatest challenge in this situation is therefore the loss of end-user trust due to insufficient explainability.

NEW QUESTION: 174

To ensure ethical and responsible AI use, which AI usage policy metric is MOST important to monitor?

- A.** Number of policy violations
- B.** Number of AI projects reviewed for compliance
- C.** Frequency of policy consultations by employees
- D.** Frequency of policy reviews and updates

Answer: ([SHOW ANSWER](#))

AAISM states the most meaningful policy performance metric is how often employees consult AI policies, which reflects:

- awareness
- practical adoption
- reliance on policy guidance
- safe decision-making behavior

NEW QUESTION: 175

Which of the following should be done FIRST when developing an acceptable use policy for generative AI?

- A. Determine the scope and intended use of AI
- B. Review AI regulatory requirements
- C. Consult with risk management and legal
- D. Review existing company policies

Answer: ([SHOW ANSWER](#))

According to the AAISM framework, the first step in drafting an acceptable use policy is defining the scope and intended use of the AI system. This ensures that governance, regulatory considerations, risk assessments, and alignment with organizational policies are all tailored to the specific applications and functions the AI will serve. Once scope and intended use are clearly defined, legal, regulatory, and risk considerations can be systematically applied. Without this step, policies risk being generic and misaligned with business objectives.

NEW QUESTION: 176

Which of the following is the MOST effective control to prevent malicious modification of AI chatbot outputs?

- A. Creating a list of disallowed words and phrases in prompts
- B. Conducting mandatory end-user AI training on safe and secure prompting practices
- C. Establishing role-based access control (RBAC) for internal access
- D. Implementing a governance AI model trained on datasets of known attack examples

Answer: ([SHOW ANSWER](#))

A governance AI model trained on known attack patterns can dynamically detect and block malicious attempts to manipulate chatbot outputs, including prompt injection and adversarial input techniques. This provides adaptive, real-time protection that is more effective than static filtering or user training alone.

NEW QUESTION: 177

An organization is conducting an AI risk assessment and wants to establish risk thresholds for adversarial threats such as evasion attacks and data poisoning. Which of the following is MOST important to be approved by senior management?

- A. The AI training plan focusing on human safety
- B. The extent of human-in-the-loop oversight
- C. The backup policy to prevent loss of data
- D. The acceptable level of model performance degradation

Answer: ([SHOW ANSWER](#))

Senior management should approve the acceptable level of model performance degradation because it defines the organization's risk tolerance for adversarial threats and operational impact. Establishing these thresholds is a governance decision that directly affects business risk acceptance and security posture.

NEW QUESTION: 178

Which of the following is the MOST significant risk during the inference phase of the AI model life cycle?

- A. Undetected erroneous or biased decisions
- B. Increased computational costs for retraining
- C. Inability to measure precision and recall
- D. Lack of an audit trail for model version control

Answer: ([SHOW ANSWER](#))

During the inference phase, the AI model is actively making decisions or predictions that affect real-world outcomes. Undetected erroneous or biased decisions present the most significant risk because they can directly impact users, operations, compliance, and organizational trust.

NEW QUESTION: 179

A CISO must provide KPIs for the organization's newly deployed AI chatbot. Which metrics are BEST?

- A. Response time and throughput
- B. Error rate and bias detection
- C. Customer effort score and user retention
- D. Explainability and F1 score

Answer: ([SHOW ANSWER](#))

AAISM recommends that AI KPIs should emphasize:

- Error rates --measure correctness and reliability
- Bias detection metrics --assess fairness and harm risk

These are the core governance KPIs for AI systems interacting with end users.

NEW QUESTION: 180

When implementing a generative AI system, which of the following approaches will BEST prevent misalignment between the corporate risk appetite and tolerance?

- A. Ensuring effective AI key performance indicators (KPIs)
- B. Performing an AI impact assessment
- C. Creating and maintaining an AI risk register
- D. Establishing and monitoring acceptable levels of AI system risk

Answer: ([SHOW ANSWER](#))

AAISM governance guidance specifies that alignment between AI system adoption and organizational risk appetite is achieved by defining and monitoring acceptable levels of system risk. This ensures that generative AI operations remain within boundaries approved by leadership and compliance frameworks. While KPIs track performance, they do not ensure alignment with risk tolerance. AI impact assessments help identify risks but do not maintain continuous oversight. A risk register records risks but does not dynamically enforce acceptable thresholds.

The most effective governance approach is to establish and monitor acceptable AI system risk levels.

NEW QUESTION: 181

A retailer's third-party large language model (LLM) for customer queries has a high likelihood of prompt injection with critical impact if personally identifiable information (PII) is leaked, exceeding the organization's risk tolerance. Which of the following would BEST reduce this risk to an acceptable level?

- A. Deploy a private LLM instance with reduced external exposure.
- B. Apply data minimization and filtering controls on model inputs.
- C. Enable continuous monitoring to detect malicious activity.
- D. Require third-party security testing and reporting of findings.

Answer: (SHOW ANSWER)

Prompt injection exploits inputs to manipulate the model into exposing sensitive data. Applying data minimization and input filtering directly reduces the likelihood of PII being included in prompts or processed in a way that could be exposed, thereby lowering the risk to an acceptable level.

Valid AAISM Dumps shared by EduDump.com for Helping Passing AAISM Exam!

EduDump.com now offer the **newest AAISM exam dumps**, the EduDump.com AAISM exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AAISM dumps with Test Engine here:

<https://www.edudump.com/exams/ISACA/AAISM/premium/> (432 Q&As Dumps, **35%OFF**

Special Discount Code: freecram)

NEW QUESTION: 182

Which of the following would BEST help to prevent attackers from exploiting an AI system to infer personal data?

- A. Conducting periodic audits of AI system use
- B. Restricting the number of queries per user
- C. Monitoring AI system behavior
- D. Limiting access to training data

Answer: (SHOW ANSWER)

Restricting the number of queries per user helps reduce the likelihood of inference attacks, where attackers repeatedly query an AI system to extract or reconstruct sensitive personal information from the model's responses. Limiting query volume directly constrains the attacker's ability to gather enough information for successful inference.

NEW QUESTION: 183

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: ([SHOW ANSWER](#))

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.

NEW QUESTION: 184

Differential privacy BEST helps to mitigate vulnerabilities in machine learning (ML) models by:

- A. preventing adversarial attacks by introducing noise directly into model weights during training.
- B. removing outliers and anomalies from the dataset to reduce the risk of overfitting.
- C. adding controlled noise to the training data, limiting the amount of information any single data point can reveal.
- D. encrypting the model's predictions to ensure data confidentiality and privacy during inference.

Answer: ([SHOW ANSWER](#))

Differential privacy protects sensitive information by adding controlled noise during data processing or training, ensuring that the contribution of any individual data point cannot be easily identified or inferred from the model's outputs. This reduces the risk of privacy leakage and inference attacks.

NEW QUESTION: 185

Which of the following stakeholders would be MOST helpful to include on an incident response team investigating AI-specific incidents involving unusual or incorrect decisions?

- A. Security expert
- B. Model owner
- C. Data owner
- D. Privacy expert

Answer: ([SHOW ANSWER](#))

The model owner has the deepest understanding of how the AI model was designed, trained, and behaves. This expertise is essential for investigating unusual or incorrect decisions, identifying root causes, and determining whether the issue stems from model logic, data, or configuration.

NEW QUESTION: 186

Which BEST addresses hallucination risk in AI systems?

- A. Human oversight
- B. Recursive chunking
- C. Automated output validation
- D. Content enrichment

Answer: ([SHOW ANSWER](#))

AAISM states that human oversight is the strongest control for hallucination risks, especially in high-impact decisions. Humans validate outputs, correct errors, and override unsafe model responses.

NEW QUESTION: 187

When selecting an AI model for a life insurance organization to enhance fraud detection, which factor is MOST critical to ensure secure and reliable deployment?

- A. Robustness
- B. Accuracy
- C. Explainability
- D. Adaptability

Answer: ([SHOW ANSWER](#))

AAISM emphasizes robustness as the key requirement for fraud-detection systems because they must resist adversarial manipulation, data poisoning, spoofing, and input tampering.

NEW QUESTION: 188

The BEST reason for inviting independent reviewers to evaluate the output of an AI application is to:

- A. improve accountability.
- B. improve transparency.
- C. mitigate internal bias.
- D. enhance governance.

Answer: ([SHOW ANSWER](#))

Independent reviewers provide an external perspective that helps identify and reduce biases that internal teams may overlook due to familiarity or shared assumptions, leading to more objective evaluation of AI outputs.

NEW QUESTION: 189

A pharmaceutical organization uses an AI solution from a third-party vendor to analyze confidential research and development data. Which of the following controls is MOST essential to ensure the vendor's AI solution complies with the organization's security requirements?

- A. Communicate security requirements to the vendor.
- B. Minimize sensitive data being used by the vendor's AI solution.
- C. Implement a private instance of the solution.
- D. Review the vendor's independent security audit report.

Answer: ([SHOW ANSWER](#))

Reviewing the vendor's independent security audit report provides objective assurance that the vendor's AI solution and related controls have been evaluated against recognized security standards and requirements. This is essential for validating compliance with the organization's security expectations when handling confidential research data.

NEW QUESTION: 190

An organization implements AI-driven behavioral analytics to identify insider threats and compromised credentials. Which of the following is the MOST significant security risk to be addressed by this deployment?

- A. Manipulation by adversaries resulting in inaccurate threat identification
- B. Complexities in threat identification due to rapid evolution of the local IT environment
- C. Increased potential for password attacks and malicious prompt injection
- D. Latency in responding to incidents due to AI alert processing

Answer: ([SHOW ANSWER](#))

Adversaries can attempt to manipulate behavioral analytics models by injecting misleading patterns or adapting their behavior to evade detection, which can result in inaccurate threat identification and undermine the effectiveness of the system.

NEW QUESTION: 191

An organization implementing a large language model (LLM) application notices significant and unexpected cost increases due to excessive computational resource usage. Which vulnerability is MOST likely in need of mitigation?

- A. Excessive agency
- B. Sensitive information disclosure
- C. System prompt leakage
- D. Unbounded consumption

Answer: ([SHOW ANSWER](#))

AAISM highlights unbounded consumption (token/payment exhaustion, unmetered tool calls, prompt bombs) as a key LLM risk affecting cost and availability. Controls include request quotas, max tokens, rate-limits, budget guards, circuit breakers, and cost-aware routing.

NEW QUESTION: 192

A third-party provider is proposing to implement an AI algorithm for access control monitoring to enhance the efficiency of the security control process. Which of the following would be the GREATEST concern for the provider?

- A. Possible revision of service level agreements (SLAs) due to difficulties in predicting AI performance
- B. Challenges in estimating resources required for the upgrade and justifying AI control costs
- C. Difficulty in recruiting technical professionals to analyze anomalous outputs from AI platforms
- D. Potential conflicts with conventional control monitoring systems already in production

Answer: ([SHOW ANSWER](#))

AI-driven access control monitoring can produce variable and less predictable performance outcomes compared to traditional deterministic systems. This creates significant concern for the provider because service level agreements may need revision if AI behavior affects reliability, accuracy, or response expectations.

NEW QUESTION: 193

When designing self-driving vehicles, conformity assessment is PRIMARILY needed to:

- A. assess compliance against AI regulations and frameworks.
- B. verify that the AI system operates reliably.
- C. test for security flaws before the AI system goes live.
- D. protect against unfair decisions.

Answer: ([SHOW ANSWER](#))

Conformity assessment is primarily conducted to evaluate whether an AI system complies with applicable regulations, standards, and governance frameworks. For self-driving vehicles, this is critical because they are high-risk AI systems subject to strict regulatory and safety requirements.

NEW QUESTION: 194

Which of the following is the MOST effective way to mitigate the risk of deepfake attacks?

- A. Relying on human judgment for oversight
- B. Limiting employee access to AI tools
- C. Validating the provenance of the data source
- D. Using a general-purpose large language model (LLM) to detect fraud

Answer: ([SHOW ANSWER](#))

AAISM study content identifies validating the provenance of data sources as the most effective way to counter deepfake risks. Provenance validation ensures that content is authentic, verifiable, and traceable, preventing malicious synthetic media from being trusted as legitimate. Human oversight helps but cannot reliably detect sophisticated fakes. Limiting tool access reduces exposure but does not prevent external attacks. General-purpose LLMs are not optimized for fraud detection. The strongest control is verifying the origin and authenticity of data before acceptance.

NEW QUESTION: 195

Which of the following is the BEST mitigation technique for a model inversion attack?

- A. Regularly updating and retraining the model
- B. Establishing input monitoring and sanitization procedures
- C. Deploying access controls to restrict unauthorized disclosures
- D. Implementing privacy-preserving methods

Answer: ([SHOW ANSWER](#))

Privacy-preserving methods such as differential privacy limit the amount of information a model can reveal about its training data, directly reducing the risk of reconstructing sensitive inputs through model inversion attacks.

NEW QUESTION: 196

Which of the following security framework elements BEST helps to safeguard the integrity of outputs generated by AI algorithms?

- A. Risk exposure due to bias in AI outputs is kept within an acceptable range
- B. Ethical standards are incorporated into security awareness programs
- C. Management is prepared to disclose AI system architecture to stakeholders
- D. Responsibility is defined for legal actions related to AI regulatory requirements

Answer: (SHOW ANSWER)

According to AAISM technical controls, the element of security frameworks that directly safeguards output integrity is ensuring that bias-related risks are maintained within acceptable ranges. This ensures that AI results remain consistent, fair, and reliable, preserving trust in outputs. Ethical awareness, architectural disclosure, and legal responsibilities are governance practices but do not directly secure output integrity. Output integrity is primarily protected through bias management and ongoing evaluation of fairness and accuracy.

Valid AAISM Dumps shared by EduDump.com for Helping Passing AAISM Exam!

EduDump.com now offer the **newest AAISM exam dumps**, the EduDump.com AAISM exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AAISM dumps with Test Engine here:

<https://www.edudump.com/exams/ISACA/AAISM/premium/> (**432** Q&As Dumps, **35%OFF**

Special Discount Code: freecram)

NEW QUESTION: 197

An organization engages its end users to be involved in the AI life cycle to validate security controls. Which of the following MOST effectively supports decision-making regarding ethical and safety controls?

- A. Engaging independent expert assistance
- B. Providing context and understanding potential risk
- C. Defining and documenting technical controls
- D. Updating the software development life cycle methodology

Answer: (SHOW ANSWER)

Providing context and understanding potential risk enables end users to make informed judgments about how AI systems may impact safety and ethics in real-world scenarios, directly supporting effective decision-making for appropriate controls.

NEW QUESTION: 198

A bank is training an AI credit scoring model on a vast amount of sensitive customer data. Which of the following is the MOST critical security risk in this training process?

- A. Data poisoning to corrupt training outcomes
- B. Algorithmic bias leading to discriminatory lending
- C. Model inversion or membership inference attacks
- D. Resources needed to retrain and validate the AI model

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 199

Which of the following BEST strengthens information security controls around the use of generative AI applications?

- A. Ensuring controls exceed industry benchmarks
- B. Monitoring AI outputs against policy
- C. Implementing a kill switch
- D. Validating AI model training data

Answer: ([SHOW ANSWER](#))

AAISM identifies continuous monitoring of AI outputs--especially generative outputs--as the most effective security control, ensuring that violations, unsafe responses, data leakage, and policy-breaking behavior are detected and corrected.

NEW QUESTION: 200

An organization is deploying an automated AI cybersecurity system. Which of the following would be the MOST effective strategy to minimize human error and improve overall security?

- A. Conducting periodic penetration testing
- B. Using historical data to train AI detection software
- C. Utilizing machine learning (ML) algorithms to ensure responsible use
- D. Implementing manual monitoring of potential alerts

Answer: ([SHOW ANSWER](#))

Training detection models on relevant, representative historical data improves signal quality, reduces false positives, and automates triage--directly lowering human workload and error rates (e.g., alert fatigue, missed correlations). Penetration testing is valuable but episodic and does not systematically reduce day-to-day operator error. "Ensure responsible use" is a governance aim, not a concrete method to cut human error in detection. Manual monitoring increases reliance on human judgment and is prone to inconsistency.

Valid AAISM Dumps shared by EduDump.com for Helping Passing AAISM Exam!

EduDump.com now offer the **newest AAISM exam dumps**, the EduDump.com AAISM exam questions have been updated and answers have been corrected get the **newest**

EduDump.com AAISM dumps with Test Engine here:

<https://www.edudump.com/exams/ISACA/AAISM/premium/> (432 Q&As Dumps, 35%OFF

Special Discount Code: **freecram**)