Question #1 *Topic 1*

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. Retrain the model regularly to handle package and library updates.

C. Scan the packages and libraries for malware prior to installation.

D. Use the latest version of all libraries from public repositories.

Suggested Answer: *C*

Community vote distribution

C (100%)

sachinmaverick 2 weeks, 4 days ago

Selected Answer: C

Scanning will also take into consideration the SBOM.

upvoted 1 times

Question #2 *Topic 1*

An organization plans to apply an AI system to its business, but developers find it difficult to predict system results due to lack of visibility to the inner workings of the AI model. Which of the following is the GREATEST challenge associated with this situation?

A. Assigning a risk owner who is responsible for system uptime and performance

B. Continuing operations to meet expected AI security requirements

C. Determining average turnaround time for AI transaction completion

D. Gaining the trust of end users through explainability and transparency

Suggested Answer: *D*

Community vote distribution

D (100%)

sachinmaverick 2 weeks, 4 days ago

Selected Answer: D

Agree with D

upvoted 1 times

zufyozirke 2 months ago

Selected Answer: D

Agree with D

upvoted 1 times

zufyozirke 2 months ago

Selected Answer: D

Agree with D

upvoted 1 times

Question #3 *Topic 1*

Which of the following is MOST important to consider when validating a third-party AI tool?

A. Terms and conditions

B. Roundtable testing

C. Right to audit

D. Industry analysis and certifications

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #4 *Topic 1*

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.

Suggested Answer: *C*

vaavoom 3 weeks, 6 days ago

Selected Answer: C

Reasoning: In ISACA risk management frameworks, the ultimate justification for any security or compliance decision is based on risk appetite. If an organization performs a formal risk assessment and determines that the potential legal, financial, or reputational impact of noncompliance falls within its pre-defined acceptable risk levels, that serves as the primary justification for the decision. Other options (like industry use or lack of cost determination) are not formal risk management justifications.

upvoted 1 times

Question #5 *Topic 1*

Which of the following is the MOST important consideration when deciding how to compose an AI red team?

A. Resource availability

B. Time-to-market constraints

C. Skills matrix

D. AI use cases

Suggested Answer: *C*

vaavoom 3 weeks, 6 days ago

Selected Answer: C

Reasoning: AI red teaming is a highly specialized field that requires a unique combination of skills, including expertise in adversarial machine learning, prompt injection, data poisoning, and bias testing. A skills matrix is the most critical tool for composing a team because it ensures that all necessary technical and domain-specific competencies are represented. While use cases (Option D) guide the scope of the testing, the composition of the team itself must prioritize the specialized talent required to execute those tests effectively.

upvoted 1 times

Question #6 *Topic 1*

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. Ensure the AI technologies are included in the asset inventory.

C. Assess risk levels based on risk appetite and regulatory requirements.

D. Identify vulnerabilities related to the technologies in use.

Suggested Answer: *B*

Community vote distribution

alanlah 2 months, 1 week ago

Selected Answer: B

Before an AI steering committee can categorize AI technologies by risk, they must first ensure the list of technologies is complete and properly captured in the organization's asset inventory.

In governance and risk management best practices (AAISM, NIST AI RMF, ISO/IEC 42001), the first step is always establishing visibility.

upvoted 2 times

Question #7 *Topic 1*

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. Data asset validation

C. Human-in-the-loop

D. AI impact analysis

Suggested Answer: *C*

Community vote distribution

C (100%)

alanlah 2 months, 1 week ago

Selected Answer: C

When deploying an AI solution to generate treatment regimens—a high-risk, safety-critical use case—the most effective way to reduce the risk of hallucinations is to ensure that a qualified human expert reviews, approves, or overrides AI outputs before they influence patient care.

A human-in-the-loop (HITL) mitigates hallucination risks by:

Catching incorrect, implausible, or clinically unsafe recommendations

Ensuring AI outputs align with medical standards, guidelines, and clinical judgment

Providing accountability and oversight, especially when real-world data introduces noise, bias, or edge cases

This is strongly aligned with FDA, EMA, and ISO/IEC 42001 expectations for high-risk AI.

upvoted 1 times

Question #8 *Topic 1*

Which of the following should be the PRIMARY consideration for an organization concerned about liabilities associated with unforeseen behavior from agentic AI systems?

A. Model dependencies

B. Approved base models

C. Acceptable risk level

D. Accountability model

Suggested Answer: *D*

Community vote distribution

D (100%)

alanlah 2 months, 1 week ago

Selected Answer: D

For organizations concerned about liability from unforeseen or autonomous behavior in agentic AI systems, the PRIMARY consideration is having a clear accountability model that defines:

Who is responsible for decisions made by the AI

Who approves and oversees agentic behaviors

Escalation paths when the system acts outside expectations

Legal and compliance ownership

Roles for monitoring, intervention, and remediation

Agentic AI amplifies risks because systems may act autonomously, pursue goals in unintended ways, or take actions without direct prompts. Regulators (e.g., EU AI Act, ISO/IEC 42001, NIST AI RMF) emphasize that accountability and governance are the foundation of managing liability in autonomous AI systems.

upvoted 1 times

Question #9 *Topic 1*

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 minimization

B. Data augmentation

C. Data classification

D. Data discovery

Suggested Answer: *A*

Community vote distribution

C (100%)

alanlah 2 months, 1 week ago

Selected Answer: C

When an organization acquires training data from multiple sources to build an LLM, the CISO's primary concerns will be:
- Sensitivity of each dataset
- Presence of personal data (PII/PHI)
- Required security controls
- Compliance obligations (e.g., GDPR, HIPAA, PDPA)
- How the data should be handled, stored, protected, and restricted

Data classification directly addresses these concerns by labeling datasets based on:
- Sensitivity (public / internal / confidential / restricted)
- Presence of personal or regulated data
- Security handling requirements
- Retention and access control needs

This enables the CISO to apply the appropriate controls before the data enters the training pipeline.

upvoted 4 times

Question #10 *Topic 1*

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. Model validation

B. Control self-assessment (CSA)

C. Explainable decision-making

D. Key performance indicator (KPI)

Suggested Answer: *A*

Community vote distribution

sachinmaverick 2 weeks, 4 days ago

Selected Answer: D

Measuring the error rate (incorrect predictions vs. total predictions) is a standard way to evaluate the operational effectiveness and business value of a system over time.

upvoted 1 times

alanlah 2 months, 1 week ago

Selected Answer: D

Analyzing incorrect predictions vs. total predictions is essentially measuring:

Error rate, or

Model accuracy

These are classic model performance KPIs used to evaluate whether an AI/ML model is still effective and worth keeping in production.

This aligns with operational monitoring and lifecycle management—not validation, explainability, or CSA.

upvoted 2 times

Question #11 *Topic 1*

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

A. The amount of data in the model

B. The rate of drift in the model

C. The accuracy rate of the model

D. The response time of the model

Suggested Answer: *B*

Community vote distribution

B (100%)

sachinmaverick 2 weeks, 4 days ago

Selected Answer: B

KRI is forward looking and KPI is backward looking.

upvoted 1 times

alanlah 2 months, 1 week ago

Selected Answer: B

This is the most critical Key Risk Indicator (KRI) because it directly measures the degradation of model performance over time due to changes in real-world data or relationships. Model drift signals that the AI system may be producing increasingly unreliable or biased outputs, which can lead to operational failures, financial losses, or harm. Monitoring drift enables proactive interventions (e.g., retraining) to mitigate these risks before they materialize.

upvoted 1 times

Question #12 *Topic 1*

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. Implement resilient payment approval processes.

C. Mandate that payments be sent only once per week.

D. Issue a security policy on deepfakes.

Suggested Answer: *B*

Community vote distribution

B (100%)

alanlah 2 months, 1 week ago

Selected Answer: B

Deepfake-enabled payment diversion attacks succeed not because employees lack awareness, but because the payment process is vulnerable to social engineering and impersonation.

The strongest protection is robust, multi-factor, out-of-band approval workflows, such as:

Dual-control or multi-party payment authorization

Independent verification for unusual or urgent requests

Non-spoofable verification channels (e.g., call-back to a known number)

Threshold-based approvals and segregation of duties

Strict exception handling governance

These controls prevent fraudulent payments even if a deepfake is convincing.

This aligns with leading practices from NIST, SWIFT CSP, and financial fraud prevention standards.

upvoted 1 times

Question #13 *Topic 1*

Which of the following technologies can be used to manage deepfake risk?

A. Adaptive authentication

B. Multi-factor authentication (MFA)

C. Systematic data tagging

D. Blockchain

Suggested Answer: *D*

wmoler1972 2 days ago

Selected Answer: C

Deepfake risk =content provenance-systematic data tagging

Were Blockchain only really supports provenance

upvoted 1 times

Question #14 *Topic 1*

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. Engage with brand monitoring services to track social media activity.

B. Conduct a detailed forensic analysis to identify the source of the deepfake.

C. Provide clarifying information in a pre-approved public statement.

D. Conduct employee awareness training on recognizing deepfake videos and audio.

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #15 *Topic 1*

Which of the following BEST reduces the risk of exposing sensitive data through the output of large language models (LLMs) in applications?

A. Enforcing least privilege access

B. Conducting adversarial testing

C. Encrypting data in transit and at rest

D. Implementing data sanitization techniques

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Question #16 *Topic 1*

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. Fine-tuning the AI model to decrease hallucinations

C. Monitoring the system for misuse cases

D. Implementing a secondary AI system to confirm images

Suggested Answer: *A*

vaavoom 3 weeks, 6 days ago

Selected Answer: A

Alterations in facial appearance refer to techniques like disguises, makeup, or adversarial patterns designed to evade or spoof the system (evasion attacks). The most effective preventive measure is improving the model's robustness during training by exposing it to a wide variety of facial variations (e.g., different lighting, angles, accessories, ethnicities, and simulated alterations). This data augmentation increases variance, making the system less susceptible to such compromises.

upvoted 1 times

Question #17 *Topic 1*

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

A. model cards.

B. security by design.

C. vendor monitoring.

D. an accountability model.

Suggested Answer: *A*

vaavoom 3 weeks, 6 days ago

Selected Answer: A

Model cards are standardized documentation frameworks (popularized by responsible AI practices) that detail a model's intended use, limitations, performance metrics, potential biases, ethical considerations, and risks. They promote transparency and accountability, directly supporting ethical and responsible deployment of in-house developed AI products.

D is vague and not a standard term in responsible AI governance compared to model cards.

upvoted 1 times

Question #18 *Topic 1*

Which of the following metrics BEST evaluates the ability of a model to correctly identify all true positive instances?

A. F1 score

B. Specificity

C. Precision

D. Recall

Suggested Answer: *D*

vaavoom 3 weeks, 6 days ago

Selected Answer: D

Recall (also known as sensitivity) measures the ability of a model to find all the positive samples.

upvoted 1 times

Question #19 *Topic 1*

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.

Suggested Answer: *D*

vaavoom 3 weeks, 6 days ago

Selected Answer: D

A Privacy Impact Assessment (PIA), also known as a Data Protection Impact Assessment (DPIA) in some frameworks, is fundamentally a systematic process to evaluate and analyze how personal data is collected, processed, stored, shared, and protected throughout the AI system's lifecycle. Its primary goal is to identify privacy risks arising from personal data handling and determine appropriate safeguards or mitigations.

upvoted 1 times

Question #20 *Topic 1*

Which of the following will BEST reduce data bias in machine learning (ML) algorithms?

A. Utilizing unstructured data sets

B. Adopting a more simplified model

C. Diversifying the model training data

D. Securing the model training data

Suggested Answer: *C*

vaavoom 3 weeks, 5 days ago

Selected Answer: C

Data bias in ML primarily arises from unrepresentative, incomplete, or skewed training datasets (e.g., underrepresenting certain groups, demographics, or scenarios), leading to unfair or inaccurate model outputs. The most direct and effective way to mitigate this is by diversifying the training data — collecting and using representative, inclusive, and varied sources to ensure balanced coverage across relevant populations, features, and conditions. This is a foundational pre-processing step emphasized in responsible AI practices and aligns with ISACA's AI governance guidance on addressing bias through high-quality, diverse data.

upvoted 1 times

Question #21 *Topic 1*

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

A. Consult with risk management and legal.

B. Review AI regulatory requirements.

C. Determine the scope and intended use of AI.

D. Review existing company policies.

Suggested Answer: *C*

vaavoom 3 weeks, 5 days ago

Selected Answer: C

Developing an effective acceptable use policy (AUP) for generative AI starts with clearly defining the scope (what tools, use cases, users, and activities are covered) and intended uses (permissible vs. prohibited scenarios). This foundational step ensures the policy is tailored to the organization's specific AI applications, risks, and objectives before consulting experts, reviewing regulations, or examining existing policies. It provides the context needed for everything else to align properly.

upvoted 1 times

Question #22 *Topic 1*

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

A. Using open-source data repositories

B. Reviewing AI model cards

C. Performing AI data augmentation

D. Incorporating data from search content

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #23 *Topic 1*

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.

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #24 *Topic 1*

Which of the following is MOST important for an organization to consider when implementing a preventive security safeguard into a new AI product?

A. Penetration testing

B. Input sanitization

C. Model output monitoring

D. Differential privacy

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #25 *Topic 1*

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. Mandating that vendors share source code and AI documentation with the contracting party

B. Requiring vendors to monitor their adherence to ethics and security standards

C. Conducting regular audits of vendor processes and adherence to AI development guidelines

D. Allowing vendors to self-attest ethical AI compliance and implement benchmark monitoring

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #26 *Topic 1*

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. Data poisoning

B. Evasion attack

C. Privilege escalation

D. Model inversion

Suggested Answer: *A*

vaavoom 3 weeks ago

Selected Answer: A

the attack persists, not just one time prompt

upvoted 1 times

Question #27 *Topic 1*

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. Protecting trade secrets in AI technologies

C. Determining the rightful ownership of AI-generated creations

D. Establishing licensing frameworks for AI-generated works

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #28 *Topic 1*

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. Restrict the model's decision-making criteria to objective financial metrics only.

C. Train the system to provide advisory results with final decisions made by human experts.

D. Integrate a mechanism for customers to appeal decisions directly within the system.

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #29 *Topic 1*

Which of the following security framework elements BEST helps to safeguard the integrity of outputs generated by AI algorithms?

A. Management is prepared to disclose AI system architecture to stakeholders.

B. Ethical standards are incorporated into security awareness programs.

C. Risk exposure due to bias in AI outputs is kept within an acceptable range.

D. Responsibility is defined for legal actions related to AI regulatory requirements.

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #30 *Topic 1*

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

A. AI threat modeling

B. Differential privacy

C. Cybersecurity-oriented red teaming

D. Model ensemble techniques

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #31 *Topic 1*

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

A. Establishing a comprehensive AI risk assessment framework

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. Benchmarking against peer organizations' AI risk strategies

Suggested Answer: *C*

vaavoom 3 weeks ago

Selected Answer: A

im going with A, its the most left-shifted answer. Cant have the others without A.

upvoted 1 times

Question #32 *Topic 1*

Which of the following is MOST important to monitor in order to ensure the effectiveness of an organization's AI vendor management program?

A. Vendor results in compliance training programs

B. Vendor participation in industry AI research

C. Vendor reviews of external AI threat reports

D. Vendor compliance with AI-related requirements

Suggested Answer: *D*

vaavoom 3 weeks ago

Selected Answer: D

agreed

upvoted 1 times

Question #33 *Topic 1*

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. Return to an earlier phase in the AI life cycle.

B. Take the AI model offline.

C. Adjust the hyperparameters of the AI model.

D. Create an emergency change request to correct the issue.

Suggested Answer: *B*

zvmail2 2 weeks, 5 days ago

Selected Answer: A

A - will allow to address root cause

B - not practical in real world

C- - this should not be done in flight

E - may not be effective until root cause is determined

upvoted 1 times

vaavoom 3 weeks ago

Selected Answer: B

if its got that mad cow disease ya gotta per 'er down

upvoted 1 times

Question #34 *Topic 1*

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

A. cause information integrity issues.

B. cause information to become unavailable.

C. breach the confidentiality of information.

D. produce unexpected data that could lead to bias.

Suggested Answer: *A*

vaavoom 3 weeks ago

Selected Answer: D

keyword is ETHICAL. only D relates to that directly. Answer the question in a vacuum, don't extrapolate.

upvoted 1 times

Question #35 *Topic 1*

To ensure AI tools do not jeopardize ethical principles, it is MOST important to validate that:

A. stakeholders have approved alignment with company values.

B. AI tools are evaluated by the privacy department before implementation.

C. outputs of AI tools do not perpetuate adverse biases.

D. the organization has implemented a responsible development policy.

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #36 *Topic 1*

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. Replacing human analysis with automated AI decision-making processes

C. Assigning AI-enabled tools to triage non-critical alerts to preserve SOC resources

D. Using AI-enabled tools exclusively to classify all types of security incidents

Suggested Answer: *A*

q65hrrhh2b 2 months ago

asdasofjo

upvoted 1 times

Question #37 *Topic 1*

When implementing a generative AI system, which of the following approaches will BEST prevent misalignment between the corporate risk appetite and tolerance?

A. Creating and maintaining an AI risk register

B. Establishing and monitoring acceptable levels of AI system risk

C. Performing an AI impact assessment

D. Ensuring effective AI key performance indicators (KPIs)

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #38 *Topic 1*

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

A. Adopting AI-specific regulations

B. Classifying and labeling AI systems

C. Ensuring human oversight

D. Reverse engineering the models

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #39 *Topic 1*

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. Review log information that records how data was collected.

B. Disclose service provider policies to declare compliance with regulations.

C. Activate filtering logic to exclude intellectual property flags.

D. Appoint a data steward specialized in AI to strengthen security governance.

Suggested Answer: *B*

vaavoom 3 weeks, 4 days ago

Selected Answer: C

The primary source of lawsuits against generative AI service providers (e.g., OpenAI, Anthropic, Meta, Stability AI) stems from unauthorized use of copyrighted internet-sourced data for model training, leading to claims of direct infringement. Activating proactive filtering logic (e.g., content moderation, opt-out mechanisms, or exclusion based on known IP/copyright indicators like robots.txt flags, known copyrighted works, or metadata) directly prevents ingestion of high-risk data, reducing exposure at the source. This is a technical, preventive control that aligns with best practices for mitigating IP-related litigation risks in generative AI.

upvoted 1 times

Question #40 *Topic 1*

Which of the following types of testing can MOST effectively mitigate prompt hacking?

A. Adversarial

B. Input

C. Load

D. Regression

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #41 *Topic 1*

An organization recently introduced a generative AI chatbot that can interact with users and answer their queries. Which of the following would BEST mitigate hallucination risk identified by the risk team?

A. Performing model testing and validation

B. Ensuring model developers have been trained in AI risk

C. Fine-tuning the foundational model

D. Training the foundational model on large data sets

Suggested Answer: *A*

vaavoom 3 weeks, 4 days ago

Selected Answer: C

fine tune fix halluc

upvoted 1 times

Question #42 *Topic 1*

An organization plans to implement a new AI system. Which of the following is the MOST important factor in determining the level of risk monitoring activities required?

A. The organization's risk appetite

B. The organization's risk tolerance

C. The organization's number of AI system users

D. The organization's compensating controls

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #43 *Topic 1*

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

A. Malicious insider threats

B. Clean desk policy

C. Authentication controls

D. Social engineering

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Question #44 *Topic 1*

Which of the following BEST ensures the integrity of data sets used to train AI models?

A. Collection and retention of only necessary data sets

B. Tracking and verification of data sets via cryptographic controls

C. Clear documentation of data sources, types used, and processing steps

D. Appropriate storage of data sets according to documented classification processes

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #45 *Topic 1*

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. Right to audit

C. Independent certification

D. Continuous data monitoring

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #46 *Topic 1*

Which of the following controls BEST mitigates the risk of bias in AI models?

A. Regular data reconciliation

B. Diverse data sourcing strategies

C. Robust access control techniques

D. Cryptographic hash functions

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #47 *Topic 1*

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

A. Implementing an automated data cataloging tool that integrates with all organizational data repositories

B. Creating a centralized team to oversee the classification of data used in AI projects

C. Conducting quarterly audits of AI data sets for anomalies and missing metadata

D. Establishing a manual process to categorize data based on business needs and regulatory compliance

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #48 *Topic 1*

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. Applying statistical methods to address missing data and reduce bias

B. Increasing the frequency of model retraining with the existing data set

C. Tuning model hyperparameters to increase performance and accuracy

D. Deleting outlier data points to prevent unusual values impacting the model

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #49 *Topic 1*

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

A. Offering detailed analyses of technical risk and vulnerabilities

B. Demonstrating governance, risk, and compliance (GRC) for external stakeholders

C. Outlining the acceptable levels of risk for AI-related initiatives

D. Providing a record of past AI-related incidents for audits

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #50 *Topic 1*

Which of the following AI system vulnerabilities is MOST easily exploited by adversaries?

A. Weak controls for access to the AI model

B. Lack of protection against denial of service (DoS) attacks

C. Inaccurate generalizations from new data by the AI model

D. Inability to detect input modifications causing inappropriate AI outputs

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #51 *Topic 1*

Which of the following is the MOST important consideration for an organization that has decided to adopt AI to leverage its competitive advantage?

A. Develop a business case for the procurement of AI monitoring tools.

B. Develop a comprehensive risk management process to address AI-related issues.

C. Develop a comprehensive strategic roadmap for AI integration.

D. Develop internal training programs on AI governance, risk, and compliance (GRC).

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #52 *Topic 1*

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 threshold

B. Risk appetite

C. Risk register

D. Risk tolerance

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #53 *Topic 1*

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

A. Rely on the developer's enforcement mechanisms.

B. Implement responsible development training and awareness.

C. Establish AI model life cycle policy and procedures.

D. Limit model usage to predefined scenarios specified by the developer.

Suggested Answer: *C*

vaavoom 3 weeks, 3 days ago

Selected Answer: C

A is wrong, it's C. The most comprehensive and proactive approach per ISACA AAISM guidance (focused on AI Governance and Program Management + Risk Management domains) is establishing formal AI model lifecycle policies and procedures. This covers the full end-to-end process — from selection, evaluation, approval, deployment, monitoring, retraining, and decommissioning — ensuring acceptable use is embedded throughout, risks are continuously assessed, and internal accountability is maintained. This aligns with frameworks like NIST AI RMF and ISO/IEC 42001, which stress lifecycle governance for responsible AI integration.

upvoted 1 times

zufyozirke 2 months ago

Agreed with A

upvoted 1 times

Question #54 *Topic 1*

Which of the following key risk indicators (KRIs) is MOST relevant when evaluating the effectiveness of an organization's AI risk management program?

A. Percentage of critical business systems with AI components

B. Number of AI-related training requests submitted

C. Number of AI models deployed into production

D. Percentage of AI project in compliance

Suggested Answer: *D*

vaavoom 3 weeks, 3 days ago

Selected Answer: D

D (percentage of AI projects in compliance) is the strongest effectiveness indicator because it directly shows whether the risk management framework is achieving its core goal: ensuring AI initiatives meet governance, policy, regulatory, ethical, and security requirements. High compliance rates demonstrate that risks are being systematically addressed throughout project lifecycles. This aligns with AAISM's emphasis on monitoring program outcomes and control effectiveness in the AI Risk Management domain.

upvoted 1 times

Question #55 *Topic 1*

Which of the following information is MOST important to include in a centralized AI inventory?

A. Ownership and accountability of AI systems

B. Foundation model and package registry

C. AI model use cases

D. Training data sets

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #56 *Topic 1*

Personal data used to train AI systems can BEST be protected by:

A. anonymizing personal data.

B. hashing personal data.

C. erasing personal data after training.

D. ensuring the quality of personal data.

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #57 *Topic 1*

Embedding unique identifiers into AI models would BEST help with:

A. preventing unauthorized access.

B. detecting adversarial attacks.

C. tracking ownership.

D. eliminating AI system biases.

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #58 *Topic 1*

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

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #59 *Topic 1*

An organization has requested a developer to apply AI algorithms to existing modules in order to improve customer service quality. At this stage, which of the following should be considered FIRST?

A. IT management may need to revise the service agreement if AI behavior cannot be predefined.

B. The organization may need to explain the performance of the applied AI algorithm.

C. Project sponsors may need to agree on a phased approach in order to ensure safe release.

D. The developer may need to be held accountable for business inquiries raised by customers.

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #60 *Topic 1*

Which of the following BEST describes how supervised learning models help reduce false positives in cybersecurity threat detection?

A. They dynamically generate new labeled data sets.

B. They analyze patterns in data to group legitimate activity from actual threats.

C. They learn from historical labeled data.

D. They use real-time feature engineering to automatically adjust decision boundaries.

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #61 *Topic 1*

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

A. Data validation

B. Intrusion detection

C. Digital watermarking

D. Data set restoration

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #62 *Topic 1*

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

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Question #63 *Topic 1*

An automotive manufacturer uses AI-enabled sensors on machinery to monitor variables such as vibration, temperature, and pressure. Which of the following BEST demonstrates how this approach contributes to operational resilience?

A. Scheduling repairs for critical equipment based on real-time condition monitoring

B. Conducting monthly manual reviews of maintenance schedules

C. Performing regular maintenance based on manufacturer recommendations

D. Automating equipment repairs without any human intervention

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #64 *Topic 1*

When an attacker uses synthetic data to reverse engineer an organization's AI model, it is an example of which of the following types of attack?

A. Prompt

B. Poisoning

C. Distillation

D. Inversion

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #65 *Topic 1*

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. Insufficient rate limiting for APIs

B. Data format incompatibility

C. Lack of application vulnerability scanning

D. Inadequate controls over parameters

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Question #66 *Topic 1*

Which of the following is the GREATEST benefit of implementing an AI tool to safeguard sensitive data and prevent unauthorized access?

A. Timely initiation of incident response

B. Reduced number of false positives

C. Timely analysis of endpoint activities

D. Reduced need for data classification

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #67 *Topic 1*

Which of the following would BEST help mitigate vulnerabilities associated with hidden triggers in generative AI models?

A. Monitoring model outputs and suspicious patterns to detect trigger activations

B. Regularly retraining the model using a diverse data set

C. Applying differential privacy and masking sensitive patterns in the training data

D. Incorporating adversarial training to expose and neutralize potential triggers

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Question #68 *Topic 1*

Which of the following is the MOST important course of action prior to placing an in-house developed AI solution into production?

A. Deploy a prototype of the solution.

B. Perform a privacy, security, and compliance gap analysis.

C. Obtain senior management sign-off.

D. Perform testing, evaluation, validation, and verification.

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Question #69 *Topic 1*

Which of the following is a key risk indicator (KRI) for an AI system used for threat detection?

A. Number of training epochs

B. Number of layers in the neural network

C. Training time of the model

D. Number of system overrides by cyber analysts

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Question #70 *Topic 1*

The PRIMARY benefit of implementing moderation controls in generative AI applications is that it can:

A. filter out harmful or inappropriate content.

B. increase the model's ability to generate diverse and creative content.

C. ensure the generated content adheres to privacy regulations.

D. optimize the model's response time.

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #71 *Topic 1*

Which of the following should be a PRIMARY consideration when defining recovery point objectives (RPOs) and recovery time objectives (RTOs) for generative AI solutions?

A. Prioritizing computational efficiency over data integrity to minimize downtime

B. Maintaining consistent hardware configurations to prevent discrepancies during model restoration

C. Preserving the most recent versions of data models to avoid inaccuracies in functionality

D. Ensuring the backup system can restore training data sets within the defined RTO window

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #72 *Topic 1*

Which of the following is the MOST effective use of AI in incident response?

A. Streamlining incident response testing

B. Automating incident response triage

C. Ensuring chain of custody

D. Improving incident response playbook

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #73 *Topic 1*

When documenting information about machine learning (ML) models, which of the following artifacts BEST helps enhance stakeholder trust?

A. Model card

B. Model prototyping

C. Hyperparameters

D. Data quality controls

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #74 *Topic 1*

Which of the following BEST represents a combination of quantitative and qualitative metrics that can be used to comprehensively evaluate AI transparency?

A. AI explainability reports and bias metrics

B. AI system availability and downtime metrics

C. AI ethical impact and user feedback metrics

D. AI model complexity and accuracy metrics

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #75 *Topic 1*

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

A. Limiting employee access to AI tools

B. Validating the provenance of the data source

C. Relying on human judgment for oversight

D. Using a general-purpose large language model (LLM) to detect fraud

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #76 *Topic 1*

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

A. Data anonymization

B. Bias minimization

C. Transparent disclosure and informed consent

D. Access controls and secure storage solutions

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #77 *Topic 1*

A retail organization implements an AI-driven recommendation system that utilizes customer purchase history. Which of the following is the BEST way for the organization to ensure privacy and comply with regulatory standards?

A. Establishing a governance committee to oversee AI privacy practices

B. Storing customer data indefinitely to ensure the AI model has a complete history

C. Maintaining a register of legal and regulatory requirements for privacy

D. Conducting quarterly retraining of the AI model to maintain the accuracy of recommendations

Suggested Answer: *C*

vaavoom 3 weeks, 2 days ago

Selected Answer: A

tough question between A and C, but i'll go with A as the the more overarching option.

upvoted 1 times

Question #78 *Topic 1*

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. Explaining the overall benefit of the system to stakeholders

B. Establishing clear lines of responsibility for AI model outputs

C. Identifying hallucinations returned by AI models

D. Determining the aggregate risk of the system

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Question #79 *Topic 1*

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. Data poisoning

B. Availability attack

C. Deepfake

D. Model inversion

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #80 *Topic 1*

Which of the following BEST enables an organization to maintain visibility to its AI usage?

A. Measuring the impact of AI implementation using key performance indicators (KPIs)

B. Maintaining a comprehensive inventory of AI systems and business units that leverage them

C. Maintaining a monthly dashboard that captures all AI vendors

D. Ensuring the board approves the policies and standards that define corporate AI strategy

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #81 *Topic 1*

A PRIMARY objective of responsibly providing AI services is to:

A. ensure the confidentiality and integrity of data processed by AI models.

B. build trust for decisions and predictions made by AI models.

C. enable AI models to operate autonomously.

D. improve the ability of AI models to learn from new data.

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #82 *Topic 1*

When integrating AI for innovation, which of the following can BEST help an organization manage security risk?

A. Evaluating compliance requirements

B. Re-evaluating the risk appetite

C. Adopting a phased approach

D. Seeking third-party advice

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #83 *Topic 1*

A model producing contradictory outputs based on highly similar inputs MOST likely indicates the presence of:

A. evasion attacks.

B. poisoning attacks.

C. membership inference.

D. model exfiltration.

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #84 *Topic 1*

An attacker crafts inputs to a large language model (LLM) to exploit output integrity controls. Which of the following types of attacks is this an example of?

A. Evasion

B. Jailbreaking

C. Remote code execution

D. Prompt injection

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Question #85 *Topic 1*

Which of the following MOST effectively minimizes the attack surface when securing AI agent components during their development and deployment?

A. Schedule periodic manual code reviews.

B. Implement compartmentalization with least privilege enforcement.

C. Consolidate event logs for correlation and centralized analysis.

D. Deploy pre-trained models directly into production.

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #86 *Topic 1*

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 exposure of personal information may lead to a decline in public trust.

C. The publicly available output of the model may include false or defamatory statements about individuals.

D. The output may reveal information about individuals or groups without their knowledge.

Suggested Answer: *D*

qt28vjyx2b 2 months ago

asdasofjo

upvoted 1 times

Question #87 *Topic 1*

Which of the following is the GREATEST risk inherent to implementing generative AI?

A. Inadequate return on investment (ROI)

B. Lack of employee training

C. Potential intellectual property violations

D. Unidentified asset vulnerabilities

Suggested Answer: *C*

Currently there are no comments in this discussion, be the first to comment!

Question #88 *Topic 1*

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

A. Implement real-time monitoring of key risk indicators (KRIs) for AI systems.

B. Implement a risk dashboard for visualizing and tracking AI-related risk over time.

C. Develop standardized risk reporting templates for different stakeholder groups.

D. Establish an automated alert system for threshold breaches in risk metrics.

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #89 *Topic 1*

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. Unwillingness of large AI companies to accept updated terms

C. Insufficient legal team experience with AI

D. Inability to sufficiently identify shadow AI within the organization

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #90 *Topic 1*

Which of the following is the BEST reason to immediately disable an AI system?

A. Overly detailed model outputs

B. Excessive model drift

C. Insufficient model training

D. Slow model performance

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!