

# 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

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 3

Which of the following is the BEST mitigation control for membership inference attacks on AI systems?

- A. Model ensemble techniques
- B. AI threat modeling
- C. Differential privacy
- D. Cybersecurity-oriented red teaming

**Answer: C**

#### NEW QUESTION 4

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 5

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 6

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 7

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 8

From a risk perspective, which of the following is the MOST important step when implementing an adoption strategy for AI systems?

- A. Benchmarking against peer organizations?? AI risk strategies
- B. Implementing a robust risk analysis methodology tailored to AI-specific tasks
- C. Conducting an AI risk assessment and updating the enterprise risk register
- D. Establishing a comprehensive AI risk assessment framework

**Answer: C**

#### NEW QUESTION 9

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 10

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 10

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 11

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 16

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 20

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 24

An organization is facing a deepfake attack intended to manipulate stock prices. The organization's crisis communication plan has been activated. Which of the following is MOST important to include in the initial response?

- A. Conduct employee awareness training on recognizing deepfake videos and audio
- B. Provide clarifying information in a pre-approved public statement
- C. Conduct a detailed forensic analysis to identify the source of the deepfake
- D. Engage with brand monitoring services to track social media activity

**Answer: B**

#### **NEW QUESTION 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 29**

What BEST protects trade secrets related to AI technologies during their life cycle?

- A. Enforcing trademark rights
- B. Restricting access to sensitive data
- C. Patenting AI algorithms and data
- D. Watermarking AI output

**Answer: B**

#### **NEW QUESTION 31**

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 32**

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 37**

An aerospace manufacturing company that prioritizes accuracy and security has decided to use generative AI to enhance operations. Which of the following large language model (LLM) adoption plans BEST aligns with the company's risk appetite?

- A. Developing a public LLM to automate critical functions
- B. Purchasing an LLM dataset on the open market
- C. Contracting LLM access from a reputable third-party provider
- D. Developing a private LLM to automate non-critical functions

**Answer: D**

#### **NEW QUESTION 40**

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 43**

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 44**

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 45**

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 46**

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 47**

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 49**

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 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 58**

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 63**

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 68**

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 71**

An organization is deploying an automated AI cybersecurity system. Which strategy MOST effectively minimizes human error and improves security?

- A. Manual monitoring of alerts
- B. Using historical data to train detection software
- C. Utilizing machine learning algorithms to ensure responsible use
- D. Conducting periodic penetration testing

**Answer: B**

**NEW QUESTION 75**

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 77**

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 79**

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 80**

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 84**

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 86**

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 87**

A financial institution plans to deploy an AI system to provide credit risk assessments for loan applications. Which of the following should be given the HIGHEST priority in the system's design to ensure ethical decision-making and prevent bias?

- A. Regularly update the model with new customer data to improve prediction accuracy.
- B. Integrate a mechanism for customers to appeal decisions directly within the system.
- C. Train the system to provide advisory outputs with final decisions made by human experts.
- D. Restrict the model's decision-making criteria to objective financial metrics only.

**Answer: C**

**NEW QUESTION 90**

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 93**

A financial services firm received a regulatory fine after a vendor switched its chatbot's AI model without due diligence, resulting in unethical investment advice to the firm's clients. Which of the following controls should be implemented by the firm to BEST prevent recurrence of this scenario?

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

**Answer: D**

**NEW QUESTION 98**

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 101**

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 102**

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 107**

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

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

**Answer: A**

**NEW QUESTION 108**

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 109**

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

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

**Answer: B**

**NEW QUESTION 114**

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 117**

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 121**

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 122**

An AI system that supports critical processes has deviated from expected performance and is producing biased outcomes. Which of the following is the BEST course of action?

- A. Retrain the model with a new and expanded dataset
- B. Perform a root cause analysis to identify mitigation steps
- C. Conduct audits of the data and the model
- D. Activate the model kill switch

**Answer: B**

**NEW QUESTION 125**

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 127**

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 130**

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 134**

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 136**

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 139**

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 143**

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 147**

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 148**

Which of the following should be done FIRST when developing an acceptable use policy for generative AI?

- A. Determine the scope and intended use of AI
- B. Review AI regulatory requirements
- C. Consult with risk management and legal
- D. Review existing company policies

**Answer: A**

**NEW QUESTION 152**

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 153**

A global organization has experienced multiple incidents of staff copying confidential data into public chatbots and acting on the model outputs. Which of the following is MOST important to reduce short-term risk when launching an AI security awareness initiative?

- A. Blocking access to public large language models (LLMs) at the network perimeter
- B. Requiring employees to complete an annual generic phishing and deepfake awareness module
- C. Delivering role-based and scenario-driven AI security training mapped to policy and job functions
- D. Publishing an AI acceptable use policy and collecting e-signatures of employees

**Answer: C**

**NEW QUESTION 154**

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 156**

A military contractor discovered that its large language model (LLM) is at high risk of being targeted by advanced persistent threat (APT) actors seeking to exploit the model to access confidential information. Which of the following attacks is the HIGHEST priority to protect against?

- A. Model inversion
- B. Data poisoning
- C. Unauthorized tuning
- D. Model distillation

**Answer: A**

**NEW QUESTION 160**

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 162**

Which of the following is the BEST approach for minimizing risk when integrating acceptable use policies for AI foundation models into business operations?

- A. Limit model usage to predefined scenarios specified by the developer
- B. Rely on the developer's enforcement mechanisms
- C. Establish AI model life cycle policy and procedures
- D. Implement responsible development training and awareness

**Answer: C**

**NEW QUESTION 163**

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 165**

A viral video shows a blurry person making claims about a product safety issue. The video has random low-quality sections. This MOST likely represents what threat?

- A. Hallucinations
- B. Model drift
- C. Data poisoning
- D. Deepfake

**Answer: D**

**NEW QUESTION 169**

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 171**

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 175**

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 178**

An organization using an AI model for financial forecasting identifies inaccuracies caused by missing data. Which of the following is the MOST effective data cleaning technique to improve model performance?

- A. Increasing the frequency of model retraining with the existing data set
- B. Applying statistical methods to address missing data and reduce bias
- C. Deleting outlier data points to prevent unusual values impacting the model

D. Tuning model hyperparameters to increase performance and accuracy

**Answer: B**

**NEW QUESTION 179**

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 180**

Which of the following is the MOST important factor to consider when selecting industry frameworks to align organizational AI governance with business objectives?

- A. Risk tolerance
- B. Risk threshold
- C. Risk register
- D. Risk appetite

**Answer: D**

**NEW QUESTION 184**

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 187**

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 190**

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 193**

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 198**

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 200**

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 201**

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 203**

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 207**

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 209**

When evaluating a new AI tool for intrusion prevention, which is MOST important to ensure fit within the existing program architecture?

- A. Ensure automated response orchestration
- B. Prioritize real-time anomaly detection
- C. Confirm tool capabilities align with control objectives
- D. Select a tool that integrates with the SIEM

**Answer: C**

**NEW QUESTION 210**

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 211**

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 214**

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 218**

Which of the following BEST describes an adversarial attack on an AI model?

- A. Attacking underlying hardware
- B. Providing inputs that mislead the model into incorrect predictions
- C. Reverse-engineering the model using social engineering
- D. Conducting denial-of-service attacks on AI APIs

**Answer:** B

**NEW QUESTION 221**

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 226**

Which of the following datasets is used to tune hyperparameters?

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

**Answer:** A

**NEW QUESTION 231**

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 232**

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

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

**Answer:** D

**NEW QUESTION 237**

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 239**

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 242**

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 243**

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 246**

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 250**

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 253**

.....

## **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](#)**