Isaca

Exam Questions AAISM

ISACA Advanced in AI Security Management (AAISM) Exam

NEW QUESTION 1
Which testing technique is BEST for determining how an AI model makes decisions?

A. Red team
B. Black box
C. White box
D. Blue team

Answer: C

NEW QUESTION 2
Which of the following is the BEST mitigation control for membership inference attacks on AI systems?

A. Model ensemble techniques
B. AI threat modeling
C. Differential privacy
D. Cybersecurity-oriented red teaming

Answer: C

NEW QUESTION 3
Which attack type is MOST likely to cause model drift?

A. Model stealing
B. Perfect knowledge
C. Data poisoning
D. Membership inference

Answer: C

NEW QUESTION 4
Which of the following should be a PRIMARY consideration when defining recovery point objectives (RPOs) and recovery time objectives (RTOs) for generative AI solutions?

A. Preserving the most recent versions of data models to avoid inaccuracies in functionality
B. Prioritizing computational efficiency over data integrity to minimize downtime
C. Ensuring the backup system can restore training data sets within the defined RTO window
D. Maintaining consistent hardware configurations to prevent discrepancies during model restoration

Answer: C

NEW QUESTION 5
A SaaS-based LLM system has risks including prompt injection, data poisoning, and model exfiltration. What is the BEST way to ensure consistent risk treatment?

A. Apply control baselines from a recognized industry standard
B. Implement an AI threat control matrix mapping threats to controls and assurance
C. Focus on post-deployment red teaming
D. Rely on vendor audit reports and SLAs

Answer: B

NEW QUESTION 6
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: C

NEW QUESTION 7
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: D

NEW QUESTION 8
Which of the following strategies is the MOST effective way to protect against AI data poisoning?

A. Ensuring the model is trained on diverse data sources
B. Increasing model complexity
C. Using robust data validation techniques and anomaly detection
D. Incorporating more features and data into model training

Answer: C

NEW QUESTION 9
A retail organization implements an AI-driven recommendation system that utilizes customer purchase history. Which of the following is the BEST way for the organization to ensure privacy and comply with regulatory standards?

A. Conducting quarterly retraining of the AI model to maintain the accuracy of recommendations
B. Maintaining a register of legal and regulatory requirements for privacy
C. Establishing a governance committee to oversee AI privacy practices
D. Storing customer data indefinitely to ensure the AI model has a complete history

Answer: B

NEW QUESTION 10
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: D

NEW QUESTION 10
Which of the following controls BEST mitigates the risk of bias in AI models?

A. Robust access control techniques
B. Regular data reconciliation
C. Cryptographic hash functions
D. Diverse data sourcing strategies

Answer: D

NEW QUESTION 15
An aerospace manufacturing company that prioritizes accuracy and security has decided to use generative AI to enhance operations. Which of the following large language model (LLM) adoption plans BEST aligns with the company??s risk appetite?

A. Developing a public LLM to automate critical functions
B. Purchasing an LLM dataset on the open market
C. Contracting LLM access from a reputable third-party provider
D. Developing a private LLM to automate non-critical functions

Answer: D

NEW QUESTION 19
An organization is adopting an agentic AI solution from an external vendor to support internal IT operations. Which of the following provides the MOST reliable and independently verifiable evidence of implemented security controls?

A. Industry benchmarking peer review
B. Third-party audit reports
C. Internal red-team testing reports
D. General AI security whitepapers

Answer: B

NEW QUESTION 21
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: D

NEW QUESTION 25

Which of the following is MOST important for an organization to consider when implementing a preventive security safeguard into a new AI product?

A. Input sanitization
B. Model output monitoring
C. Penetration testing
D. Differential privacy

Answer: A

NEW QUESTION 26
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

NEW QUESTION 29
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: C

NEW QUESTION 31
An organization plans to use an open-source foundational AI model. Which of the following is MOST important for the AI governance committee to consider when approving its use?

A. Confidential data leakage
B. AI model accuracy
C. AI model support
D. Employee privacy rights

Answer: A

NEW QUESTION 36
An organization needs large data sets to perform application testing. Which of the following would BEST fulfill this need?

A. Reviewing AI model cards
B. Incorporating data from search content
C. Using open-source data repositories
D. Performing AI data augmentation

Answer: C

NEW QUESTION 38
An organization is deploying an automated AI cybersecurity system. Which strategy MOST effectively minimizes human error and improves security?

A. Manual monitoring of alerts
B. Using historical data to train detection software
C. Utilizing machine learning algorithms to ensure responsible use
D. Conducting periodic penetration testing

Answer: B

NEW QUESTION 40
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: C

NEW QUESTION 43
Which of the following BEST represents a combination of quantitative and qualitative metrics that can be used to comprehensively evaluate AI transparency?

A. AI system availability and downtime metrics

B. AI model complexity and accuracy metrics
C. AI explainability reports and bias metrics
D. AI ethical impact and user feedback metrics

Answer: D

NEW QUESTION 46
The PRIMARY ethical concern of generative AI is that it may:

A. Produce unexpected data that could lead to bias
B. Cause information integrity issues
C. Cause information to become unavailable
D. Breach the confidentiality of information

Answer: B

NEW QUESTION 50
When preparing for an AI incident, which of the following should be done FIRST?

A. Implement a communication channel to report AI incidents
B. Establish a cross-functional incident response team with AI knowledge
C. Establish recovery processes for AI system models and data sets
D. Create containment and eradication procedures for AI-related incidents

Answer: B

NEW QUESTION 53
When using AI as part of incident response, which of the following BEST ensures the automation aligns with regulatory and governance obligations?

A. Use deep learning models to autonomously classify all incidents
B. Train the AI incident response platform to mirror legacy response workflows and log containment
C. Apply anomaly detection models to filter incoming threats and automate containment
D. Implement a tiered automation strategy where severity ratings inform the need for human oversight

Answer: D

NEW QUESTION 54
An organization decides to use an anomaly-based intrusion detection system (IDS) integrated with a generative adversarial network–enabled AI tool. The integrated tool would MOST effectively detect intrusions by leveraging:

A. synthetic intrusion data to train the tool??s components
B. validation data sets to enable highly realistic AI decisions
C. automated rule creation to increase model performance
D. classified real intrusion data based on labeled data

Answer: A

NEW QUESTION 57
Within an incident handling process, which of the following would BEST help restore end user trust with an AI system?

A. The AI model prioritizes incidents based on business impact
B. AI is being used to monitor incident detection and alerts
C. The AI model??s outputs are validated by team members
D. Remediation of the AI system based on lessons learned

Answer: C

NEW QUESTION 59
An organization is looking to purchase an AI application from a vendor but is concerned about the security of its data. Which of the following is the MOST effective way to address this concern?

A. Mandate an AI security audit by an external auditor before procurement
B. Initiate discussions between the organization??s and the vendor??s legal teams
C. Ensure vendors disclose how the application uses the organization??s data
D. Assess the vendor??s publicly available AI usage policy

Answer: C

NEW QUESTION 64
An organization utilizes AI-enabled mapping software to plan routes for delivery drivers. A driver following the AI route drives the wrong way down a one-way street, despite numerous signs. Which of the following biases does this scenario demonstrate?

A. Selection
B. Reporting
C. Confirmation

D. Automation

Answer: D

NEW QUESTION 67
A financial services firm received a regulatory fine after a vendor switched its chatbot??s AI model without due diligence, resulting in unethical investment advice to the firm??s clients. Which of the following controls should be implemented by the firm to BEST prevent recurrence of this scenario?

A. Master services agreement
B. Shared responsibility model
C. Data minimization
D. Change management

Answer: D

NEW QUESTION 72
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: A

NEW QUESTION 74
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: A

NEW QUESTION 77
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: C

NEW QUESTION 80
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: C

NEW QUESTION 81
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: A

NEW QUESTION 86
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: B

NEW QUESTION 87
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: B

NEW QUESTION 92
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. Establish IP ownership guidelines with third parties
B. Require opt-out provisions for data usage
C. Establish policies and awareness training for acceptable AI use
D. Rely on the AI provider??s independent audit reports

Answer: C

NEW QUESTION 94
A large financial institution is integrating a third-party AI solution into its fraud detection system. Which is the BEST way to reduce AI vendor/supply chain risk?

A. Conduct annual vulnerability assessments after integration
B. Establish contractual agreements requiring evidence of secure development practices
C. Use isolated virtual environments to validate integration
D. Focus on performance testing

Answer: B

NEW QUESTION 95
Which phase of the AI data life cycle presents the GREATEST inherent risk?

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

Answer: D

NEW QUESTION 96
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: B

NEW QUESTION 97
A global organization has experienced multiple incidents of staff copying confidential data into public chatbots and acting on the model outputs. Which of the following is MOST important to reduce short-term risk when launching an AI security awareness initiative?

A. Blocking access to public large language models (LLMs) at the network perimeter
B. Requiring employees to complete an annual generic phishing and deepfake awareness module
C. Delivering role-based and scenario-driven AI security training mapped to policy and job functions
D. Publishing an AI acceptable use policy and collecting e-signatures of employees

Answer: C

NEW QUESTION 99
A military contractor discovered that its large language model (LLM) is at high risk of being targeted by advanced persistent threat (APT) actors seeking to exploit the model to access confidential information. Which of the following attacks is the HIGHEST priority to protect against?

A. Model inversion
B. Data poisoning
C. Unauthorized tuning
D. Model distillation

Answer: A

NEW QUESTION 102
The PRIMARY purpose of adopting and implementing AI architecture as part of an organizational AI program is to:

A. ensure the development of powerful, efficient, and scalable AI systems
B. deploy fast and cost-efficient AI systems for rapidly changing environments
C. align the system components of AI with the business goals of the organization
D. provide a basis for identification of threats and vulnerabilities

Answer: C

NEW QUESTION 105
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: B

NEW QUESTION 109
Which of the following key risk indicators (KRIs) is MOST relevant when evaluating the effectiveness of an organization??s AI risk management program?

A. Number of AI models deployed into production
B. Percentage of critical business systems with AI components
C. Percentage of AI projects in compliance
D. Number of AI-related training requests submitted

Answer: C

NEW QUESTION 112
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: C

NEW QUESTION 116
An organization decides to contract a vendor to implement a new set of AI libraries. Which of the following is MOST important to address in the master service agreement to protect data used during the AI training process?

A. Data pseudonymization
B. Continuous data monitoring
C. Independent certification
D. Right to audit

Answer: D

NEW QUESTION 117
Which of the following reviews MUST be conducted as part of an AI impact assessment?

A. Testing, evaluation, validation, and verification
B. Evaluation of model reproducibility
C. Security control self-assessment (CSA)
D. Identification of environmental and societal consequences

Answer: D

NEW QUESTION 122
Which of the following is the MOST important factor to consider when selecting industry frameworks to align organizational AI governance with business objectives?

A. Risk tolerance
B. Risk threshold
C. Risk register
D. Risk appetite

Answer: D

NEW QUESTION 123
When robust input controls cannot prevent prompt injections in an LLM, what is the BEST compensating control?

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

Answer: D

NEW QUESTION 128
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: C

NEW QUESTION 131
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

NEW QUESTION 133
A financial organization relies on AI-based identity verification and fraud detection services. Which of the following BEST integrates AI security risk into the business continuity plan (BCP)?

A. Using explainable AI to document decision paths
B. Periodic retraining using pre-labeled data
C. Including AI model supporting infrastructure in disaster recovery scenarios
D. Duplicating AI microservices across multiple availability zones

Answer: C

NEW QUESTION 137
Which of the following is the MAIN objective of the operational phase of AI life cycle management?

A. Monitor model performance
B. Align the model to business needs
C. Optimize the model??s algorithms
D. Obtain end-user feedback on the model

Answer: A

NEW QUESTION 140
A preliminary risk assessment of a SaaS-based large language model (LLM) business support system has identified prompt injection, data poisoning, and model exfiltration as material threats. Which of the following is the BEST approach to ensure risks are treated consistently?

A. Implementing an AI threat control matrix that maps threats to specific controls and assurance activities
B. Applying control baselines from a recognized industry standard to AI components
C. Relying on vendor independent audit reports and service level agreements (SLAs) as evidence of AI risk coverage
D. Focusing resources on post-deployment red teaming and deferring control selection until post go-live feedback is received

Answer: A

NEW QUESTION 143
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

NEW QUESTION 144
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: D

NEW QUESTION 146
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: B

NEW QUESTION 148
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: D

NEW QUESTION 152
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: C

NEW QUESTION 153
Which of the following controls BEST mitigates the inherent limitations of generative AI models?

A. Ensuring human oversight
B. Adopting AI-specific regulations
C. Classifying and labeling AI systems
D. Reverse engineering the models

Answer: A

NEW QUESTION 154
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

AAISM Practice Exam Features:

* AAISM Questions and Answers Updated Frequently

* AAISM Practice Questions Verified by Expert Senior Certified Staff

* AAISM Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AAISM Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year