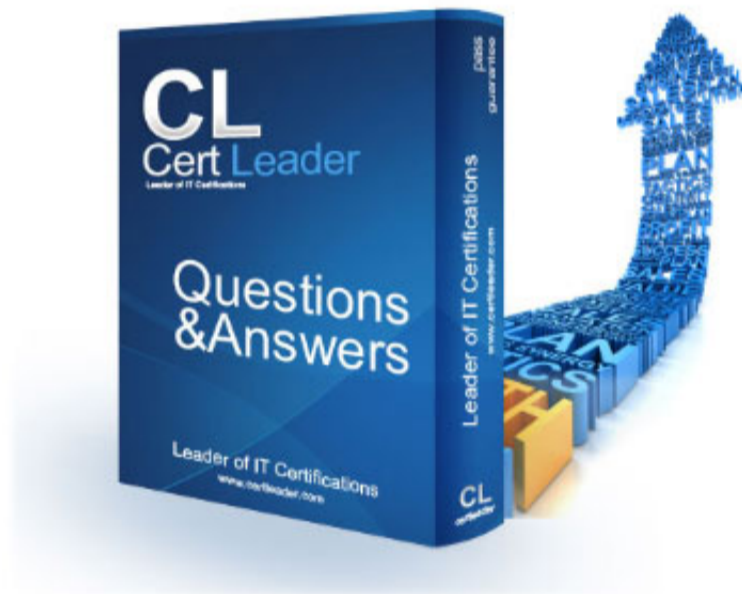


AAISM Dumps

ISACA Advanced in AI Security Management (AAISM) Exam

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



NEW QUESTION 1

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 2

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 3

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 4

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 5

A SaaS-based LLM system has risks including prompt injection, data poisoning, and model exfiltration. What is the BEST way to ensure consistent risk treatment?

- A. Apply control baselines from a recognized industry standard
- B. Implement an AI threat control matrix mapping threats to controls and assurance
- C. Focus on post-deployment red teaming
- D. Rely on vendor audit reports and SLAs

Answer: B

NEW QUESTION 6

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 7

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 8

The PRIMARY goal of data poisoning attacks is to:

- A. compromise the confidentiality of output data from the model
- B. compromise the confidentiality of model input data
- C. manipulate the behavior of the model during development
- D. undermine the integrity of the AI system's outputs

Answer: D

NEW QUESTION 9

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 10

The PRIMARY purpose of adopting and implementing AI architecture within an organizational AI program is to:

- A. Deploy fast and cost-efficient AI systems
- B. Provide a basis for identifying threats and vulnerabilities
- C. Align AI system components with business goals
- D. Ensure powerful and scalable AI systems

Answer: C

NEW QUESTION 10

Which of the following MOST effectively secures ongoing stakeholder support for AI initiatives?

- A. Quantifying and communicating the value of AI solutions
- B. Conducting periodic staff training
- C. Addressing and optimizing AI-related risk
- D. Developing and monitoring an AI strategic roadmap

Answer: A

NEW QUESTION 14

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 17

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 19

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 23

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 26

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 30

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 33

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 35

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 39

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 43

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 45

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 48

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 50

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 53

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 55

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 56

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 57

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

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

Answer: D

NEW QUESTION 60

An organization plans to use an open-source foundational AI model. Which of the following is MOST important for the AI governance committee to consider when approving its use?

- A. Confidential data leakage
- B. AI model accuracy
- C. AI model support
- D. Employee privacy rights

Answer: A

NEW QUESTION 64

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

- A. Delay implementation until more data scientists are hired
- B. Increase budgets for AI certifications
- C. Update the security program to include cross-functional AI-specific responsibilities
- D. Transition responsibilities to external consultants

Answer: C

NEW QUESTION 67

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 70

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 74

When using AI as part of incident response, which of the following BEST ensures the automation aligns with regulatory and governance obligations?

- A. Use deep learning models to autonomously classify all incidents
- B. Train the AI incident response platform to mirror legacy response workflows and log containment
- C. Apply anomaly detection models to filter incoming threats and automate containment
- D. Implement a tiered automation strategy where severity ratings inform the need for human oversight

Answer: D

NEW QUESTION 75

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 79

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 80

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 85

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 90

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 99

A CISO must provide KPIs for the organization's newly deployed AI chatbot. Which metrics are BEST?

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

Answer: B

NEW QUESTION 102

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 105

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 110

Which of the following metrics BEST evaluates the ability of a model to correctly identify all true positive instances?

- A. F1 score
- B. Recall
- C. Precision
- D. Specificity

Answer: B

NEW QUESTION 114

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 119

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 120

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 122

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 124

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

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

- A. Complying with applicable AI-related regulations
- B. Ensuring ethical use of AI technologies in projects
- C. Establishing an AI governance committee
- D. Protecting enterprise information used by AI projects

Answer: C

NEW QUESTION 131

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 134

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 138

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 139

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 141

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 144

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 145

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 150

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 155

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 158

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 161

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 165

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

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

Answer: C

NEW QUESTION 166

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 169

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 171

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 174

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 178

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 182

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 183

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 188

When deriving statistical information from AI systems, which source of risk is MOST important to address?

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

Answer: D

NEW QUESTION 192

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 194

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 198

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 202

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 207

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 211

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 216

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 220

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 222

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 225

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 230

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 232

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 234

An organization is implementing AI agent development across engineering teams. What should AI-specific training focus on?

- A. Prompt injection, agent memory control, insecure tool execution
- B. Dataset bias, explainability, fairness
- C. Output moderation, hallucination handling, policy alignment
- D. API abuse, data leakage, third-party plug-in risk

Answer: A

NEW QUESTION 236

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 238

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 239

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

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

Answer: A

NEW QUESTION 242

Which of the following is the MOST effective action an organization can take to address data security risk when using generative AI features in an application?

- A. Rely on the AI provider's independent third-party audit reports for assurance
- B. Establish policies and awareness training for acceptable use of AI
- C. Require opt-out provisions for data usage in service agreements
- D. Establish guidelines and best practices with third parties for intellectual property ownership

Answer: C

NEW QUESTION 246

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 249

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 253

Which of the following is the BEST control for preventing deepfakes?

- A. Output provenance verification
- B. Regular AI risk assessment
- C. AI governance policies
- D. System input validation

Answer: A

NEW QUESTION 257

An organization plans to apply an AI system to its business, but developers find it difficult to predict system results due to lack of visibility to the inner workings of the AI model. Which of the following is the GREATEST challenge associated with this situation?

- A. Gaining the trust of end users through explainability and transparency
- B. Assigning a risk owner who is responsible for system uptime and performance
- C. Determining average turnaround time for AI transaction completion
- D. Continuing operations to meet expected AI security requirements

Answer: A

NEW QUESTION 262

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 264

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 268

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 272

To ensure the ethical and responsible use of AI, which of the following AI usage policy metrics is MOST important for an organization to monitor?

- A. Frequency of policy consultations by employees
- B. Number of reported policy violations
- C. Number of AI projects that have undergone policy compliance review
- D. Frequency of policy reviews and updates

Answer: C

NEW QUESTION 275

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 276

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 281

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 282

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 285

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 286

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 291

Which of the following datasets is used to tune hyperparameters?

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

Answer: A

NEW QUESTION 293

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 298

Which of the following is the GREATEST risk inherent to implementing generative AI?

- A. Lack of employee training
- B. Unidentified asset vulnerabilities
- C. Inadequate return on investment (ROI)
- D. Potential intellectual property violations

Answer: D

NEW QUESTION 299

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 302

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 305

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 307

Which of the following is the MOST critical success factor for an AI implementation project?

- A. Developing and using model cards
- B. Ensuring AI risk is captured in the risk register
- C. Mapping data throughout the life cycle
- D. Obtaining senior management buy-in

Answer: D

NEW QUESTION 312

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 314

Embedding unique identifiers into AI models would BEST help with:

- A. Preventing unauthorized access
- B. Tracking ownership
- C. Eliminating AI system biases
- D. Detecting adversarial attacks

Answer: B

NEW QUESTION 318

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 323

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 327

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 329

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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