

Isaca

Exam Questions AAISM

ISACA Advanced in AI Security Management (AAISM) Exam



NEW QUESTION 1

An organization has requested a developer to apply AI algorithms to existing modules in order to improve customer service quality. At this stage, which of the following should be considered FIRST?

- A. The developer may need to be held accountable for business inquiries raised by customers
- B. IT management may need to revise the service agreement if AI behavior cannot be predefined
- C. Project sponsors may need to agree on a phased approach in order to ensure safe release
- D. The organization may need to explain the performance of the applied AI algorithm

Answer: B

NEW QUESTION 2

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 3

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 4

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 5

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 6

An organization is facing a deepfake attack intended to manipulate stock prices. The organization's crisis communication plan has been activated. Which of the following is MOST important to include in the initial response?

- A. Conduct employee awareness training on recognizing deepfake videos and audio
- B. Provide clarifying information in a pre-approved public statement
- C. Conduct a detailed forensic analysis to identify the source of the deepfake
- D. Engage with brand monitoring services to track social media activity

Answer: B

NEW QUESTION 7

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 8

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 9

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: B

NEW QUESTION 10

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 10

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 13

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 18

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 19

Which of the following is the MOST important consideration when an organization is adopting generative AI for personalized advertising?

- A. Fraud risk
- B. Reputational risk
- C. Commercial risk
- D. Regulatory risk

Answer: D

NEW QUESTION 21

A regulator warns of increased risk of AI re-identification attacks on anonymized datasets. What should the information security manager do FIRST?

- A. Assume anonymization is permanent and continue operations
- B. Immediately delete anonymized datasets and suspend AI services
- C. Implement a monitoring program including privacy audits and adversarial testing
- D. Establish strong access controls for services using anonymized data

Answer: C

NEW QUESTION 25

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 27

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: C

NEW QUESTION 28

A financial organization is concerned about AI data poisoning. Which control BEST mitigates this risk?

- A. Implementing a break-glass policy
- B. Transparency with customers about data sources
- C. Using training data from multiple sources
- D. Delivering AI-specific security awareness training

Answer: C

NEW QUESTION 31

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: B

NEW QUESTION 34

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 38

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

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

Answer: A

NEW QUESTION 42

A financial organization uses AI to detect potential fraudulent activities but is concerned about the impact of potential data poisoning. Which of the following controls would BEST mitigate this risk?

- A. Being transparent with customers about the data sources
- B. Implementing an updated and tested break-glass policy
- C. Delivering AI-specific security awareness training

D. Using training data from multiple sources

Answer: D

NEW QUESTION 45

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 47

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

NEW QUESTION 49

Which of the following is MOST important to monitor in order to ensure the effectiveness of an organization's AI vendor management program?

- A. Vendor compliance with AI-related requirements
- B. Vendor reviews of external AI threat reports
- C. Vendor results in compliance training programs
- D. Vendor participation in industry AI research

Answer: A

NEW QUESTION 50

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 53

Which of the following is the PRIMARY purpose of a dedicated AI system policy?

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

Answer: C

NEW QUESTION 56

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 58

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 62

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: C

NEW QUESTION 67

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: B

NEW QUESTION 70

Which of the following would BEST ensure a proper business continuity plan (BCP) is in place for an AI solution?

- A. Enhancing monitoring and detection of model failures and anomalies
- B. Implementing access controls to protect the AI system from unauthorized use
- C. Testing the AI infrastructure failover mechanisms
- D. Increasing the detail of AI solution backup and restoration processes

Answer: C

NEW QUESTION 72

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 74

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 78

An organization is reviewing an AI application to determine whether it is still needed. Engineers have been asked to analyze the number of incorrect predictions against the total number of predictions made. Which of the following is this an example of?

- A. Control self-assessment (CSA)
- B. Model validation
- C. Key performance indicator (KPI)
- D. Explainable decision-making

Answer: C

NEW QUESTION 81

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 82

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: C

NEW QUESTION 84

An AI research team is developing a natural language processing model that relies on several open-source libraries. Which of the following is the team's BEST course of action to ensure the integrity of the software packages used?

- A. Maintain a list of frequently used libraries to ensure consistent application in projects
- B. Scan the packages and libraries for malware prior to installation
- C. Use the latest version of all libraries from public repositories
- D. Retrain the model regularly to handle package and library updates

Answer: B

NEW QUESTION 88

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 92

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 93

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 97

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 100

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 105

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 107

Which of the following information is MOST important to include in a centralized AI inventory?

- A. Ownership and accountability of AI systems
- B. AI model use cases
- C. Training data sets
- D. Foundation model and package registry

Answer: A

NEW QUESTION 108

Secure aggregation enhances the security of federated learning systems by:

- A. Processing client updates in isolation to reduce the risk of exposing sensitive information
- B. Applying differential privacy techniques to mask sensitive information in training data
- C. Encrypting individual model updates during transmission to ensure only the server can access the data
- D. Ensuring individual client contributions remain confidential even if the server is compromised

Answer: D

NEW QUESTION 109

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

NEW QUESTION 114

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 118

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: B

NEW QUESTION 120

Which of the following is the MOST effective use of AI-enabled tools in a security operations center (SOC)?

- A. Employing AI-enabled tools to reduce false negatives by detecting subtle attack patterns
- B. Using AI-enabled tools exclusively to classify all types of security incidents
- C. Replacing human analysis with automated AI decision-making processes
- D. Assigning AI-enabled tools to triage non-critical alerts to preserve SOC resources

Answer: A

NEW QUESTION 122

How can an organization best remain compliant when decommissioning an AI system that recorded patient data?

- A. Perform a post-destruction risk assessment
- B. Ensure backups are tested and access controls are audited
- C. Update governance policies based on lessons learned
- D. Ensure a certificate of destruction is received and archived

Answer: D

NEW QUESTION 124

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: C

NEW QUESTION 129

Which of the following would MOST effectively ensure an organization developing AI systems has comprehensive data classification and inventory management?

- A. Creating a centralized team to oversee the classification of data used in AI projects
- B. Conducting quarterly audits of AI data sets for anomalies and missing metadata
- C. Establishing a manual process to categorize data based on business needs and regulatory compliance
- D. Implementing an automated data cataloging tool that integrates with all organizational data repositories

Answer: D

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 138

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: C

NEW QUESTION 141

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 142

Within an incident handling process, which of the following would BEST help restore end-user trust in an AI system?

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

Answer: A

NEW QUESTION 143

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 147

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 148

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 151

.....

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](#)