

Isaca

Exam Questions AAISM

ISACA Advanced in AI Security Management (AAISM) Exam



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

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 3

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 4

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 5

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 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

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 8

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

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

Answer: B

NEW QUESTION 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

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 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 15

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 20

An organization implementing a large language model (LLM) application notices significant and 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. System prompt leakage
- D. Unbounded consumption

Answer: D

NEW QUESTION 23

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 28

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 32

Which of the following is the MOST likely cause of model drift?

- A. Data poisoning
- B. Perfect knowledge
- C. Membership inference
- D. Model stealing

Answer: A

NEW QUESTION 34

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 36

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 37

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 41

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 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 44

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

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 54

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 57

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 61

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 64

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 66

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 69

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 71

An organization needs large data sets to perform application testing. Which of the following would BEST fulfill this need?

- A. Reviewing AI model cards
- B. Incorporating data from search content
- C. Using open-source data repositories
- D. Performing AI data augmentation

Answer: C

NEW QUESTION 76

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 81

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 85

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 89

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 94

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

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

Answer: B

NEW QUESTION 98

Which of the following is the MOST important consideration when an organization is adopting generative AI for personalized advertising?

- A. Fraud risk
- B. Reputational risk
- C. Commercial risk
- D. Regulatory risk

Answer: D

NEW QUESTION 103

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 108

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 112

Within an incident handling process, which of the following would BEST help restore end user trust with an AI system?

- A. The AI model prioritizes incidents based on business impact
- B. AI is being used to monitor incident detection and alerts
- C. The AI model's outputs are validated by team members
- D. Remediation of the AI system based on lessons learned

Answer: C

NEW QUESTION 114

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 118

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 121

How can an organization BEST protect itself from payment diversions caused by deepfake attacks impersonating management?

- A. Require mandatory deepfake detection training for all employees
- B. Mandate that payments be sent only once per week
- C. Issue a security policy on deepfakes
- D. Implement resilient payment approval processes

Answer: D

NEW QUESTION 122

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 125

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 126

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 127

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 129

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 133

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 135

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 139

Which of the following should be the PRIMARY objective of implementing differential privacy techniques in AI models leveraging fraud detection systems?

- A. Enhancing the accuracy of predictions to desired levels
- B. Increasing model training speed for an efficient launch
- C. Protecting individual data contributions while allowing statistical analysis
- D. Reducing computational resources required for the model training phase

Answer: C

NEW QUESTION 142

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

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

Answer: A

NEW QUESTION 143

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 145

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 147

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 151

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 153

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 157

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 160

Which of the following BEST addresses risk associated with hallucinations in AI systems?

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

Answer: D

NEW QUESTION 163

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 164

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 167

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

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

Answer: D

NEW QUESTION 171

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 172

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 174

An organization is commissioning a third-party AI system using sensitive data. Which metric is MOST important to consider?

- A. Accessibility rating
- B. Model response time
- C. Accuracy thresholds
- D. Service availability

Answer: C

NEW QUESTION 175

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 177

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 181

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 185

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 186

Which of the following would BEST ensure a proper business continuity plan (BCP) is in place for an AI solution?

- A. Enhancing monitoring and detection of model failures and anomalies
- B. Implementing access controls to protect the AI system from unauthorized use
- C. Testing the AI infrastructure failover mechanisms
- D. Increasing the detail of AI solution backup and restoration processes

Answer: C

NEW QUESTION 187

Which phase of the AI data life cycle presents the GREATEST inherent risk?

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

Answer: D

NEW QUESTION 190

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 193

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 194

A large financial services organization is integrating a third-party AI solution into its critical fraud detection system. Which of the following is the BEST way for the organization to reduce risk associated with AI vendor and supply chain dependencies?

- A. Conducting annual vulnerability assessments of the fraud detection system after integration
- B. Focusing on performance testing to ensure the solution meets operational requirements
- C. Establishing contractual agreements requiring vendors to provide evidence of secure development practices
- D. Implementing isolated virtual environments to validate the integration of the fraud detection system with the solution

Answer: C

NEW QUESTION 197

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

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

Answer: C

NEW QUESTION 199

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 200

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

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

Answer: C

NEW QUESTION 201

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 202

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 203

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 208

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 210

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 211

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 214

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 215

Which of the following is the MOST critical key risk indicator (KRI) for an AI system?

- A. The accuracy rate of the model
- B. The amount of data in the model
- C. The response time of the model
- D. The rate of drift in the model

Answer: D

NEW QUESTION 216

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

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

Answer: B

NEW QUESTION 217

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 220

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 224

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 229

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 233

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

An organization's CIO provided the AI steering committee with a list of AI technologies in use and tasked them with categorizing the technologies by risk. Which of the following should the committee do FIRST?

- A. Begin grouping similar AI products and solutions together
- B. Identify vulnerabilities related to the technologies in use
- C. Ensure the AI technologies are included in the asset inventory
- D. Assess risk levels based on risk appetite and regulatory requirements

Answer: C

NEW QUESTION 239

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 244

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

- A. Zero Trust
- B. Sandboxing

- C. Siloing
- D. Containerization

Answer: C

NEW QUESTION 245

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 249

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 253

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 257

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 262

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 265

Which of the following would MOST effectively ensure an organization developing AI systems has comprehensive data classification and inventory management?

- A. Creating a centralized team to oversee the classification of data used in AI projects
- B. Conducting quarterly audits of AI data sets for anomalies and missing metadata
- C. Establishing a manual process to categorize data based on business needs and regulatory compliance
- D. Implementing an automated data cataloging tool that integrates with all organizational data repositories

Answer: D

NEW QUESTION 266

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 267

Which of the following BEST enables an organization to strengthen information security controls around the use of generative AI applications?

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

Answer: B

NEW QUESTION 272

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 273

Which of the following is the MAIN objective of the operational phase of AI life cycle management?

- A. Monitor model performance
- B. Align the model to business needs
- C. Optimize the model's algorithms
- D. Obtain end-user feedback on the model

Answer: A

NEW QUESTION 274

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

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

Answer: A

NEW QUESTION 278

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 282

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 287

A vendor switched its chatbot's AI model without due diligence, causing unethical investment advice. What control BEST prevents this scenario?

- A. Master services agreement
- B. Change management
- C. Shared responsibility model
- D. Data minimization

Answer: B

NEW QUESTION 292

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 294

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 296

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 300

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions 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

100% Actual & Verified — Instant Download, Please Click
[Order The AAISM Practice Test Here](#)