

AAISM Dumps

ISACA Advanced in AI Security Management (AAISM) Exam

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



NEW QUESTION 1

Which testing technique is BEST for determining how an AI model makes decisions?

- A. Red team
- B. Black box
- C. White box
- D. Blue team

Answer: C

NEW QUESTION 2

Which 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

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 5

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 6

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 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

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

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

Answer: C

NEW QUESTION 9

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

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

Answer: B

NEW QUESTION 10

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

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 13

A data scientist creating categories and training the algorithm on large data sets is an example of which type of AI model learning technique?

- A. Reinforcement
- B. Unsupervised
- C. Machine learning (ML)
- D. Supervised

Answer: D

NEW QUESTION 18

An organization is designing an AI-based credit risk assessment system integrating sensitive financial data. Which option BEST supports security-by-design?

- A. Integrating differential privacy mechanisms into model training
- B. Applying threat modeling specific to AI components before deployment
- C. Segmenting AI services across containers
- D. Restricting access to AI models using IP allow lists

Answer: B

NEW QUESTION 21

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 23

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 28

Which of the following is BEST for analyzing true positives, true negatives, false positives, and false negatives produced by an AI model?

- A. Hyperparameter tuning
- B. Precision
- C. Confusion matrix
- D. Recall

Answer: C

NEW QUESTION 31

An organization is implementing AI agent development across multiple engineering teams. Which of the following is the MOST important focus of AI-specific security training for developers?

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

Answer: A

NEW QUESTION 36

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 40

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 45

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 49

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 53

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 55

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 58

In the context of generative AI, which of the following would be the MOST likely goal of penetration testing during a red-teaming exercise?

- A. Generate outputs that are unexpected using adversarial inputs
- B. Stress test the model's decision-making process
- C. Degrade the model's performance for existing use cases
- D. Replace the model's outputs with entirely random content

Answer: A

NEW QUESTION 63

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 65

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 69

Which of the following is the BEST way to reduce the risk of misuse of an AI agent that has access to critical data and systems?

- A. Validate agent compliance with output restrictions
- B. Allow users to configure the agent for productivity
- C. Prohibit users from manipulating agent behavior
- D. Limit human review of AI decisions

Answer: A

NEW QUESTION 71

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 73

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 78

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 82

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 87

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

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 94

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 98

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 99

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 102

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 107

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 112

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 117

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 121

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 125

An organization concerned about the ethical and responsible use of a newly developed AI product should consider implementing:

- A. Model cards
- B. Vendor monitoring
- C. An accountability model
- D. Security by design

Answer: C

NEW QUESTION 128

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 129

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 132

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 133

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 134

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 139

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 141

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 145

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 149

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 151

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 152

Which AI data management technique involves creating validation and test data?

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

Answer: B

NEW QUESTION 153

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 156

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 158

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 163

Which of the following BEST describes the role of transparency in AI?

- A. Talking through a decision tree to better understand how the algorithm made each of its choices
- B. Publishing AI mechanisms, data sources, and decision-making processes while making them openly available
- C. Explaining the AI system in an understandable and logical way so reasons for decisions can be given
- D. Persuading someone that the AI tool in use is beneficial and operates as expected

Answer: C

NEW QUESTION 167

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 168

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 169

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 172

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 177

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 178

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

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

Answer: C

NEW QUESTION 182

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 186

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 188

Which area of intellectual property law presents the GREATEST challenge in determining copyright protection for AI-generated content?

- A. Enforcing trademark rights associated with AI systems
- B. Determining the rightful ownership of AI-generated creations
- C. Protecting trade secrets in AI technologies
- D. Establishing licensing frameworks for AI-generated works

Answer: B

NEW QUESTION 191

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 194

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 197

An organization has implemented a natural language processing model to respond to customer questions when personnel are not available. A pre-implementation security assessment revealed attackers could access sensitive company data through a chat interface injection attack. Which of the following is the BEST way to prevent this attack?

- A. Ensuring continuous monitoring and data tagging
- B. Manually reviewing AI model outputs
- C. Implementing input validation and templates
- D. Conducting regular information security audits

Answer: C

NEW QUESTION 200

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 203

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 204

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 206

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 209

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 212

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 214

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 217

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 218

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 219

Which of the following BEST describes the role of risk documentation in an AI governance program?

- A. Providing a record of past AI-related incidents for audits
- B. Outlining the acceptable levels of risk for AI-related initiatives
- C. Offering detailed analyses of technical risk and vulnerabilities
- D. Demonstrating governance, risk, and compliance (GRC) for external stakeholders

Answer: B

NEW QUESTION 222

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 227

Which of the following reviews **MUST** be conducted as part of an AI impact assessment?

- A. Testing, evaluation, validation, and verification
- B. Evaluation of model reproducibility
- C. Security control self-assessment (CSA)
- D. Identification of environmental and societal consequences

Answer: D

NEW QUESTION 230

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 233

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 238

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 240

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 244

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 248

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 253

During red-team testing of an AI system used to make lending decisions, which of the following techniques BEST simulates a data poisoning attack?

- A. Inputting encrypted data into the model
- B. Adding noise to output predictions
- C. Stealing model weights from a deployed API
- D. Corrupting training data sets to manipulate outcomes

Answer: D

NEW QUESTION 256

An organization 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 260

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 261

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 265

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 266

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 268

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 269

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 274

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

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

Answer: D

NEW QUESTION 279

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 283

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

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

Answer: D

NEW QUESTION 288

Which of the following methods provides the MOST effective protection against model inversion attacks?

- A. Using adversarial training
- B. Reducing the model's complexity
- C. Implementing regularization output
- D. Increasing the number of training iterations

Answer: C

NEW QUESTION 293

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 297

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 298

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 301

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 302

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

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

Answer: C

NEW QUESTION 306

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 308

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 313

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 316

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 319

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 322

Which of the following controls BEST mitigates the inherent limitations of generative AI models?

- A. Ensuring human oversight
- B. Adopting AI-specific regulations
- C. Classifying and labeling AI systems
- D. Reverse engineering the models

Answer: A

NEW QUESTION 323

.....

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>