



experience experts who have worked in this line more than 10 years.

## CompTIA SecAI+ Certification Exam Sample Questions (Q13-Q18):

### NEW QUESTION # 13

An AI security team must assess the probability of an attack on its new system and the impact associated with such an attack. Which of the following threat-modeling resources best addresses the threat landscape for machine learning (ML)?

- A. Massachusetts Institute of Technology (MIT) risk repository
- **B. MITRE Adversarial Threat Landscape for AI Systems (ATLAS)**
- C. Open Worldwide Application Security Project (OWASP)
- D. Common Vulnerabilities and Exposures (CVE) AI working group

**Answer: B**

Explanation:

MITRE ATLAS is specifically designed to capture adversarial tactics, techniques, and procedures (TTPs) targeting machine learning systems. It helps organizations assess both the probability and impact of AI/ML-related attacks, making it the most relevant threat-modeling resource.

### NEW QUESTION # 14

A security consultant must summarize the impact of posture management on a machine learning (ML) use case. Which of the following is the most appropriate reference for this purpose?

- A. Generative adversarial network (GAN)
- **B. National Institute of Standards and Technology (NIST) AI Risk Management Framework (RMF)**
- C. Organization for Economic Co-operation and Development (OECD) standards
- D. European Union AI Act

**Answer: B**

Explanation:

Basic Concept: Security posture management for AI systems involves assessing and improving the overall security state of AI deployments, including identifying risks, implementing controls, and maintaining ongoing compliance. Appropriate frameworks provide structure for this assessment. CompTIA SecAI+ Study Guide identifies NIST AI RMF as the primary framework for AI risk and posture management.

Why B is Correct: The NIST AI Risk Management Framework provides comprehensive, actionable guidance for managing and improving AI security and risk posture across the entire AI lifecycle. It includes the GOVERN, MAP, MEASURE, and MANAGE functions that directly address posture management activities including risk identification, assessment, and control implementation for ML use cases. Its technical depth and ML-specific guidance make it ideal for this summarization task.

Why A is Wrong: OECD standards provide high-level policy principles for AI governance at an international level. They lack the technical specificity and operational guidance needed to summarize posture management impact on a specific ML use case.

Why C is Wrong: The EU AI Act is a regulatory compliance framework establishing legal requirements for AI systems. While it addresses risk management, its focus is on legal compliance rather than technical posture management guidance for ML systems.

Why D is Wrong: A Generative Adversarial Network is an AI architecture for generating synthetic data, not a framework or standard. It has no relevance as a reference for AI security posture management.

### NEW QUESTION # 15

A manufacturing company wants to use AI within its operations to improve the efficiency and accuracy of its processes. Which of the following should the organization do first to enable adoption and achieve the business objectives?

- A. Select a large language model (LLM).
- B. Achieve International Organization for Standardization (ISO) 42001 certification.
- C. Introduce a generative adversarial network (GAN).
- **D. Hire a data and AI architect.**

**Answer: D**

Explanation:

The first step in adopting AI to meet business objectives is to establish the right expertise. A data and AI architect can design the overall strategy, infrastructure, and data pipelines needed for effective AI integration, ensuring alignment with operational goals before selecting specific models or certifications.

#### NEW QUESTION # 16

An organization recently created a custom model that integrates with a language model (LLM). The developer notices that the application programming interface (API) costs have increased.

Which of the following is the best control to reduce cost?

- A. Adjusting token limits
- B. Increasing central processing unit (CPU) and memory
- C. Reducing the model size
- D. Implementing prompt templates

**Answer: A**

Explanation:

Basic Concept: LLM API pricing is primarily based on token consumption - the number of tokens processed in both input prompts and output responses. Controlling token usage is the most direct lever for managing and reducing LLM API costs. CompTIA SecAI+ Study Guide covers AI cost management and resource controls under securing AI systems.

Why D is Correct: Adjusting token limits directly caps the maximum number of tokens used per request for both input and output.

By setting appropriate token limits, the organization prevents excessively long prompts or verbose responses from consuming unnecessary tokens, directly translating to lower API costs and providing hard budget control.

Why A is Wrong: Prompt templates standardize how queries are structured, which can indirectly improve efficiency. However, they do not enforce a hard cap on token usage and cannot prevent costs from escalating with large volumes or verbose responses.

Why B is Wrong: Increasing CPU and memory addresses computational infrastructure performance on the client side. LLM API costs are billed by the API provider based on token usage, not on the client's hardware resources.

Why C is Wrong: Reducing model size means using a smaller, less powerful model version. While this may lower cost per token, it is a model selection decision, not an ongoing operational control that can be adjusted to manage cost in real time.

#### NEW QUESTION # 17

Which of the following responsible AI standards refers to a principle that clearly states the reasons behind the decisions for a particular conclusion?

- A. Accountability
- B. Auditability
- C. Explainability
- D. Transparency

**Answer: C**

Explanation:

Basic Concept: Responsible AI encompasses several key principles governing how AI systems should behave to be trustworthy and ethical. These principles are distinct but related. Understanding their precise definitions is essential for CompTIA SecAI+ Domain 4 governance questions.

Why D is Correct: Explainability in responsible AI means the AI system can clearly articulate the specific reasons, factors, and logic that led to a particular decision or output. It answers the question "why did the AI make this specific decision?" For example, an explainable credit scoring AI would not only give a score but also explain which factors such as payment history or credit utilization contributed most to that specific score.

This directly matches the question's description of "clearly stating reasons behind decisions." Why A is Wrong: Accountability refers to the ability to identify who is responsible for AI system decisions and their consequences. It addresses ownership and responsibility assignment rather than explaining the reasoning behind specific decisions.

Why B is Wrong: Auditability refers to the ability to examine and verify an AI system's decisions, processes, and outputs through systematic review. It enables after-the-fact verification but does not mean the system itself explains its reasoning.

Why C is Wrong: Transparency refers to openness about how an AI system works at a general level, including its purpose, capabilities, limitations, and the data it was trained on. It is broader than explainability and does not specifically address articulating reasons for individual decisions.

