



Isaca

Exam Questions AAISM

ISACA Advanced in AI Security Management (AAISM) Exam

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

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 2

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 3

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 4

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 5

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

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

Answer: C

NEW QUESTION 6

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 7

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 8

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 9

What is the PRIMARY purpose of a dedicated AI management system policy?

- A. Minimizing environmental impact
- B. Optimizing AI model accuracy
- C. Complying with external regulations
- D. Providing a framework to set AI objectives

Answer: D

NEW QUESTION 10

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 10

When addressing privacy concerns related to AI, what is the GREATEST significance of user consent?

- A. It prevents unauthorized access to data
- B. It enables deletion/modification of personal data
- C. It allows the organization to process user data in the AI system
- D. It helps detect bias and ensure fairness

Answer: C

NEW QUESTION 12

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 15

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 17

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 20

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 22

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 26

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 30

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 34

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 36

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 40

The PRIMARY reason to conduct a privacy impact assessment (PIA) on an AI system is to:

- A. Identify applicable regulations
- B. Determine whether personal data is poisoned
- C. Build customer confidence
- D. Analyze how personal data is handled

Answer: D

NEW QUESTION 44

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 47

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 49

When deriving statistical information generated by AI systems, which of the following types of risk is MOST important to address?

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

Answer: A

NEW QUESTION 53

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 55

What BEST ensures a proper business continuity plan (BCP) for an AI solution?

- A. Enhancing monitoring for model failure
- B. Testing AI infrastructure failover mechanisms
- C. Implementing access controls
- D. Increasing backup restoration detail

Answer: B

NEW QUESTION 57

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 60

AI developers often find deep learning systems difficult to explain PRIMARILY because:

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

Answer: B

NEW QUESTION 65

As organizations increasingly rely on vendors to develop AI systems, which of the following is the MOST effective way to monitor vendors and ensure compliance with ethical and security standards?

- A. Conducting regular audits of vendor processes and adherence to AI development guidelines
- B. Requiring vendors to monitor their adherence to ethics and security standards
- C. Mandating that vendors share source code and AI documentation with the contracting party
- D. Allowing vendors to self-attest ethical AI compliance and implement benchmark monitoring

Answer: A

NEW QUESTION 69

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 74

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 77

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 78

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 80

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 81

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 82

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 86

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 89

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 94

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 96

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 100

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 103

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 105

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 109

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 114

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 117

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 120

During the deployment of a generative AI platform, a risk assessment highlighted threats such as data leakage and prompt manipulation. Which of the following is the BEST way to ensure appropriate control selection?

- A. Rely primarily on vendor-provided security features and seek third-party certifications
- B. Map identified AI threats to enterprise control catalogs and integrate AI-specific safeguards where gaps exist
- C. Apply AI-specific controls from external frameworks without customization and initiate monitoring to expedite compliance
- D. Postpone control selection until deployment and address risk through enhanced monitoring

Answer: B

NEW QUESTION 121

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 125

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 130

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 131

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 134

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 137

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 140

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 145

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 147

Which of the following BEST strengthens information security controls around the use of generative AI applications?

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

Answer: B

NEW QUESTION 151

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 152

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 154

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 155

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 160

Which of the following is the GREATEST concern when a vendor enables generative AI features for an organization's critical system?

- A. Access to the model
- B. Proposed regulatory enhancements
- C. Security monitoring and alerting
- D. Bias and ethical practices

Answer: A

NEW QUESTION 164

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 168

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 169

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 170

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 173

Secure aggregation enhances the security of federated learning systems by:

- A. Processing client updates in isolation to reduce the risk of exposing sensitive information
- B. Applying differential privacy techniques to mask sensitive information in training data
- C. Encrypting individual model updates during transmission to ensure only the server can access the data
- D. Ensuring individual client contributions remain confidential even if the server is compromised

Answer: D

NEW QUESTION 176

Which AI model is BEST suited to ensure explainability in an HR department's pre-screening tool for candidate resumes?

- A. Support vector machine
- B. Neural network
- C. Decision tree
- D. Gradient boosting machine

Answer: C

NEW QUESTION 180

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 181

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 183

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 186

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 187

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 188

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 189

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 193

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 197

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 202

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 205

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 210

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 214

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 219

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 223

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 227

Which of the following controls BEST mitigates the risk of data poisoning?

- A. Data set restoration
- B. Data validation
- C. Digital watermarking
- D. Intrusion detection

Answer: B

NEW QUESTION 228

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 232

During the creation of a new large language model (LLM), an organization procured training data from multiple sources. Which of the following is MOST likely to address the CISO's security and privacy concerns?

- A. Data augmentation
- B. Data minimization
- C. Data classification
- D. Data discovery

Answer: B

NEW QUESTION 233

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 238

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 241

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 244

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 246

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 249

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 254

A security assessment revealed that attackers could access sensitive company data through chat interface injection. What is the BEST mitigation?

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

Answer: C

NEW QUESTION 257

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 262

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 263

.....

Relate Links

100% Pass Your AAISM Exam with ExamBible Prep Materials

<https://www.exambible.com/AAISM-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>