



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 strategy is MOST effective for penetration testers assessing an AI model against membership inference attacks?

- A. Generating synthetic training data
- B. Analyzing AI model confidence scores
- C. Disabling model logging
- D. Measuring accuracy on the test set

Answer: B

NEW QUESTION 2

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 3

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 4

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 5

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 6

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 7

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 8

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 9

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 10

An organization develops and implements an AI-based plug-in for users that summarizes their individual emails. Which of the following is the GREATEST risk associated with this application?

- A. Lack of application vulnerability scanning
- B. Data format incompatibility
- C. Insufficient rate limiting for APIs
- D. Inadequate controls over parameters

Answer: D

NEW QUESTION 10

Which of the following MOST effectively addresses bias in generative AI models?

- A. Data minimization
- B. Data augmentation
- C. Adversarial training
- D. Fairness constraints

Answer: D

NEW QUESTION 15

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 20

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 23

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 26

Which of the following is the MOST effective way to mitigate the risk of deepfake attacks?

- A. Relying on human judgment for oversight
- B. Limiting employee access to AI tools

- C. Validating the provenance of the data source
- D. Using a general-purpose large language model (LLM) to detect fraud

Answer: C

NEW QUESTION 30

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 33

The PRIMARY ethical concern of generative AI is that it may:

- A. Produce unexpected data that could lead to bias
- B. Cause information integrity issues
- C. Cause information to become unavailable
- D. Breach the confidentiality of information

Answer: B

NEW QUESTION 37

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 42

An organization decides to use an anomaly-based intrusion detection system (IDS) integrated with a generative adversarial network-enabled AI tool. The integrated tool would MOST effectively detect intrusions by leveraging:

- A. synthetic intrusion data to train the tool's components
- B. validation data sets to enable highly realistic AI decisions
- C. automated rule creation to increase model performance
- D. classified real intrusion data based on labeled data

Answer: A

NEW QUESTION 43

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 48

A regulator warns of increased risk of AI re-identification attacks on anonymized datasets. What should the information security manager do FIRST?

- A. Assume anonymization is permanent and continue operations
- B. Immediately delete anonymized datasets and suspend AI services
- C. Implement a monitoring program including privacy audits and adversarial testing
- D. Establish strong access controls for services using anonymized data

Answer: C

NEW QUESTION 52

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 55

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 57

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 61

Which of the following is the MOST important course of action when implementing continuous monitoring and reporting for AI-based systems?

- A. Establish an automated alert system for threshold breaches in risk metrics
- B. Develop standardized risk reporting templates for different stakeholder groups
- C. Implement real-time monitoring of key risk indicators (KRIs) for AI systems
- D. Implement a risk dashboard for visualizing and tracking AI-related risk over time

Answer: C

NEW QUESTION 62

Who is responsible for implementing recommendations in a final report after an external AI compliance audit?

- A. System architects
- B. Internal auditors
- C. End users
- D. Model owners

Answer: D

NEW QUESTION 65

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 69

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 73

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 77

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 79

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 80

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 81

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 84

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 88

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 91

Which approach should an organization prioritize to effectively verify the security of its AI models?

- A. Automating vulnerability identification
- B. Developing a testing strategy including AI-specific threat modeling and adversarial attack simulations
- C. Testing team competencies in IT threat mitigation
- D. Using standard penetration testing methods

Answer: B

NEW QUESTION 92

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 93

An AI research team is developing a natural language processing model that relies on several open-source libraries. Which of the following is the team's BEST course of action to ensure the integrity of the software packages used?

- A. Maintain a list of frequently used libraries to ensure consistent application in projects
- B. Scan the packages and libraries for malware prior to installation
- C. Use the latest version of all libraries from public repositories
- D. Retrain the model regularly to handle package and library updates

Answer: B

NEW QUESTION 97

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 101

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 102

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

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

Answer: D

NEW QUESTION 107

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 110

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 112

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 115

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 117

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 120

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 124

An organization is implementing an AI-based credit assessment engine using internal and third-party customer data. Which of the following BEST aligns with data management controls for the AI life cycle?

- A. Documented procedures for data sourcing, lineage tracking, and quality validation
- B. Use of hashed identifiers to anonymize datasets used for model validation and internal analytics
- C. Encrypted isolation and dynamic access controls on training data pipelines
- D. Limitation of model training to structured data from vetted sources to minimize ingestion risk

Answer: A

NEW QUESTION 128

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 131

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 133

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 138

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 141

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 144

A preliminary risk assessment of a SaaS-based large language model (LLM) business support system has identified prompt injection, data poisoning, and model exfiltration as material threats. Which of the following is the BEST approach to ensure risks are treated consistently?

- A. Implementing an AI threat control matrix that maps threats to specific controls and assurance activities
- B. Applying control baselines from a recognized industry standard to AI components
- C. Relying on vendor independent audit reports and service level agreements (SLAs) as evidence of AI risk coverage
- D. Focusing resources on post-deployment red teaming and deferring control selection until post go-live feedback is received

Answer: A

NEW QUESTION 148

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 151

Which of the following factors is MOST important for preserving user confidence and trust in generative AI systems?

- A. Bias minimization
- B. Access controls and secure storage solutions
- C. Transparent disclosure and informed consent
- D. Data anonymization

Answer: C

NEW QUESTION 155

.....

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/>