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
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 3
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 4
An organization deploying an LLM is concerned input manipulations could compromise security. What is the MOST effective way to determine an acceptable risk threshold?

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

Answer: C

NEW QUESTION 5
Which of the following is the MOST serious consequence of an AI system correctly guessing the personal information of individuals and drawing conclusions based on that information?

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

Answer: C

NEW QUESTION 6
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 7
A retail organization implements an AI-driven recommendation system that utilizes customer purchase history. Which of the following is the BEST way for the organization to ensure privacy and comply with regulatory standards?

A. Conducting quarterly retraining of the AI model to maintain the accuracy of recommendations
B. Maintaining a register of legal and regulatory requirements for privacy
C. Establishing a governance committee to oversee AI privacy practices
D. Storing customer data indefinitely to ensure the AI model has a complete history

Answer: B

NEW QUESTION 8
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 9
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 10
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 10
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 14
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 15
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 18
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 22
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 26
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 28
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 30
An organization is deploying an automated AI cybersecurity system. Which of the following would be the MOST effective strategy to minimize human error and improve overall security?

A. Conducting periodic penetration testing
B. Using historical data to train AI detection software
C. Utilizing machine learning (ML) algorithms to ensure responsible use
D. Implementing manual monitoring of potential alerts

Answer: B

NEW QUESTION 31
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 33
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 36
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 39
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 41
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 45
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 48
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 52
AI developers often find deep learning systems difficult to explain PRIMARILY because:

A. Knowledge dynamically changes without logs
B. Neural network architectures include statistical methods not fully understood
C. Algorithms rely on probability theories
D. Training data is spread across public domains

Answer: B

NEW QUESTION 53
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 56
Which of the following recommendations would BEST help a service provider mitigate the risk of lawsuits arising from generative AI??s access to and use of internet data?

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

Answer: A

NEW QUESTION 61
The PRIMARY ethical concern of generative AI is that it may:

A. Produce unexpected data that could lead to bias
B. Cause information integrity issues
C. Cause information to become unavailable
D. Breach the confidentiality of information

Answer: B

NEW QUESTION 64
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 67
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 72
A CISO has been tasked with providing key performance indicators (KPIs) on the organization??s newly launched AI chatbot. Which of the following are the BEST metrics for the CISO to recommend?

A. Explainability and F1 score
B. Customer effort score and user retention rate
C. Response time and throughput
D. Error rate and bias detection

Answer: D

NEW QUESTION 77
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 78
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 81
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 85
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 86
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 91
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 94
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 98
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 100
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 101
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 102
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 105
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 107
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 108
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 112
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 115
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 116
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 120
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 123
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 126
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 128
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 129
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 131
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 132
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 137
Which of the following is the MOST effective strategy for penetration testers assessing the security of an AI model against membership inference attacks?

A. Disabling AI model logging to reduce noise during testing
B. Measuring AI model accuracy on the test set
C. Analyzing AI model confidence scores to indicate training data
D. Generating synthetic data to replace the training data

Answer: C

NEW QUESTION 139
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 143
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 148
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 150
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 154
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 155
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 158
Which defense is MOST effective against cyberattacks that alter input data to avoid detection?

A. Enhancing model robustness through adversarial training
B. Restricting access to internal model parameters
C. Conducting periodic monitoring of decisions
D. Applying differential privacy to training data

Answer: A

NEW QUESTION 159
After deployment, an AI model??s output begins to drift outside of the expected range. Which of the following is the development team??s BEST course of action?

A. Take the AI model offline
B. Adjust the hyperparameters of the AI model
C. Create an emergency change request to correct the issue
D. Return to an earlier phase in the AI life cycle

Answer: D

NEW QUESTION 162

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 164
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 167
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 169
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 173
Which of the following controls would BEST help to prevent data poisoning in AI models?

A. Increasing the size of the training data set
B. Implementing a strict data validation mechanism
C. Establishing continuous monitoring
D. Regularly updating the foundational model

Answer: B

NEW QUESTION 178
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 181
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 186
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 188
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 193
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 198
Which of the following BEST ensures the integrity of data sets used to train AI models?

A. Collection and retention of only necessary data sets
B. Tracking and verification of data sets via cryptographic controls
C. Appropriate storage of data sets according to documented classification processes
D. Clear documentation of data sources, types used, and processing steps

Answer: B

NEW QUESTION 202
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 207
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 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 212
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 217
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 221
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 223
Which of the following is MOST important for effective AI risk management?

A. Utilization of best practice AI risk management frameworks
B. Internal stakeholder participation in AI risk management processes
C. Risk measurement during an early stage of the AI system life cycle
D. Creation of separate risk management processes for AI-specific risk

Answer: C

NEW QUESTION 227
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 231
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 236
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 241
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 244
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 249
Which of the following datasets is used to tune hyperparameters?

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

Answer: A

NEW QUESTION 250
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 252
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 256
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 257
Which of the following would BEST help to prevent the compromise of a facial recognition AI system through the use of alterations in facial appearance?

A. Enhancing training data to increase variance
B. Monitoring the system for misuse cases
C. Fine-tuning the AI model to decrease hallucinations
D. Implementing a secondary AI system to confirm images

Answer: A

NEW QUESTION 261
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 263
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 266
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 271
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 273
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 277
.....

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