Exam Questions AAISM

ISACA Advanced in AI Security Management (AAISM) Exam

https://www.2passeasy.com/dumps/AAISM/

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
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 5
In a new supply chain management system, AI models used by participating parties are interactively connected to generate advice in support of management decision making. Which of the following is the GREATEST challenge related to this architecture?

A. Establishing clear lines of responsibility for AI model outputs
B. Identifying hallucinations returned by AI models
C. Determining the aggregate risk of the system
D. Explaining the overall benefit of the system to stakeholders

Answer: A

NEW QUESTION 6
Which of the following should be a PRIMARY consideration when defining recovery point objectives (RPOs) and recovery time objectives (RTOs) for generative AI solutions?

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

Answer: C

NEW QUESTION 7
When evaluating a new AI tool for intrusion prevention, which of the following is the MOST important consideration to ensure the tool fits within the existing program architecture?

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

Answer: A

NEW QUESTION 8
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 9
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 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 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 12
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 13
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 15
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 19
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 20
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 22
What is the PRIMARY purpose of a dedicated AI management system policy?

A. Minimizing environmental impact
B. Optimizing AI model accuracy
C. Complying with external regulations
D. Providing a framework to set AI objectives

Answer: D

NEW QUESTION 24
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 27
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 30
Which of the following approaches BEST helps reduce model bias?

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

Answer: A

NEW QUESTION 35
An aerospace manufacturing company that prioritizes accuracy and security has decided to use generative AI to enhance operations. Which of the following large language model (LLM) adoption plans BEST aligns with the company??s risk appetite?

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

Answer: D

NEW QUESTION 36
An organization is adopting an agentic AI solution from an external vendor to support its internal IT operations. To evaluate the security posture of this system, which of the following provides the MOST reliable and independently verifiable evidence of implemented security controls?

A. Internal red team testing reports
B. Industry benchmarking peer review
C. General AI security whitepapers

D. Third-party audit reports

Answer: D

NEW QUESTION 41
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 46
An organization is evaluating a SaaS-based HR system that uses AI for resume vetting. Which control is MOST important?

A. Inclusion of diverse and representative training data
B. Availability of backups
C. Vendor conformity assessments
D. Encryption and isolation of customer data

Answer: A

NEW QUESTION 48
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 50
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 52
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
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 61
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
The PRIMARY reason to conduct a privacy impact assessment (PIA) on an AI system is to:

A. Identify applicable regulations
B. Determine whether personal data is poisoned
C. Build customer confidence
D. Analyze how personal data is handled

Answer: D

NEW QUESTION 71
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 73
When documenting information about machine learning (ML) models, which of the following artifacts BEST helps enhance stakeholder trust?

A. Hyperparameters
B. Data quality controls
C. Model card
D. Model prototyping

Answer: C

NEW QUESTION 75
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 77
When deriving statistical information generated by AI systems, which of the following types of risk is MOST important to address?

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

Answer: A

NEW QUESTION 81
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 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 84
Employees are regularly using open-source generative AI without guidance. What should be the CISO??s GREATEST concern?

A. Model hallucinations
B. Data leakage
C. Lack of monitoring
D. Policy violations

Answer: B

NEW QUESTION 85
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 88
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 93
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 98
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 101
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 103
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 107
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 108
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 113
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 117
Which of the following is the MOST important course of action when implementing continuous monitoring and reporting for AI-based systems?

A. Establish an automated alert system for threshold breaches in risk metrics
B. Develop standardized risk reporting templates for different stakeholder groups
C. Implement real-time monitoring of key risk indicators (KRIs) for AI systems
D. Implement a risk dashboard for visualizing and tracking AI-related risk over time

Answer: C

NEW QUESTION 120
Which of the following security framework elements BEST helps to safeguard the integrity of outputs generated by AI algorithms?

A. Risk exposure due to bias in AI outputs is kept within an acceptable range
B. Ethical standards are incorporated into security awareness programs
C. Management is prepared to disclose AI system architecture to stakeholders
D. Responsibility is defined for legal actions related to AI regulatory requirements

Answer: A

NEW QUESTION 122
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 123
Which of the following is the MOST effective way to identify and address security risk in an AI model?

A. Assign staff to review AI model outputs for accuracy
B. Conduct threat modeling to identify vulnerabilities and possible attack methods
C. Encrypt the training data and model parameters to prevent unauthorized access
D. Add more data to the model to increase its accuracy and reduce errors

Answer: B

NEW QUESTION 128
Which of the following types of data is used to tune hyperparameters?

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

Answer: A

NEW QUESTION 133
An organization is designing an AI-based credit risk assessment system that will integrate with sensitive financial datasets. Which of the following would BEST support the implementation of security-by-design principles in the AI system??s architecture?

A. Segmenting AI services across containers to manage resource constraints
B. Restricting access to AI models using IP allow lists to reduce public exposure
C. Integrating differential privacy mechanisms into model training to limit data leakage
D. Applying threat modeling specific to AI components before deployment

Answer: D

NEW QUESTION 136
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 141
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 142
Which of the following actions BEST enables the evaluation of bias during an AI impact assessment?

A. Assessing the AI system??s training data to ensure it represents all relevant end-user groups
B. Comparing the AI system??s output against historical data benchmarks
C. Analyzing the AI system??s reaction time under peak workload conditions
D. Measuring the AI system??s performance processing speed under predefined varying workloads

Answer: A

NEW QUESTION 143
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 144
When an attacker uses synthetic data to reverse engineer an organization??s AI model, it is an example of which of the following types of attack?

A. Distillation
B. Inversion
C. Prompt
D. Poisoning

Answer: B

NEW QUESTION 147
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 148
Which AI data management technique involves creating validation and test data?

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

Answer: B

NEW QUESTION 153
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 158
Which of the following BEST enables an organization to maintain visibility to its AI usage?

A. Ensuring the board approves the policies and standards that define corporate AI strategy
B. Maintaining a monthly dashboard that captures all AI vendors
C. Maintaining a comprehensive inventory of AI systems and business units that leverage them
D. Measuring the impact of AI implementation using key performance indicators (KPIs)

Answer: C

NEW QUESTION 159
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 162
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 164
A large corporation has received an influx of sophisticated credential-phishing emails and wants to leverage an AI solution to detect and quarantine these messages before they reach employees. Which of the following blue-team AI features is BEST suited to this task?

A. Large language model (LLM)
B. Natural language processing (NLP)
C. Natural language generation (NLG)
D. Retrieval-augmented generation (RAG)

Answer: B

NEW QUESTION 166
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 170
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 172
Which of the following is the GREATEST benefit of performing AI security risk assessments?

A. Appropriate privacy risk controls are implemented for AI models
B. The appropriate level of funding is secured for AI security risk
C. The risk register is updated with the latest AI risk
D. Risk prioritization decisions are made for AI security

Answer: D

NEW QUESTION 176
When implementing a generative AI system, which of the following approaches will BEST prevent misalignment between the corporate risk appetite and tolerance?

A. Ensuring effective AI key performance indicators (KPIs)
B. Performing an AI impact assessment
C. Creating and maintaining an AI risk register
D. Establishing and monitoring acceptable levels of AI system risk

Answer: D

NEW QUESTION 178
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 179
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 180
Which phase of the AI data life cycle presents the GREATEST inherent risk?

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

Answer: D

NEW QUESTION 185
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 188
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 191
A programmer suspects an AI system is inferring sensitive user information. What is the BEST action?

A. Inform the governance panel
B. Suggest fine-tuning
C. Conduct a code review
D. Alert the CIO

Answer: A

NEW QUESTION 195
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 197
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 202
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 205
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 207
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 209
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 212
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 213
Which of the following types of testing can MOST effectively mitigate prompt hacking?

A. Load
B. Input
C. Regression
D. Adversarial

Answer: D

NEW QUESTION 216
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 218
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 222
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 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 227
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 228
Which of the following would BEST protect trade secrets related to AI technologies during their life cycle?

A. Patenting AI algorithms along with data sets
B. Enforcing trademark rights in AI systems
C. Introducing watermarks when generating AI output
D. Restricting access to sensitive data

Answer: D

NEW QUESTION 230
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 233
Security and assurance requirements for AI systems should FIRST be embedded in the:

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

Answer: A

NEW QUESTION 235
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 239
When preparing for an AI incident, which of the following should be done FIRST?

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

Answer: B

NEW QUESTION 243
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 247
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 251
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 256
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 261
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 265
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 270
Which of the following BEST describes the role of model cards in AI solutions?

A. They are primarily used to visualize the performance of AI models
B. They are used to automatically fine-tune AI models by adjusting hyperparameters based on user feedback
C. They provide a standardized way to document the training data and AI model use cases
D. They help developers create synthetic data and train AI models

Answer: C

NEW QUESTION 275
Which of the following approaches BEST helps to reduce model bias?

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

Answer: D

NEW QUESTION 279
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 284
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 285
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 288
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 290
Which of the following BEST describes an adversarial attack on an AI model?

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

Answer: B

NEW QUESTION 294
Implementing which of the following would MOST effectively address bias in generative AI models?

A. Data augmentation
B. Data minimization
C. Adversarial training
D. Fairness constraints

Answer: D

NEW QUESTION 295
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 300
During the creation of a new large language model (LLM), an organization procured training data from multiple sources. Which of the following is MOST likely to address the CISO's security and privacy concerns?

A. Data augmentation
B. Data minimization
C. Data classification
D. Data discovery

Answer: B

NEW QUESTION 305
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 306
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 310
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 314
What is the GREATEST concern when a vendor enables generative AI features for an organization??s critical system?

A. Security monitoring and alerting
B. Bias and ethical practices
C. Proposed regulatory enhancements
D. Access to the model

Answer: D

NEW QUESTION 317
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 319
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 321
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 326
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 329
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 333
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 335
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 338
Which of the following is the MOST important course of action prior to placing an in-house developed AI solution into production?

A. Perform a privacy, security, and compliance gap analysis
B. Deploy a prototype of the solution
C. Obtain senior management sign-off
D. Perform testing, evaluation, validation, and verification

Answer: D

NEW QUESTION 339
.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AAISM Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AAISM Product From:

https://www.2passeasy.com/dumps/AAISM/

Money Back Guarantee

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