



Isaca

Exam Questions AAISM

ISACA Advanced in AI Security Management (AAISM) Exam

NEW QUESTION 1

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 2

The PRIMARY purpose of adopting and implementing AI architecture within an organizational AI program is to:

- A. Deploy fast and cost-efficient AI systems
- B. Provide a basis for identifying threats and vulnerabilities
- C. Align AI system components with business goals
- D. Ensure powerful and scalable AI systems

Answer: C

NEW QUESTION 3

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: B

NEW QUESTION 4

Which of the following MOST effectively secures ongoing stakeholder support for AI initiatives?

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

Answer: A

NEW QUESTION 5

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 6

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: A

NEW QUESTION 7

Which of the following approaches BEST helps reduce model bias?

- A. Ensuring diversity in training data sources
- B. Utilizing a more complex architecture
- C. Decreasing frequency of model updates
- D. Increasing the number of labels per instance

Answer: A

NEW QUESTION 8

An organization is adopting an agentic AI solution from an external vendor to support its internal IT operations. To evaluate the security posture of this system, which of the following provides the MOST reliable and independently verifiable evidence of implemented security controls?

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

Answer: D

NEW QUESTION 9

When creating a use case for an AI model that provides sensitive decisions affecting end users, which of the following is the GREATEST benefit of using model cards?

- A. Ethical considerations of the model are documented
- B. Technical instructions for model deployment are created
- C. Data collection requirements are reduced
- D. Model type selection is documented

Answer: A

NEW QUESTION 10

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 10

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: D

NEW QUESTION 11

During red-team testing of an AI system used for lending decisions, which technique BEST simulates a data poisoning attack?

- A. Adding noise to output predictions
- B. Stealing model weights
- C. Inputting encrypted data
- D. Corrupting training datasets to manipulate outcomes

Answer: D

NEW QUESTION 16

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: A

NEW QUESTION 20

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: B

NEW QUESTION 22

As organizations increasingly rely on vendors to develop AI systems, which of the following is the MOST effective way to monitor vendors and ensure compliance with ethical and security standards?

- A. Conducting regular audits of vendor processes and adherence to AI development guidelines
- B. Requiring vendors to monitor their adherence to ethics and security standards
- C. Mandating that vendors share source code and AI documentation with the contracting party
- D. Allowing vendors to self-attest ethical AI compliance and implement benchmark monitoring

Answer: A

NEW QUESTION 25

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: D

NEW QUESTION 29

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

- A. Attacking the underlying hardware of the AI system
- B. Providing inputs that mislead the AI model into incorrect predictions
- C. Reverse engineering the AI model using social engineering techniques
- D. Conducting denial-of-service (DoS) attacks against AI APIs

Answer: B

NEW QUESTION 30

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 35

Which of the following is a key risk indicator (KRI) for an AI system used for threat detection?

- A. Number of training epochs
- B. Training time of the model
- C. Number of layers in the neural network
- D. Number of system overrides by cyber analysts

Answer: D

NEW QUESTION 36

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

- A. Disconnecting primary model training clusters to test retraining workflow during extended outages
- B. Simulating denial of service (DoS) attacks against AI APIs to evaluate detection capabilities
- C. Running simulated data loss scenarios by erasing test records from the AI system's feature store
- D. Monitoring model performance metrics during failover and recovery to assess system stability

Answer: D

NEW QUESTION 41

Personal data used to train AI systems can BEST be protected by:

- A. Erasing personal data after training
- B. Ensuring the quality of personal data
- C. Anonymizing personal data
- D. Hashing personal data

Answer: C

NEW QUESTION 43

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

- A. Assign staff to review AI model outputs for accuracy
- B. Conduct threat modeling to identify vulnerabilities and possible attack methods
- C. Encrypt the training data and model parameters to prevent unauthorized access

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

Answer: B

NEW QUESTION 45

Which of the following types of data is used to tune hyperparameters?

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

Answer: A

NEW QUESTION 47

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: B

NEW QUESTION 52

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 54

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: C

NEW QUESTION 58

An organization is commissioning a third-party AI system using sensitive data. Which metric is MOST important to consider?

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

Answer: C

NEW QUESTION 60

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: B

NEW QUESTION 61

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 65

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: D

NEW QUESTION 67

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: D

NEW QUESTION 70

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

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

Answer: C

NEW QUESTION 74

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: A

NEW QUESTION 76

An organization has implemented a natural language processing model to respond to customer questions when personnel are not available. A pre-implementation security assessment revealed attackers could access sensitive company data through a chat interface injection attack. Which of the following is the BEST way to prevent this attack?

- A. Ensuring continuous monitoring and data tagging
- B. Manually reviewing AI model outputs
- C. Implementing input validation and templates
- D. Conducting regular information security audits

Answer: C

NEW QUESTION 81

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: B

NEW QUESTION 84

Which of the following mitigation control strategies would BEST reduce the risk of introducing hidden backdoors during model fine-tuning via third-party components?

- A. Leveraging open-source models and packages
- B. Performing threat modeling and integrity checks
- C. Disabling runtime logs during model training
- D. Implementing unsupervised learning methods

Answer: B

NEW QUESTION 88

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: A

NEW QUESTION 92

Which of the following BEST describes the role of risk documentation in an AI governance program?

- A. Providing a record of past AI-related incidents for audits
- B. Outlining the acceptable levels of risk for AI-related initiatives
- C. Offering detailed analyses of technical risk and vulnerabilities
- D. Demonstrating governance, risk, and compliance (GRC) for external stakeholders

Answer: B

NEW QUESTION 93

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: A

NEW QUESTION 98

Security and assurance requirements for AI systems should FIRST be embedded in the:

- A. Model design phase
- B. Model training phase
- C. Model testing phase
- D. Model deployment phase

Answer: A

NEW QUESTION 101

Cybersecurity teams should FIRST be embedded in the:

- A. Model testing phase
- B. Model deployment phase
- C. Model training phase
- D. Model design phase

Answer: D

NEW QUESTION 105

Which of the following is the MOST effective use of AI in incident response?

- A. Streamlining incident response testing
- B. Automating incident response triage
- C. Improving incident response playbook
- D. Ensuring chain of custody

Answer: B

NEW QUESTION 110

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 112

Which strategy BEST ensures generative AI tools do not expose company data?

- A. Conducting an independent AI data audit
- B. Implementing a solution prohibiting input of sensitive data
- C. Testing AI tools before implementation

D. Ensuring AI tools comply with local regulations

Answer: B

NEW QUESTION 117

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: A

NEW QUESTION 120

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: A

NEW QUESTION 124

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

- A. Attacking underlying hardware
- B. Providing inputs that mislead the model into incorrect predictions
- C. Reverse-engineering the model using social engineering
- D. Conducting denial-of-service attacks on AI APIs

Answer: B

NEW QUESTION 128

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: D

NEW QUESTION 129

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: B

NEW QUESTION 132

Which of the following datasets is used to tune hyperparameters?

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

Answer: A

NEW QUESTION 136

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: C

NEW QUESTION 138

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: A

NEW QUESTION 143

A vendor switched its chatbot's AI model without due diligence, causing unethical investment advice. What control BEST prevents this scenario?

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

Answer: B

NEW QUESTION 145

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 146

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

- A. Conducting regular security audits
- B. Manually reviewing AI model outputs
- C. Implementing input validation and templates
- D. Ensuring continuous monitoring and tagging

Answer: C

NEW QUESTION 150

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 153

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions 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

100% Actual & Verified — Instant Download, Please Click
[Order The AAISM Practice Test Here](#)