Isaca

Exam Questions AAISM

ISACA Advanced in AI Security Management (AAISM) Exam

NEW QUESTION 1
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 2
Which of the following is the BEST mitigation control for membership inference attacks on AI systems?

A. Model ensemble techniques
B. AI threat modeling
C. Differential privacy
D. Cybersecurity-oriented red teaming

Answer: C

NEW QUESTION 3
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 4
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 5
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 6
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 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
An organization is updating its vendor arrangements to facilitate the safe adoption of AI technologies. Which of the following would be the PRIMARY challenge in delivering this initiative?

A. Failure to adequately assess AI risk
B. Inability to sufficiently identify shadow AI within the organization
C. Unwillingness of large AI companies to accept updated terms
D. Insufficient legal team experience with AI

Answer: C

NEW QUESTION 10
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
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 16
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 18
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 22

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 24
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 25
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 30
For a life insurance company deploying AI for fraud detection, which factor is MOST critical?

A. Robustness
B. Accuracy
C. Explainability
D. Adaptability

Answer: A

NEW QUESTION 34
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 39
An organization is adopting an agentic AI solution from an external vendor to support internal IT operations. Which of the following provides the MOST reliable and independently verifiable evidence of implemented security controls?

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

Answer: B

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 48
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 51
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 52
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 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 61
Which of the following MOST effectively addresses bias in generative AI models?

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

Answer: D

NEW QUESTION 63
An attack has occurred on an AI system that has been in use for two years. Which of the following would BEST mitigate the impact of the attack?

A. Monitoring AI systems for suspicious activities
B. Updating deployed training data with new adversarial data
C. Replacing the AI model with a new model that hides confidence levels
D. Implementing strict access controls to the model??s architecture

Answer: B

NEW QUESTION 65
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 68
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 71
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 75
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 78
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 81
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 85
Which of the following BEST represents a combination of quantitative and qualitative metrics that can be used to comprehensively evaluate AI transparency?

A. AI system availability and downtime metrics
B. AI model complexity and accuracy metrics
C. AI explainability reports and bias metrics
D. AI ethical impact and user feedback metrics

Answer: D

NEW QUESTION 88
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 92
An organization decides to use an anomaly-based intrusion detection system (IDS) integrated with a generative adversarial network (GAN)–enabled AI tool. The integrated tool would MOST effectively detect intrusions by leveraging:

A. Validation data sets to enable highly realistic AI decisions
B. Classified real intrusion data based on labeled data
C. Automated rule creation to increase model performance
D. Synthetic intrusion data to train the tool??s components

Answer: D

NEW QUESTION 97
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 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 103
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 105
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 107
When robust input controls are not practical on a large language model (LLM) to prevent prompt injection attacks from external threats, which of the following would be the BEST compensating control to address the risk?

A. Review and annotate the AI system's outputs
B. Implement identity and access management (IAM)
C. Conduct human reviews of the AI system's inputs
D. Fine-tune the system to validate the AI system's inputs

Answer: A

NEW QUESTION 112
Which of the following BEST ensures AI components are validated as part of disaster recovery testing?

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

Answer: D

NEW QUESTION 113
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 117
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 122
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 124
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 128
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 131
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 133
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 135
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 137
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 139
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 140
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 145
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 150
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 152
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 155
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 159
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 164
Which of the following BEST strengthens information security controls around the use of generative AI applications?

A. Ensuring controls exceed industry benchmarks
B. Monitoring AI outputs against policy
C. Implementing a kill switch
D. Validating AI model training data

Answer: B

NEW QUESTION 168
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 172
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 173
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 177
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 181
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 184
Which of the following key risk indicators (KRIs) is MOST relevant when evaluating the effectiveness of an organization??s AI risk management program?

A. Number of AI models deployed into production
B. Percentage of critical business systems with AI components
C. Percentage of AI projects in compliance

D. Number of AI-related training requests submitted

Answer: C

NEW QUESTION 185
An aerospace manufacturer prioritizing accuracy and security wants to use generative AI. Which LLM adoption plan BEST aligns with its risk appetite?

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

Answer: A

NEW QUESTION 187
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 190
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 193
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 196
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 200
When robust input controls cannot prevent prompt injections in an LLM, what is the BEST compensating control?

A. Fine-tune the system to validate inputs
B. Implement identity and access management (IAM)
C. Conduct human reviews of AI system inputs
D. Review and annotate the AI system's outputs

Answer: D

NEW QUESTION 205
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 210
Which strategy BEST ensures generative AI tools do not expose company data?

A. Conducting an independent AI data audit
B. Implementing a solution prohibiting input of sensitive data
C. Testing AI tools before implementation
D. Ensuring AI tools comply with local regulations

Answer: B

NEW QUESTION 215
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 220
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 225
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 226
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 230
An organization is implementing an AI-based credit assessment engine using internal and third-party customer data. Which of the following BEST aligns with data management controls for the AI life cycle?

A. Documented procedures for data sourcing, lineage tracking, and quality validation
B. Use of hashed identifiers to anonymize datasets used for model validation and internal analytics
C. Encrypted isolation and dynamic access controls on training data pipelines
D. Limitation of model training to structured data from vetted sources to minimize ingestion risk

Answer: A

NEW QUESTION 234
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 239
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 240
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 243
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 247
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 250
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 253
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 257
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 258
Which of the following should be the PRIMARY objective of implementing differential privacy techniques in AI models used for fraud detection systems?

A. Reducing computational resources

B. Enhancing the accuracy of predictions
C. Protecting individual data contributions while allowing statistical analysis
D. Increasing model training speed

Answer: C

NEW QUESTION 263
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 266
A critical AI system shows biased outcomes. What is the BEST course of action?

A. Activate the kill switch
B. Conduct audits of data and model
C. Perform root cause analysis to identify mitigation
D. Retrain the model with a new diverse dataset

Answer: C

NEW QUESTION 269
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 272
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 273
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 277
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 281
.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons 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