CompTIA

CY0-001

CompTIA SecAI+ v1 Exam

Questions & Answers (Demo)

Version: 4.0

Question: 1

Which of the following job roles in an organizational governance structure develops a model from business use cases?

A. Platform architect

B. AI risk analyst

C. Machine learning operations (MLOps) engineer

D. Data scientist

Answer: D

Explanation:

Basic Concept: In AI governance, each role holds distinct responsibilities. Understanding these roles is core to CompTIA SecAI+ Domain 4 (AI Governance, Risk, and Compliance).

Why D is Correct: The Data Scientist is responsible for translating business use cases into working AI/ML models. They analyze business requirements, identify the appropriate machine learning approach, and develop models that fulfill specific business objectives. According to the CompTIA SecAI+ Study Guide, data scientists bridge raw data and actionable AI solutions by building and validating models derived from business-driven needs.

Why A is Wrong: A Platform Architect designs and manages the infrastructure and technical platforms hosting AI systems. Their focus is architectural design of the environment, not model development from business use cases.

Why B is Wrong: An AI Risk Analyst identifies, evaluates, and mitigates risks associated with AI adoption. Their role is governance and risk-oriented, not model creation.

Why C is Wrong: An MLOps Engineer operationalizes, deploys, monitors, and maintains AI models in production. They take models already built by data scientists and ensure reliable operation at scale, not develop them from business use cases.

Question: 2

An administrator, who works for a financial institution, is required to implement data security controls for data at rest within AI systems that involve data disclosure.

Which of the following is the most suitable control?

A. Data lineage

B. Rate limits

C. Encryption

D. Masking

Answer: C

Explanation:

Basic Concept: Data at rest refers to inactive data stored in databases or storage media. Protecting it from unauthorized disclosure is a fundamental data security principle covered in the CompTIA SecAI+ Study Guide under securing AI data pipelines.

Why C is Correct: Encryption protects data at rest by rendering it unreadable to unauthorized parties without the appropriate decryption key. In a financial institution with sensitive data, encryption at rest (e.g., AES-256) is the primary control against data disclosure. Even if storage media is physically compromised, encrypted data remains unintelligible. CompTIA SecAI+ Exam Objectives highlight encryption as the primary confidentiality control for stored AI data.

Why A is Wrong: Data lineage tracks the origin and movement of data throughout its lifecycle. It improves traceability and auditability but does not prevent unauthorized disclosure of data at rest.

Why B is Wrong: Rate limits control the number of API requests within a time period. They protect against abuse and denial-of-service scenarios, not data-at-rest confidentiality.

Why D is Wrong: Data masking replaces sensitive values with fictitious substitutes, useful during development or testing. For actual production data at rest in AI systems handling real financial records, encryption provides stronger and more comprehensive confidentiality.

Question: 3

A security engineer needs to monitor an AI-based system for runtime operations. The engineer is mostly concerned about the visibility of internal activity.

Which of the following is the most appropriate monitoring solution?

A. Deploying a security information and event management (SIEM) tool

B. Implementing a web application firewall (WAF) with header logging

C. Relying on vendor model controls and monitoring prompt inputs

D. Enabling stack call and debugging level traces at the function level

Answer: D

Explanation:

Basic Concept: Monitoring an AI system's internal runtime behavior requires deep observability into what the system is doing at the code and function execution level, not just at the perimeter. CompTIA SecAI+ Study Guide addresses AI system observability and runtime monitoring under securing AI infrastructure.

Why D is Correct: Enabling stack call and debugging level traces at the function level provides the highest granularity of visibility into internal operations. This approach exposes what functions are called, in what order, with what inputs, and what is returned, offering genuine insight into the AI system's internal activity at runtime precisely as the engineer requires.

Why A is Wrong: A SIEM aggregates and correlates log and event data from multiple sources. While useful for security alerting, it does not inherently provide visibility into internal function-level operations of an AI model at runtime.

Why B is Wrong: A WAF with header logging monitors and filters HTTP traffic at the application boundary. It captures external request and response data, not the AI system's internal runtime mechanics.

Why C is Wrong: Relying on vendor controls and monitoring prompt inputs is a passive, externally-focused approach. It provides no visibility into intermediate computations or internal operations within the AI model itself.

Question: 4

Which of the following should an auditor reference when reviewing a company's human resources AI systems for legal non-compliance?

A. Organization for Economic Cooperation and Development (OECD) standard

B. National Institute of Standards and Technology (NIST) AI Risk Management Framework (RMF)

C. European Union (EU) AI Act

D. International Organization for Standardization (ISO)

Answer: C

Explanation:

Basic Concept: Various regulatory frameworks govern AI use in different contexts. For auditing legal compliance in high-risk AI applications such as employment and HR, binding regulatory legislation takes precedence over voluntary standards. CompTIA SecAI+ Exam Objectives cover AI governance and compliance frameworks under Domain 4.

Why C is Correct: The EU AI Act is the world's first comprehensive, legally binding AI regulation. It explicitly classifies AI systems used in employment, worker management, and recruitment as high-risk AI systems, subjecting them to strict compliance requirements including conformity assessments, transparency obligations, and human oversight mandates. An auditor reviewing HR AI for legal non-compliance must reference this binding legislation.

Why A is Wrong: The OECD AI Principles are non-binding international guidelines promoting responsible AI. They offer policy guidance but carry no legal enforcement power for compliance auditing.

Why B is Wrong: The NIST AI RMF is a voluntary, risk management-focused framework. It is not a legal compliance standard and cannot be used to assess legal non-compliance.

Why D is Wrong: ISO standards such as ISO 42001 are voluntary international best practice standards. They are not legal compliance instruments with enforceable penalties for HR AI systems.

Question: 5

An airline corporation wants to implement a chatbot application using a large language model (LLM) so its customers can ask questions and receive answers about flight details and have the option to upload files.

Which of the following security controls should the airline use to protect against malicious input and unauthorized use beyond the service-level agreement? (Choose two.)

A. Prompt guardrails

B. Role-based access controls

C. Firewall rules

D. Model token quotas

Answer: A, D

Explanation:

Basic Concept: LLM-based chatbots accepting user-uploaded files face two critical risk categories: malicious input injection and resource or cost abuse. CompTIA SecAI+ Study Guide highlights prompt security controls and resource management as key defensive layers for public-facing LLM applications.

Why A is Correct: Prompt guardrails intercept and filter user inputs and model outputs, blocking malicious prompts, prompt injection attempts, and harmful file content before affecting model behavior. Since users can upload files, guardrails are essential for sanitizing and validating that content before processing.

Why D is Correct: Model token quotas directly limit how much of the LLM's processing capacity a user can consume. This prevents abuse beyond the SLA, including denial-of-wallet attacks or resource exhaustion through excessively large inputs or repeated requests.

Why B is Wrong: Role-based access controls manage who can access what resources. While useful for internal systems, they do not address malicious input content or enforce LLM resource consumption limits for a public-facing chatbot.

Why C is Wrong: Firewall rules operate at the network layer and can block unauthorized IPs or ports but cannot inspect or filter the semantic content of prompts or control token-level LLM usage.