

AAISM Dumps

ISACA Advanced in AI Security Management (AAISM) Exam

<https://www.certleader.com/AAISM-dumps.html>



NEW QUESTION 1

Which strategy is MOST effective for penetration testers assessing an AI model against membership inference attacks?

- A. Generating synthetic training data
- B. Analyzing AI model confidence scores
- C. Disabling model logging
- D. Measuring accuracy on the test set

Answer: B

NEW QUESTION 2

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 3

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 4

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 5

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 6

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 7

A post-incident investigation finds that an AI-powered anti-money laundering system inadvertently allowed suspicious transactions because certain risk signals were disabled to reduce false positives. Which of the following governance failures does this BEST demonstrate?

- A. Lack of sufficient computing resources for the AI system
- B. Insufficient model validation and change control processes
- C. Excessive reliance on external consultants for model design
- D. Absence of metrics and dashboard for analysts

Answer: B

NEW QUESTION 8

When evaluating a third-party AI service provider, which of the following master services agreement provisions is MOST critical for managing security risk?

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

Answer: A

NEW QUESTION 9

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 10

Which of the following employee awareness topics would MOST likely be revised to account for AI-enabled cyber risk?

- A. Clean desk policy
- B. Social engineering
- C. Malicious insider threats
- D. Authentication controls

Answer: B

NEW QUESTION 10

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 12

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 16

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 17

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 20

A data scientist creating categories and training the algorithm on large data sets is an example of which type of AI model learning technique?

- A. Reinforcement
- B. Unsupervised
- C. Machine learning (ML)
- D. Supervised

Answer: D

NEW QUESTION 23

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

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

Answer: B

NEW QUESTION 26

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 29

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 34

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 36

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 37

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 41

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 43

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 44

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

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 49

When addressing privacy concerns related to AI systems, which of the following is the GREATEST significance of user consent for an organization?

- A. It helps the organization detect biases and ensure fairness
- B. It enables users to delete and modify their personal data
- C. It prevents unauthorized access to data within the AI system
- D. It allows the organization to process user data in the AI system

Answer: D

NEW QUESTION 51

In the context of generative AI, which of the following would be the MOST likely goal of penetration testing during a red-teaming exercise?

- A. Generate outputs that are unexpected using adversarial inputs
- B. Stress test the model's decision-making process
- C. Degrade the model's performance for existing use cases
- D. Replace the model's outputs with entirely random content

Answer: A

NEW QUESTION 56

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 57

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 59

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 62

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 66

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 67

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 70

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

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

Answer: D

NEW QUESTION 74

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 77

What BEST ensures a proper business continuity plan (BCP) for an AI solution?

- A. Enhancing monitoring for model failure
- B. Testing AI infrastructure failover mechanisms
- C. Implementing access controls
- D. Increasing backup restoration detail

Answer: B

NEW QUESTION 78

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 80

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 82

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 87

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 90

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 92

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 94

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 96

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 97

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 98

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 101

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 106

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 111

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 112

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 113

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 117

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 120

Which of the following should be included in an AI acceptable use policy?

- A. AI training data requirements
- B. Data collection and storage processes
- C. Ethical and legal compliance standards
- D. AI monitoring requirements

Answer: C

NEW QUESTION 122

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 127

Who is responsible for implementing recommendations in a final report after an external AI compliance audit?

- A. System architects
- B. Internal auditors
- C. End users
- D. Model owners

Answer: D

NEW QUESTION 130

A data scientist creating categories and training an algorithm on large data sets is performing which learning technique?

- A. Supervised
- B. Reinforcement
- C. Unsupervised
- D. Machine learning (ML)

Answer: A

NEW QUESTION 135

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 136

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

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

Answer: A

NEW QUESTION 141

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

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

Answer: B

NEW QUESTION 145

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 148

Which of the following MOST effectively minimizes the attack surface when securing AI agent components during their development and deployment?

- A. Deploy pre-trained models directly into production.
- B. Consolidate event logs for correlation and centralized analysis.
- C. Schedule periodic manual code reviews.
- D. Implement compartmentalization with least privilege enforcement.

Answer: D

NEW QUESTION 149

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

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

Answer: D

NEW QUESTION 154

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 157

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 162

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

- A. Complying with applicable AI-related regulations
- B. Ensuring ethical use of AI technologies in projects

- C. Establishing an AI governance committee
- D. Protecting enterprise information used by AI projects

Answer: C

NEW QUESTION 164

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 165

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 169

A post-incident investigation finds that an AI-powered anti-money laundering system inadvertently allowed suspicious transactions because certain risk signals were disabled to reduce false positives. Which of the following governance failures does this BEST demonstrate?

- A. Lack of sufficient computing resources for the AI system
- B. Excessive reliance on external consultants for model design
- C. Absence of metrics and dashboards for analysts
- D. Insufficient model validation and change control processes

Answer: D

NEW QUESTION 171

After implementing a third-party generative AI tool, an organization learns about new regulations related to how organizations use AI. Which of the following would be the BEST justification for the organization to decide not to comply?

- A. The AI tool is widely used within the industry
- B. The AI tool is regularly audited
- C. The risk is within the organization's risk appetite
- D. The cost of noncompliance was not determined

Answer: C

NEW QUESTION 173

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 178

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 179

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 183

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 184

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 185

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

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

Answer: C

NEW QUESTION 188

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 189

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

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

Answer: B

NEW QUESTION 192

Which of the following is the BEST approach for minimizing risk when integrating acceptable use policies for AI foundation models into business operations?

- A. Limit model usage to predefined scenarios specified by the developer
- B. Rely on the developer's enforcement mechanisms
- C. Establish AI model life cycle policy and procedures
- D. Implement responsible development training and awareness

Answer: C

NEW QUESTION 194

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 197

An organization uses an AI tool to scan social media for product reviews. Fraudulent social media accounts begin posting negative reviews attacking the organization's product. Which type of AI attack is MOST likely to have occurred?

- A. Model inversion
- B. Deepfake
- C. Availability attack
- D. Data poisoning

Answer: C

NEW QUESTION 201

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 202

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

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

Answer: C

NEW QUESTION 206

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 208

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 210

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 213

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 218

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 222

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

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

Answer: A

NEW QUESTION 224

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 225

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 229

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 230

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

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

Answer: B

NEW QUESTION 234

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 238

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 241

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 245

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 248

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 250

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 251

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 252

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 253

An organization plans to use AI to analyze the shopping patterns of its customers to predict interests and send targeted, customized marketing emails. Which of the following should be done FIRST?

- A. Obtain customer consent
- B. Train the marketing department

- C. Update the terms of service
- D. Verify customer email addresses

Answer: A

NEW QUESTION 258

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 262

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

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

Answer: A

NEW QUESTION 263

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 266

Which of the following approaches BEST enables the separation of sensitive and shareable data to prevent an AI chatbot from inadvertently disclosing confidential information?

- A. Zero Trust
- B. Sandboxing
- C. Siloing
- D. Containerization

Answer: C

NEW QUESTION 270

A health services organization is developing a proprietary generative AI chatbot to assist patients with medical devices. Which of the following should be the organization's HIGHEST priority?

- A. Maximizing neural network size
- B. Tuning algorithms used in the AI model
- C. Maximizing the amount of training data
- D. Selecting the appropriate training data

Answer: D

NEW QUESTION 273

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 278

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

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

D. The organization's compensating controls

Answer: C

NEW QUESTION 283

To ensure AI tools do not jeopardize ethical principles, it is MOST important to validate that:

- A. The organization has implemented a responsible development policy
- B. Outputs of AI tools do not perpetuate adverse biases
- C. Stakeholders have approved alignment with company values
- D. AI tools are evaluated by the privacy department before implementation

Answer: B

NEW QUESTION 287

Which of the following strategies BEST ensures generative AI tools do not expose company data?

- A. Conducting an independent AI data audit
- B. Testing AI tools before implementation
- C. Implementing a solution to prohibit the input of sensitive data
- D. Ensuring AI tools are compliant with local regulations

Answer: C

NEW QUESTION 292

Which BEST addresses hallucination risk in AI systems?

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

Answer: A

NEW QUESTION 297

Which of the following controls BEST mitigates the risk of data poisoning?

- A. Data set restoration
- B. Data validation
- C. Digital watermarking
- D. Intrusion detection

Answer: B

NEW QUESTION 299

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

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

Answer: A

NEW QUESTION 302

Which of the following is MOST important to ensure security throughout the AI data life cycle?

- A. Leveraging selected open-source models
- B. Conducting periodic data reviews
- C. Restricting use of data in third-party models
- D. Maintaining a complete inventory with data lineage records

Answer: D

NEW QUESTION 307

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 308

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 309

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

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

Answer: D

NEW QUESTION 314

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

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

Answer: D

NEW QUESTION 319

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 322

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 323

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 326

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 329

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 334

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AAISM Exam with Our Prep Materials Via below:

<https://www.certleader.com/AAISM-dumps.html>