

Isaca

Exam Questions AAISM

ISACA Advanced in AI Security Management (AAISM) Exam



NEW QUESTION 1

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: D

NEW QUESTION 2

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: C

NEW QUESTION 3

From a risk perspective, which of the following is the MOST important step when implementing an adoption strategy for AI systems?

- A. Benchmarking against peer organizations' AI risk strategies
- B. Implementing a robust risk analysis methodology tailored to AI-specific tasks
- C. Conducting an AI risk assessment and updating the enterprise risk register
- D. Establishing a comprehensive AI risk assessment framework

Answer: C

NEW QUESTION 4

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 5

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

- A. The exposure of personal information may result in litigation
- B. The publicly available output of the model may include false or defamatory statements about individuals
- C. The output may reveal information about individuals or groups without their knowledge
- D. The exposure of personal information may lead to a decline in public trust

Answer: C

NEW QUESTION 6

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 7

A global organization experienced multiple incidents of staff pasting confidential data into public chatbots. Which action is MOST important to reduce short-term risk?

- A. Deliver role-based, scenario-driven AI security training mapped to job functions
- B. Require employees to complete an annual generic phishing and deepfake module
- C. Publish an AI acceptable use policy and collect signatures
- D. Block access to public LLMs at the network perimeter

Answer: A

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 14

Which of the following would BEST help mitigate vulnerabilities associated with hidden triggers in generative AI models?

- A. Regularly retraining the model using a diverse data set
- B. Applying differential privacy and masking sensitive patterns in the training data
- C. Incorporating adversarial training to expose and neutralize potential triggers
- D. Monitoring model outputs and suspicious patterns to detect trigger activations

Answer: C

NEW QUESTION 19

An organization is deploying a large language model (LLM) and is concerned that input manipulations may compromise its integrity. Which of the following is the MOST effective way to determine an acceptable risk threshold?

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

Answer: D

NEW QUESTION 21

Which of the following is the BEST reason to immediately disable an AI system?

- A. Excessive model drift
- B. Slow model performance
- C. Overly detailed model outputs
- D. Insufficient model training

Answer: A

NEW QUESTION 26

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: B

NEW QUESTION 31

Which of the following is BEST for analyzing true positives, true negatives, false positives, and false negatives produced by an AI model?

- A. Hyperparameter tuning
- B. Precision
- C. Confusion matrix
- D. Recall

Answer: C

NEW QUESTION 36

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: A

NEW QUESTION 38

When addressing privacy concerns related to AI, what is the GREATEST significance of user consent?

- A. It prevents unauthorized access to data
- B. It enables deletion/modification of personal data
- C. It allows the organization to process user data in the AI system
- D. It helps detect bias and ensure fairness

Answer: C

NEW QUESTION 42

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: B

NEW QUESTION 47

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: C

NEW QUESTION 48

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: C

NEW QUESTION 53

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: C

NEW QUESTION 56

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 57

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

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

Answer: D

NEW QUESTION 59

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 62

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

NEW QUESTION 67

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

- A. Delay implementation until more data scientists are hired
- B. Increase budgets for AI certifications
- C. Update the security program to include cross-functional AI-specific responsibilities
- D. Transition responsibilities to external consultants

Answer: C

NEW QUESTION 68

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 70

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 74

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 77

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 79

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 83

Which of the following recommendations would BEST help a service provider mitigate the risk of lawsuits arising from generative AI's access to and use of internet data?

- A. Activate filtering logic to exclude intellectual property flags
- B. Disclose service provider policies to declare compliance with regulations
- C. Appoint a data steward specialized in AI to strengthen security governance
- D. Review log information that records how data was collected

Answer: A

NEW QUESTION 87

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 88

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 90

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: D

NEW QUESTION 94

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 95

Which of the following is MOST important to consider when validating a third-party AI tool?

- A. Terms and conditions
- B. Right to audit
- C. Industry analysis and certifications
- D. Roundtable testing

Answer: B

NEW QUESTION 100

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 104

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

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

Answer: D

NEW QUESTION 106

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 109

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 112

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 116

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

NEW QUESTION 118

Which of the following is the GREATEST benefit of implementing an AI tool to safeguard sensitive data and prevent unauthorized access?

- A. Timely analysis of endpoint activities
- B. Timely initiation of incident response
- C. Reduced number of false positives
- D. Reduced need for data classification

Answer: C

NEW QUESTION 121

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

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

Answer: A

NEW QUESTION 122

Which of the following should be the MOST important consideration when conducting an AI impact assessment?

- A. Achieve business objectives
- B. Effect on employee retention
- C. Security awareness training
- D. Reputation of the organization

Answer: A

NEW QUESTION 126

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

NEW QUESTION 127

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 132

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 137

An AI system that supports critical processes has deviated from expected performance and is producing biased outcomes. Which of the following is the BEST course of action?

- A. Retrain the model with a new and expanded dataset
- B. Perform a root cause analysis to identify mitigation steps
- C. Conduct audits of the data and the model

D. Activate the model kill switch

Answer: B

NEW QUESTION 138

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 139

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 142

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: A

NEW QUESTION 146

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: A

NEW QUESTION 147

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 151

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 154

A viral video shows a blurry person making claims about a product safety issue. The video has random low-quality sections. This MOST likely represents what threat?

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

Answer: D

NEW QUESTION 158

When deriving statistical information from AI systems, which source of risk is MOST important to address?

- A. Presence of hallucinations
- B. Incomplete outputs
- C. Lack of data normalization
- D. Systemic bias in data sets

Answer: D

NEW QUESTION 159

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 163

An organization plans to leverage AI in the software development process to speed up coding. Which of the following should the information security manager do FIRST?

- A. Conduct an impact assessment
- B. Train developers to verify AI output
- C. Update the security policy to include AI controls
- D. Perform a cost-benefit analysis

Answer: A

NEW QUESTION 168

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 171

Which of the following is the MOST critical key risk indicator (KRI) for an AI system?

- A. The accuracy rate of the model
- B. The amount of data in the model
- C. The response time of the model
- D. The rate of drift in the model

Answer: D

NEW QUESTION 176

Which AI model is BEST suited to ensure explainability in an HR department's pre-screening tool for candidate resumes?

- A. Support vector machine
- B. Neural network
- C. Decision tree
- D. Gradient boosting machine

Answer: C

NEW QUESTION 181

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 185

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

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

Answer: A

NEW QUESTION 187

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: C

NEW QUESTION 192

When evaluating a new AI tool for intrusion prevention, which is MOST important to ensure fit within the existing program architecture?

- A. Ensure automated response orchestration
- B. Prioritize real-time anomaly detection
- C. Confirm tool capabilities align with control objectives
- D. Select a tool that integrates with the SIEM

Answer: C

NEW QUESTION 197

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 198

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: C

NEW QUESTION 203

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

NEW QUESTION 208

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 209

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 212

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 213

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

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

Answer: C

NEW QUESTION 217

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 222

Which of the following factors is MOST important for preserving user confidence and trust in generative AI systems?

- A. Bias minimization
- B. Access controls and secure storage solutions
- C. Transparent disclosure and informed consent
- D. Data anonymization

Answer: C

NEW QUESTION 223

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 227

.....

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](#)