Accelerate Your Certification with CompTIA CY0-001 Practice Questions

Last chance to prepare smart! Get your hands on free CompTIA SecAI+ v1 Exam PDF questions. Study real CY0-001 dumps with verified answers and fast-track your certification success with PrepBolt CY0-001 exam pdf questions and answers.

Thank you for Downloading CY0-001 exam PDF Demo

https://prepbolt.com/CY0-001.html

Question 1

Question Type: MultipleChoice

An administrator must conduct generative AI cost monitoring for use in the healthcare industry.

Which of the following criteria is the best way to calculate this cost?

Options:

A- Connection access and exchange gateway
B- Encryption and decryption processing
C- Storage retrieval and prompt processing
D- Catalog servicing and exchange processing

Answer:

C

Explanation:

Basic Concept: Generative AI systems in healthcare settings incur costs from multiple operational activities. Understanding the cost drivers specific to generative AI helps administrators implementaccurate cost monitoring and controls. CompTIA SecAI+ Study Guide covers AI cost management under securing AI systems.

Why C is Correct: Storage retrieval and prompt processing are the two primary cost drivers for generative AI systems in healthcare. Storage retrieval refers to the cost of querying vector databases or document stores in RAG-based AI systems to fetch relevant patient records, clinical guidelines, or historical data for context. Prompt processing encompasses the token-based cost of the LLM processing the combined retrieved content and user query to generate a response. Together these two activities represent the billable units that drive generative AI costs in healthcare RAG deployments, making them the most accurate basis for cost calculation and monitoring.

Why A is Wrong: Connection access and exchange gateway costs relate to network infrastructure and API gateway usage fees. While there may be minor costs associated with API calls, these are not the primary cost drivers for generative AI systems where the dominant expenses are computational token processing and data retrieval operations.

Why B is Wrong: Encryption and decryption processing costs relate to cryptographic operations for data security. While encryption is important for healthcare data protection under HIPAA, cryptographic processing overhead is minimal compared to the substantial token-based LLM processing and storage retrieval costs that dominate generative AI operational expenses.

Why D is Wrong: Catalog servicing and exchange processing are terms associated with data catalog

management and data exchange infrastructure. These are not recognized primary cost components of generative AI systems in healthcare, where storage retrieval and token-based prompt processing are the established cost measurement criteria.

Question 2

Question Type: MultipleChoice

A security analyst notices that regardless of user-submitted prompts, an AI model always returns unsanitized responses. These responses are then passed to multiple plug-ins. The analyst is concerned with the potential security implications.

Which of the following Open Worldwide Application Security Project (OWASP) categories addresses this vulnerability?

Options:

A- Misinformation
B- Prompt injection
C- Unbounded consumption
D- Improper output handling

Answer:

D

Explanation:

Basic Concept: OWASP has published the Top 10 vulnerabilities for Large Language Model Applications, each addressing a distinct category of LLM security risk. Understanding which OWASP category maps to specific LLM vulnerability scenarios is a key competency in the CompTIA SecAI+ Study Guide under securing AI systems.

Why D is Correct: Improper output handling (OWASP LLM02) occurs when an application passes LLM-generated outputs to downstream systems such as plug-ins, web browsers, or databases without proper validation, sanitization, or encoding. This can enable XSS, SQL injection, remote code execution, or other injection attacks against plug-ins and downstream systems. The scenario exactly matches this: unsanitized AI responses are automatically passed to multiple plug-ins, which could execute malicious content in the model's output.

Why A is Wrong: Misinformation refers to the AI generating false or misleading content that users might believe. It is a content accuracy concern related to hallucinations and false information propagation, not a vulnerability describing how model outputs are handled by downstream systems.

Why B is Wrong: Prompt injection involves crafting inputs to manipulate model behavior and override instructions. While it can be a contributing cause of unsafe outputs, the vulnerability described --- passing unsanitized outputs to plug-ins --- is specifically the output handling failure, not the injection mechanism itself.

Why C is Wrong: Unbounded consumption (OWASP LLM10) refers to resource exhaustion attacks including denial-of-wallet and denial-of-service through excessive token consumption. It addresses resource management vulnerabilities, not the security implications of passing model outputs to downstream systems.

Question 3

Question Type: MultipleChoice

Which of the following is the primary security risk when deploying AI models in production?

Options:

A- Graphics processing unit (GPU) acceleration
B- Model overfitting
C- Model encryption
D- Data exposure

Answer:

D

Explanation:

Basic Concept: When AI models are deployed in production, they interact with real data including sensitive business information, personal data, and confidential records. The intersection of AI capabilities and sensitive data creates significant security risks. CompTIA SecAI+ Exam Objectives identify data exposure as the primary production security risk for AI deployments.

Why D is Correct: Data exposure is the primary security risk in production AI deployments. AI models in production process sensitive data through queries and responses, and vulnerabilities such as prompt injection, model inversion attacks, insecure output handling, and misconfigured access controls can expose confidential training data, user PII, proprietary information, or system credentials. The consequences include regulatory violations, legal liability, and reputational damage, making data exposure the most critical ongoing security concern.

Why A is Wrong: GPU acceleration is a performance optimization technique that uses graphics processors for faster AI computation. While hardware security is important, GPU acceleration itself is

not a security risk --- it is a performance feature that does not inherently expose data.

Why B is Wrong: Model overfitting is a model quality issue where a model performs poorly on new data after memorizing training data too specifically. While it can indirectly contribute to data memorization, it is primarily a performance and generalization concern during development rather than a primary production security risk.

Why C is Wrong: Model encryption is a security control used to protect AI model weights from unauthorized access, not a risk itself. Framing a protection mechanism as a primary risk conflates controls with threats.

Question 4

Question Type: MultipleChoice

Which of the following roles best supports the implementation of AI governance, risk, and compliance (GRC)? (Choose two.)

Options:

A- Desktop specialist
B- Data scientist
C- Software developer
D- Security architect
E- Security operations center (SOC) analyst
F- Network engineer

Answer:

B, D

Explanation:

Basic Concept: AI GRC implementation requires roles that combine understanding of AI technical capabilities and limitations with security risk assessment, control design, and compliance framework expertise. Identifying which roles naturally contribute to AI GRC is essential for team design. CompTIA SecAI+ Study Guide covers AI governance role responsibilities under Domain 4.

Why B is Correct: Data Scientists possess deep understanding of AI model capabilities, limitations, data requirements, and failure modes. For GRC implementation, their technical expertise is essential for identifying AI-specific risks such as bias, model drift, and data quality issues, assessing compliance implications of model design choices, and evaluating whether AI systems meet governance requirements.

Why D is Correct: Security Architects design comprehensive security frameworks and risk management strategies. For AI GRC, they translate governance requirements into technical controls, design AI security architectures that satisfy compliance obligations, assess the risk posture of AI deployments, and ensure security principles including least privilege, defense-in-depth, and audit logging are built into AI system designs.

Why A is Wrong: Desktop specialists manage user workstation hardware and software. Their role focuses on endpoint management and user support, not on the strategic risk assessment, compliance evaluation, or technical AI governance activities required for AI GRC implementation.

Why C is Wrong: Software developers write application code. While they implement security controls when directed, they typically lack the broad risk management, compliance framework expertise, and security architecture perspective needed to lead AI GRC implementation.

Why E is Wrong: SOC analysts focus on monitoring, detecting, and responding to security incidents in operational environments. Their expertise is in reactive security operations rather than the proactive governance framework design and compliance management that AI GRC requires.

Why F is Wrong: Network engineers design and maintain network infrastructure. Their expertise is in network connectivity and protocols, not in AI system governance, risk assessment frameworks, or compliance requirements.

Question 5

Question Type: MultipleChoice

During a model validation procedure, an engineer notices that a model performs well during training but poorly during testing.

Which of the following best describes the reason?

Options:

A- Fine-tuning
B- Overfitting
C- Regularization
D- Inference

Answer:

B

Explanation:

Basic Concept: The gap between training performance and test performance is a classic indicator of a specific model quality problem. Understanding this phenomenon and its causes is fundamental to AI model development. CompTIA SecAI+ Study Guide covers overfitting under basic AI concepts and model quality.

Why B is Correct: Overfitting occurs when a model learns the training data too specifically --- memorizing noise, outliers, and specific patterns in the training set rather than learning generalizable underlying patterns. The model achieves high accuracy on training data but fails to generalize to new, unseen test data. This produces exactly the scenario described: excellent training performance combined with poor test performance. Overfitting is the quintessential cause of this training-testing performance gap.

Why A is Wrong: Fine-tuning is a training technique that adapts a pre-trained model to a new task or domain using additional training data. It is a deliberate training process, not a description of why a model's performance degrades from training to testing.

Why C is Wrong: Regularization is a training technique specifically used to prevent overfitting by adding penalties to large model weights, encouraging the model to learn simpler, more generalizable patterns. It is the solution to overfitting, not its cause.

Why D is Wrong: Inference is the process of using a trained model to make predictions on new data. It describes the operational use of a model, not a quality characteristic that explains why performance differs between training and testing phases.

Question 6

Question Type: MultipleChoice

A security analyst reviews a recently released chatbot's log and discovers that outputs sometimes include personally identifiable information (PII) from other chatbot users.

Which of the following corrective actions should the security analyst take first to resolve this issue?

Options:

A- Take the chatbot offline and restore it from a backup.
B- Disable memory from the chat history for all users.
C- Ask all users to refrain from using PII with the chatbot.
D- Require users to label the sensitivity of their requests.

Answer:

B

Explanation:

Basic Concept: When a chatbot leaks PII from one user's conversation into another user's responses, the root cause is cross-user memory contamination --- the chatbot is retaining and sharing conversation context across user sessions. Disabling the memory feature stops the active data leakage immediately. CompTIA SecAI+ Study Guide covers session memory management as a privacy control for AI chatbots.

Why B is Correct: Disabling memory from chat history for all users immediately stops the mechanism causing PII leakage between users. If the chatbot retains no cross-session memory, it cannot include information from one user's conversation in another user's response. This is the most direct, immediate corrective action that eliminates the root cause of the privacy violation without requiring additional user behavior changes or service disruption.

Why A is Wrong: Taking the chatbot offline and restoring from backup is a drastic action appropriate when the issue requires investigating a potential compromise or data breach. For a configuration issue such as cross-user memory sharing, disabling the memory feature is a more targeted and proportionate first response that addresses the root cause directly.

Why C is Wrong: Asking users to refrain from using PII relies on voluntary user behavior change and does not address the technical root cause. Users may not comply, and even if they do, previously stored PII in memory would continue to leak. This is an ineffective first corrective action.

Why D is Wrong: Requiring users to label sensitivity does not stop the chatbot from storing and sharing PII that has already been submitted. Labels inform the system about data sensitivity but do not prevent the memory mechanism from sharing labeled sensitive data across user sessions.

Question 7

Question Type: MultipleChoice

Which of the following helps end users within an organization the most in safeguarding against the risk of AI-related non-compliance?

Options:

A- AI center of excellence
B- Policies and procedures
C- Implementing data loss prevention
D- Enabling multifactor authentication (MFA) for access

Answer:

B

Explanation:

Basic Concept: End users are the employees who interact with AI systems daily and may inadvertently create compliance risks through their AI usage behaviors. Equipping users with clear guidance on acceptable and compliant AI use is the most effective way to reduce compliance violations at the user level. CompTIA SecAI+ Study Guide emphasizes policies and procedures as the foundational compliance tool for end users.

Why B is Correct: Policies and procedures directly inform end users of what AI-related behaviors are compliant, what is prohibited, and how to use AI tools safely and legally. Comprehensive AI usage policies covering acceptable use, data handling requirements, prohibited data inputs, and reporting obligations give users the knowledge they need to avoid compliance violations. Without clear policies, users cannot reliably identify compliant from non-compliant behavior.

Why A is Wrong: An AI center of excellence governs AI adoption at the organizational level, developing standards and approving use cases. While it benefits the organization overall, its governance activities are directed at organizational processes and technical standards rather than providing direct day-to-day compliance guidance to individual end users.

Why C is Wrong: Data loss prevention (DLP) technology automatically prevents the transmission of sensitive data through monitoring and blocking capabilities. While effective at preventing certain compliance violations technically, it cannot guide users on why certain behaviors are non-compliant or how to make compliant choices in situations DLP doesn't cover.

Why D is Wrong: MFA secures user authentication and prevents unauthorized account access. It is an identity security control that protects accounts, not a mechanism that helps users understand or comply with AI governance requirements.

Question 8

Question Type: MultipleChoice

Developers introduce new features to their generative AI product in an effort to stand out from the competition and offer more value to customers.

Which of the following most accurately explains the risks when enabling more functionality?

Options:

A- The risks remain the same as before the new features were added.

B- The risks increase when new features are added.
C- The risks are measured qualitatively.
D- The risks are proportional to the model's capabilities.

Answer:

D

Explanation:

Basic Concept: The relationship between AI system capabilities and security risk is a fundamental concept in AI governance. As AI models gain more functionality and capabilities, their potential for misuse, unintended consequences, and attack surface expansion grows proportionally. CompTIA SecAI+ Study Guide addresses capability-risk proportionality under AI governance.

Why D is Correct: The risks of a generative AI product are proportional to its capabilities. Each new feature expands what the model can do, which simultaneously expands what adversaries can manipulate it to do, what sensitive operations it can be directed to perform, and what unintended harm it can cause. A model that can generate text, images, execute code, and call external APIs has dramatically greater risk potential than one that can only generate text. Risk grows with capability scope.

Why A is Wrong: Risks do not remain constant when new features are added. New features introduce new attack vectors, expand the model's action space, and create new opportunities for misuse. Each addition fundamentally changes the system's risk profile.

Why B is Wrong: While risks do increase with new features, saying they simply increase does not capture the precise relationship. The increase is proportional to the nature and scope of the capabilities added, not a uniform increment for any feature addition.

Why C is Wrong: While risks can be measured qualitatively, stating that risks are measured qualitatively is a statement about measurement methodology rather than an explanation of how risks change when functionality is enabled. It does not accurately describe the relationship between capability and risk.

Thank You for trying CY0-001 PDF Demo

To try our CY0-001 practice exam software visit link below

https://prepbolt.com/CY0-001.html

Start Your CY0-001 Preparation

Use Coupon “SAVE50” for extra 50% discount on the purchase of Practice Test Software. Test your CY0-001 preparation with actual exam questions.