Exam

CY0-001

CompTIA SecAI+

Free demo	26 sample questions
Full exam	126 questions with answers and explanations
Contents	Questions with answers and explanations

Exam domains

Domain	Exam weight	Questions
D1 Basic AI concepts related to cybersecurity		4
D2 Securing AI systems		13
D3 AI-assisted security		4
D4 AI governance, risk, and compliance		5

This is a free demo from Dumpsgate: 26 of the exam's 126 questions. **Get the full CY0-001 exam** with every question, answer and explanation.

Question 1

Single choice · Domain: D1 Basic AI concepts related to cybersecurity › Identify overfitting from a model's training and testing performance

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?

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

Answer: B. Overfitting

Explanation

This question tests recognition of a basic machine-learning failure mode. A model that performs well on training data but poorly on unseen testing data has likely learned the training examples too specifically, which is overfitting.

✗ **A** Fine-tuning
Incorrect. Fine-tuning adapts a pretrained model to a particular task or dataset; it does not specifically describe the performance gap in which training results are strong but testing results are poor.
https://www.comptia.org/content/articles/what-is-artificial-intelligence

✓ **B** Overfitting
Correct. Overfitting occurs when a model learns the training data too closely, including noise or irrelevant patterns, and therefore performs poorly on unseen testing data.
https://www.comptia.org/content/articles/what-is-artificial-intelligence

✗ **C** Regularization
Incorrect. Regularization is a technique used to reduce overfitting by constraining model complexity; it is not the condition described in the question.
https://www.comptia.org/content/articles/what-is-artificial-intelligence

✗ **D** Inference
Incorrect. Inference is the process of using a trained model to generate predictions or outputs from new input data. It does not describe the training-versus-testing performance problem.
https://www.comptia.org/content/articles/what-is-artificial-intelligence

Dumpsgate

Question 2

Single choice · Domain: D1 Basic AI concepts related to cybersecurity › Distinguish unsupervised learning from other machine-learning paradigms

A data scientist is working with unlabeled data and wants to build a clustering model.

Which of the following techniques should a data scientist use?

A. Supervised learning
B. Reinforcement learning
C. Unsupervised learning
D. Semi-supervised learning

Answer: C. Unsupervised learning

Explanation

This question tests recognition that clustering uses unlabeled data to identify inherent groups or patterns, which is the defining use of unsupervised learning.

✗ **A** Supervised learning
Incorrect. Supervised learning requires labeled training data and learns a mapping between inputs and known target outputs.

https://www.comptia.org/certifications/secai

✗ **B** Reinforcement learning
Incorrect. Reinforcement learning trains an agent through interactions with an environment using feedback such as rewards or penalties; it is not the standard technique for clustering unlabeled data.

https://www.comptia.org/certifications/secai

✓ **C** Unsupervised learning
Correct. Unsupervised learning works with unlabeled data and can use clustering to group similar records or identify patterns without predefined target labels.

https://www.comptia.org/certifications/secai

✗ **D** Semi-supervised learning
Incorrect. Semi-supervised learning combines a relatively small amount of labeled data with a larger amount of unlabeled data, whereas the question specifies unlabeled data for clustering.

https://www.comptia.org/certifications/secai

Dumpsgate

Question 3

Single choice · Domain: D1 Basic AI concepts related to cybersecurity › Identify tokens as the units used to measure LLM usage and utilization

A company introduces a large language model (LLM) in an application in order to monitor for a potential denial-of-service attack.

Which of the following should the company use to measure the utilization of the LLM?

A. Token
B. Transformer
C. Chain of thoughts
D. Prompt

Answer: A. Token

Explanation

This question tests foundational LLM terminology by distinguishing tokens, which represent processed units of text and are used to measure model usage, from model architecture, reasoning processes, and input instructions.

✓ **A** Token
Correct. Tokens are the units of text processed by an LLM, so token counts can be used to measure model utilization and consumption.

https://www.comptia.org/content/articles/what-is-generative-ai

✗ **B** Transformer
Incorrect. A transformer is a neural-network architecture used by many modern LLMs; it is not the unit used to measure utilization.

https://www.comptia.org/content/articles/what-is-generative-ai

✗ **C** Chain of thoughts
Incorrect. Chain of thought refers to an intermediate reasoning process or prompting approach, not a quantitative measure of LLM utilization.

https://www.comptia.org/content/articles/what-is-generative-ai

✗ **D** Prompt
Incorrect. A prompt is the input instruction or text provided to an LLM, whereas utilization is commonly measured by the number of tokens processed.

https://www.comptia.org/content/articles/what-is-generative-ai

Dumpsgate

Question 4

Single choice · Domain: D1 Basic AI concepts related to cybersecurity › Identify the neural-network technology used to create synthetic or manipulated media

A short AI-generated video shows a celebrity's likeness talking about a fake public security event.

Which of the following was used to create this video?

A. Statistical analysis
B. Convolutional neural network
C. Machine learning (ML) classifier
D. Random forest

Answer: B. Convolutional neural network

Explanation

The question tests recognition of the AI technology commonly used to process and synthesize visual data in deepfake videos. A convolutional neural network is designed to learn spatial features from images and video, making it suitable for manipulating a celebrity's likeness.

✗ **A** Statistical analysis
Incorrect. Statistical analysis can identify patterns or anomalies but is not the primary neural-network technology used to generate a synthetic video likeness.

https://www.comptia.org/certifications/secai

✓ **B** Convolutional neural network
Correct. Convolutional neural networks process visual features and are commonly used in computer-vision applications, including the creation or manipulation of faces in deepfake video.

https://www.comptia.org/certifications/secai

✗ **C** Machine learning (ML) classifier
Incorrect. An ML classifier assigns data to categories or labels; it is generally used for detection or classification rather than generating a synthetic celebrity video.

https://www.comptia.org/certifications/secai

✗ **D** Random forest
Incorrect. A random forest is an ensemble of decision trees typically used for classification or regression, not for synthesizing realistic video likenesses.

https://www.comptia.org/certifications/secai

Free demo · full exam: 126 questions

Question 5

Single choice · Domain: D2 Securing AI systems › Identify improper output handling vulnerabilities in AI applications

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?

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

Answer: D. Improper output handling

Explanation

The question tests recognition of a vulnerability in which an AI model's unsanitized output is passed to downstream plug-ins, creating opportunities for injection or other security impacts.

✗ **A** Misinformation
Incorrect. Misinformation concerns inaccurate or misleading content, not the failure to sanitize model output before passing it to downstream components. CompTIA's SecAI+ certification covers securing AI systems and AI-related security risks.
https://www.comptia.org/en-us/certifications/secai/

✗ **B** Prompt injection
Incorrect. Prompt injection involves manipulating an AI system through crafted instructions. The described issue is that the model's responses are unsanitized and then passed to plug-ins, which is an output-handling vulnerability. CompTIA's SecAI+ certification covers AI security risks and controls.
https://www.comptia.org/en-us/certifications/secai/

✗ **C** Unbounded consumption
Incorrect. Unbounded consumption involves excessive or unrestricted use of resources, such as tokens, compute, or API calls. It does not describe unsanitized AI responses being sent to plug-ins. CompTIA's SecAI+ certification addresses risks associated with securing AI systems.
https://www.comptia.org/en-us/certifications/secai/

✓ **D** Improper output handling
Correct. Improper output handling occurs when AI-generated output is trusted or forwarded to downstream applications without appropriate validation or sanitization. Passing unsanitized responses to multiple plug-ins matches this vulnerability. CompTIA identifies securing AI systems as a core SecAI+ area.
https://www.comptia.org/en-us/certifications/secai/

Question 6

Single choice · Domain: D2 Securing AI systems › Identify data exposure as a primary production security risk for deployed AI models

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

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

Answer: D. Data exposure

Explanation

This question tests recognition of sensitive-data exposure and leakage risks when AI models are deployed in production. It focuses on protecting AI systems and their data rather than model performance or hardware configuration.

✗ **A** Graphics processing unit (GPU) acceleration
Incorrect. GPU acceleration is a performance and processing capability, not inherently a primary security risk.

https://www.comptia.org/certifications/secai

✗ **B** Model overfitting
Incorrect. Overfitting primarily affects model generalization and accuracy; it is not the primary security risk described here.

https://www.comptia.org/certifications/secai

✗ **C** Model encryption
Incorrect. Model encryption is a protective security measure intended to help prevent unauthorized access, not a security risk itself.

https://www.comptia.org/certifications/secai

✓ **D** Data exposure
Correct. AI models in production may expose sensitive training data, prompts, outputs, or connected-system information if they are improperly secured.

https://www.comptia.org/certifications/secai

Free demo · full exam: 126 questions

Question 7

Single choice · Domain: D2 Securing AI systems › Disable chatbot memory or chat-history retention to prevent cross-user PII disclosure

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?

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. Disable memory from the chat history for all users.

Explanation

This question tests the ability to apply an immediate technical control when a chatbot exposes one user's PII to another user. Disabling shared or persistent chat-history memory directly limits the mechanism causing the data leakage.

✗ **A** Take the chatbot offline and restore it from a backup.
Incorrect. Taking the chatbot offline may contain the incident, but restoring from a backup does not necessarily remove the configuration or memory behavior that caused cross-user PII disclosure. The appropriate first corrective action is to disable the problematic memory function.

https://www.comptia.org/certifications/secai

✓ **B** Disable memory from the chat history for all users.
Correct. Disabling chat-history memory for all users directly prevents the chatbot from retaining and reusing one user's information in another user's responses, reducing the immediate risk of PII leakage.

https://www.comptia.org/certifications/secai

✗ **C** Ask all users to refrain from using PII with the chatbot.
Incorrect. User guidance may reduce future exposure but does not remediate the existing cross-user data-sharing behavior, and users may still enter PII unintentionally. A technical control should be applied first.

https://www.comptia.org/certifications/secai

✗ **D** Require users to label the sensitivity of their requests.
Incorrect. Sensitivity labels do not prevent the chatbot from retaining or disclosing PII and depend on users to classify requests accurately. Disabling the memory mechanism is the direct corrective action.

https://www.comptia.org/certifications/secai

Free demo · full exam: 126 questions

Question 8

Single choice · Domain: D2 Securing AI systems › Assess how increased generative AI functionality and model capability affect security risk

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?

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. The risks are proportional to the model's capabilities.

Explanation

This question tests whether the candidate can recognize that expanding an AI product's functionality and capabilities can increase its attack surface and potential impact. The most precise relationship is that risk grows with the model's capabilities.

✗ **A** The risks remain the same as before the new features were added.
Incorrect. New functionality can introduce additional attack paths, misuse cases, dependencies, and impacts, so assuming the risk remains unchanged is unsafe.

https://www.comptia.org/content/articles/what-is-artificial-intelligence

✗ **B** The risks increase when new features are added.
Although adding features commonly increases exposure, this answer is less precise than relating risk to the model's capabilities. A feature's risk depends on what capability it adds, how it is implemented, and how it can be misused.

https://www.comptia.org/content/articles/what-is-artificial-intelligence

✗ **C** The risks are measured qualitatively.
Incorrect. Risk assessment can use both qualitative and quantitative methods; the question is about how risk changes as AI functionality and capability increase, not merely how risk is measured.

https://www.comptia.org/content/articles/what-is-artificial-intelligence

✓ **D** The risks are proportional to the model's capabilities.
Correct. More capable models and more extensive functionality can produce greater benefits but also enable more harmful actions, increase the attack surface, and magnify potential impact. Therefore, risk generally scales with the capabilities being provided.

https://www.comptia.org/content/articles/what-is-artificial-intelligence

Free demo · full exam: 126 questions

Question 9

Single choice · Domain: D2 Securing AI systems › Apply data masking to prevent sensitive information disclosure in AI-generated outputs

A healthcare company deploys an AI chatbot that implements retrieval-augmented generation (RAG) using the company's historical data set. The chatbot output contains patient information.

Which of the following is the most effective technique to mitigate this vulnerability?

A. Masking
B. Classification
C. Minimization
D. Normalization

Answer: A. Masking

Explanation

This question tests the ability to protect sensitive patient information exposed through a retrieval-augmented generation system. Masking directly redacts or obscures sensitive values before they are presented in the chatbot output.

✓ **A** Masking
Correct. Masking obscures sensitive data such as patient identifiers, reducing the risk that confidential information retrieved from the historical data set is disclosed in the chatbot response.
https://www.comptia.org/en-us/certifications/secai/

✗ **B** Classification
Incorrect. Classification identifies or labels data according to sensitivity, but labeling alone does not prevent patient information from appearing in the chatbot output.
https://www.comptia.org/en-us/certifications/secai/

✗ **C** Minimization
Incorrect. Data minimization reduces the amount of data collected or retained and can lower exposure, but it does not directly obscure patient information that the RAG system retrieves and outputs.
https://www.comptia.org/en-us/certifications/secai/

✗ **D** Normalization
Incorrect. Normalization standardizes data structure or format; it does not protect confidential patient information from disclosure.
https://www.comptia.org/en-us/certifications/secai/

Dumpsgate

Question 10

Single choice · Domain: D2 Securing AI systems › Implement authentication and access controls at the appropriate AI system layer

A security administrator must implement security controls for AI systems.

Which of the following access controls should the administrator set up first for authentication?

A. Model
B. Server
C. Data
D. Endpoint

Answer: D. Endpoint

Explanation

This question tests the ability to identify the initial authentication control point in an AI environment. Endpoint authentication should be established first so that users and devices are verified before they can access servers, models, or data.

✗ **A** Model
Incorrect. Model-level controls protect access to an AI model, but authentication should begin at the endpoint before a user or device reaches the model.

https://www.comptia.org/certifications/secai

✗ **B** Server
Incorrect. Server authentication and authorization are important, but they should follow initial endpoint authentication and verification.

https://www.comptia.org/certifications/secai

✗ **C** Data
Incorrect. Data access controls protect information after access to the environment has been authenticated and authorized; they are not the first authentication control.

https://www.comptia.org/certifications/secai

✓ **D** Endpoint
Correct. Endpoint authentication verifies the user or device at the initial access point, establishing the first control before access is granted to servers, models, or data.

https://www.comptia.org/certifications/secai

Dumpsgate

Question 11

Single choice · Domain: D2 Securing AI systems › Implement human review and approval controls for machine learning model and pipeline changes

A data scientist investigates reports that a production machine learning (ML) model no longer performs with accuracy.

The data scientist finds the following pipeline log entries:

Time	Event	User
2025-Dec-15 01:00:00 UTC	training_db cloned	svc_runner
2025-Dec-15 01:05:01 UTC	training_db 15 changes committed	svc_runner
2025-Dec-15 01:05:04 UTC	model_training job started	svc_runner
2025-Dec-15 04:18:07 UTC	training_db 10000 changes committed	iam_breakglass
2025-Dec-15 07:15:27 UTC	model_bin merged to main	svc_runner
2025-Dec-15 07:15:28 UTC	model_training job started	svc_runner
2025-Dec-15 19:15:10 UTC	model_bin merged to main	svc_runner
2025-Dec-16 08:31:54 UTC	model_bin cloned locally	devuser1

Which of the following should the security team do to mitigate future occurrences?

A. Add static code scanning tooling to the runner job.
B. Enable human review and approval workflows in the repository.
C. Retrain the model on using increased data and epochs.
D. Keep multiple copies of the model for restoration.

Answer: B. Enable human review and approval workflows in the repository.

Explanation

The question tests securing the ML development and deployment pipeline against unauthorized or insufficiently reviewed changes. The log shows a break-glass commit and model binaries merged into the main branch without evidence of human approval, so repository protection and approval workflows are the most appropriate mitigation.

✗ **A** Add static code scanning tooling to the runner job.
Incorrect. Static code scanning may identify certain source-code flaws, but it does not prevent unauthorized model merges or ensure that privileged changes receive human review and approval.

https://www.comptia.org/certifications/secai

✓ **B** Enable human review and approval workflows in the repository.
Correct. The log shows a break-glass commit followed by model binaries being merged into the main branch, indicating inadequate change control. Required human review and approval can prevent unreviewed model or pipeline changes from reaching production.

https://www.comptia.org/certifications/secai

✗ **C** Retrain the model on using increased data and epochs.
Incorrect. Increasing training data or epochs may affect model accuracy, but it does not mitigate the security and change-management weakness indicated by unauthorized or insufficiently reviewed repository activity.

https://www.comptia.org/certifications/secai

✗ **D** Keep multiple copies of the model for restoration.
Incorrect. Multiple copies support recovery and availability, but they do not prevent unauthorized commits, model merges, or unreviewed changes from being introduced into the production pipeline.

Free demo · full exam: 126 questions

https://www.comptia.org/certifications/secai

Question 12

Multiple choice · Domain: D2 Securing AI systems › Select secure protocols, cryptographic mechanisms, and privacy-preserving data transformation techniques for AI system data flows

Part 1: Use drop-down menu to select the most appropriate protocol or cipher for each system component.

Part 2: Use the drop-down menu to select the most appropriate technique to apply to the modified data.

An engineer is analyzing findings from a penetration test that indicate insufficient data encryption.

The engineer must implement data security.

Dumpsgate


Part 1
Part 2
End client
API gateway
Database
AI model
Select...
RSA-2048
3DES
AES-512
TLS 1.2
gRPC
SSL 1.0
SHA-2 512
HMAC-SHA512
RC4

Part 1
Part 2
Technique
Original data
Modified data
Select...
Anonymization
Classification
De-identification
Tokenization
Masking
{pin:"999-99-9999", name:"john doe"}
{pin:"999-99-9999", name:"john doe", sensitivity:"SECRET"}
{ip_addr:"1.2.3.4", cookie:"aKCIdkd==", name:"John Doe" uid="1111"}
{cookie:"aKCIdkd==", uid="1111"}
{pin:"999-99-XXXX", name:"john doe"}
{card_number:"111122223333444", name:"john doe", user_id:"1"}
{name:"john doe", patient_id:"10000", dob:"1980-jan-05"}
{patient_id:"10000", dob:"1980-jan-05"}

A. RSA-2048
B. 3DES
C. AES-512
D. TLS 1.2
E. gRPC
F. SSL 1.0
G. SHA-2 512
H. HMAC-SHA512
I. RC4
J. Classification
K. Anonymization
L. De-identification
M. Tokenization
N. Masking

Answer

- A. **RSA-2048**
- D. **TLS 1.2**
- G. **SHA-2 512**
- H. **HMAC-SHA512**
- J. **Classification**
- K. **Anonymization**
- L. **De-identification**
- M. **Tokenization**
- N. **Masking**

Explanation

This question tests the ability to protect communications and sensitive data moving between an end client, API gateway, database, and AI model. It also tests selecting appropriate techniques such as tokenization, masking, anonymization, de-identification, and classification.

✓ **A** RSA-2048
Correct when used for public-key operations such as key exchange or digital signatures, but it is not efficient for encrypting large data volumes directly.

https://www.comptia.org/certifications/secai

✗ **B** 3DES
Incorrect for a new implementation because 3DES is obsolete and should not be selected when modern encryption is available.

https://www.comptia.org/certifications/secai

✗ **C** AES-512
Incorrect because AES does not define a standard 512-bit key size; recognized AES key sizes are 128, 192, and 256 bits.

https://www.comptia.org/certifications/secai

✓ **D** TLS 1.2
Correct for protecting data in transit between system components such as the client, API gateway, database service, and AI service. TLS 1.3 would generally be preferred where supported, but TLS 1.2 remains a valid secure protocol when securely configured.

https://www.comptia.org/certifications/secai

✗ **E** gRPC
Incorrect as an encryption or security mechanism. gRPC is a service communication framework and must use a protection mechanism such as TLS to secure the connection.

https://www.comptia.org/certifications/secai

✗ **F** SSL 1.0
Incorrect because SSL 1.0 is obsolete and must not be used for modern protected communications.

https://www.comptia.org/certifications/secai

✓ **G** SHA-2 512
Correct when the requirement is hashing for integrity or use within a broader authentication design. SHA-512 is a one-way hash and is not a substitute for reversible encryption.

https://www.comptia.org/certifications/secai

✓ **H** HMAC-SHA512
Correct for providing message integrity and authentication when the communicating parties share a secret key. It does not provide confidentiality by itself.

https://www.comptia.org/certifications/secai

✗ **I** RC4
Incorrect because RC4 is an obsolete and insecure stream cipher that should not be used in a new implementation.

https://www.comptia.org/certifications/secai

✓ **J** Classification
Correct when the modified data adds a sensitivity label, such as Confidential or Secret, without changing the underlying data. The classification can drive handling and access-control decisions.

https://www.comptia.org/certifications/secai

✓ **K** Anonymization
Correct when identifying information is removed or transformed so that an individual cannot reasonably be identified from the resulting data.

https://www.comptia.org/certifications/secai

✓ **L** De-identification
Correct when direct identifiers such as a name, address, or patient identifier are removed or changed. Residual information may still permit re-identification in some circumstances.

https://www.comptia.org/certifications/secai

✓ **M** Tokenization
Correct when a sensitive value is replaced with a non-sensitive token while the original value is retained in a separate protected system or vault.

https://www.comptia.org/certifications/secai

✓ **N** Masking
Correct when part of a sensitive value is hidden while enough information remains visible for recognition, such as displaying only the last four digits of a card number.

https://www.comptia.org/certifications/secai

Question 13

Single choice · Domain: D2 Securing AI systems › Select compensating controls that detect and mitigate prompt injection and unsafe LLM output

A security administrator wants to prevent prompt injection attacks and ensure responses have sanitized output.

Which of the following provides a primary compensating control for these requirements?

A. Least privilege
B. Encryption
C. A large language model (LLM) firewall
D. Rate limiting

Answer: C. A large language model (LLM) firewall

Explanation

This question tests the ability to identify an AI-specific security control that filters prompts and validates or sanitizes model responses. An LLM firewall is designed to provide protections at the model interaction layer against prompt injection and harmful output.

✗ **A** Least privilege
Incorrect. Least privilege limits permissions and reduces the impact of unauthorized access, but it does not directly inspect prompts or sanitize LLM responses.
https://www.comptia.org/certifications/secai

✗ **B** Encryption
Incorrect. Encryption protects data confidentiality in transit or at rest; it does not detect prompt injection attacks or sanitize generated output.
https://www.comptia.org/certifications/secai

✓ **C** A large language model (LLM) firewall
Correct. An LLM firewall can inspect and filter prompts for injection attempts and apply controls to validate, block, or sanitize model output before it reaches the user or downstream system.
https://www.comptia.org/certifications/secai

✗ **D** Rate limiting
Incorrect. Rate limiting restricts request volume and can help mitigate abuse or resource exhaustion, but it does not primarily detect prompt injection or sanitize responses.
https://www.comptia.org/certifications/secai

Dumpsgate

Question 14

Single choice · Domain: D2 Securing AI systems › Apply model guardrails to reduce system-prompt disclosure

A SOC analyst identifies that a user extracted the full system prompt from the company's chatbot by prompting it to repeat the last query and provide the entire conversation context. Which of the following mitigations reduces the risk to the AI system?

A. Restricting the LLM's access to internal services
B. Using data version control to detect content manipulation
C. Enhancing model guardrails
D. Segregating and identifying external content

Answer: C. Enhancing model guardrails

Explanation

The question tests the ability to select a control that prevents an AI model from revealing protected system instructions and conversation context in response to prompt-based extraction attempts.

✗ **A** Restricting the LLM's access to internal services
Incorrect. Restricting access to internal services reduces risks such as excessive agency and unauthorized tool use, but it does not directly prevent the chatbot from disclosing its system prompt.

https://www.comptia.org/certifications/secai

✗ **B** Using data version control to detect content manipulation
Incorrect. Data version control helps identify changes or manipulation in datasets and other content; it is not the primary mitigation for prompt extraction or system-prompt leakage.

https://www.comptia.org/certifications/secai

✓ **C** Enhancing model guardrails
Correct. Stronger model guardrails can instruct the model not to reveal system prompts, hidden instructions, or confidential conversation context, reducing the risk of prompt extraction.

https://www.comptia.org/certifications/secai

✗ **D** Segregating and identifying external content
Incorrect. Separating and labeling external content primarily helps reduce risks from untrusted or indirect prompt injection. It does not directly address a user asking the model to disclose its system prompt.

https://www.comptia.org/certifications/secai

Free demo · full exam: 126 questions

Question 15

Single choice · Domain: D2 Securing AI systems › Select encryption at rest to protect confidential data stored by an AI system

A recent release of an AI software update exposes confidential customer information due to storage misconfiguration.

Which of the following data security controls will help maintain confidentiality despite the data leak?

A. Model encryption
B. Encryption in transit
C. Encryption in use
D. Encryption at rest

Answer: D. Encryption at rest

Explanation

This question tests identifying the encryption state that protects data stored on systems or storage media when a storage configuration exposes it. Encryption at rest preserves confidentiality even if unauthorized parties obtain the stored data.

✗ **A** Model encryption
Incorrect. Model encryption is not the specific control that protects confidential customer data stored in a misconfigured storage location.

https://www.comptia.org/content/articles/what-is-data-encryption

✗ **B** Encryption in transit
Incorrect. Encryption in transit protects data while it moves between systems, not data exposed because of a storage misconfiguration.

https://www.comptia.org/content/articles/what-is-data-encryption

✗ **C** Encryption in use
Incorrect. Encryption in use protects data while it is being processed or used in memory; it does not directly address data stored in a misconfigured repository.

https://www.comptia.org/content/articles/what-is-data-encryption

✓ **D** Encryption at rest
Correct. Encryption at rest protects stored data by making it unreadable without the required decryption key, helping maintain confidentiality if a storage misconfiguration exposes the data.

https://www.comptia.org/content/articles/what-is-data-encryption

Dumpsgate

Question 16

Single choice · Domain: D2 Securing AI systems › Implement prompt-firewall controls to detect and block suspicious or malicious AI queries

A security administrator sees suspicious queries on AI logs.

Which of the following should the administrator implement to address this issue?

A. Prompt firewalls
B. Data size
C. Rate limit
D. Agentic AI

Answer: A. Prompt firewalls

Explanation

The question tests selection of a control that protects an AI system from suspicious prompt activity. A prompt firewall can inspect AI queries and identify or block potentially malicious input before it reaches the model.

✓ **A** Prompt firewalls
Correct. Prompt firewalls inspect prompts or queries for suspicious, malicious, or policy-violating content and can block or flag the activity.

https://www.comptia.org/certifications/secai

✗ **B** Data size
Incorrect. Data size is a capacity or processing consideration, not a security control for detecting and addressing suspicious AI queries.

https://www.comptia.org/certifications/secai

✗ **C** Rate limit
Incorrect. Rate limiting restricts the number or frequency of requests and can help mitigate abuse or denial-of-service activity, but it does not inspect the content of suspicious queries.

https://www.comptia.org/certifications/secai

✗ **D** Agentic AI
Incorrect. Agentic AI describes AI that can plan and take actions, not a control for detecting or blocking suspicious queries in AI logs.

https://www.comptia.org/certifications/secai

Dumpsgate

Question 17

Single choice · Domain: D2 Securing AI systems › Select a prompt firewall to validate and filter input to an AI system

A cybersecurity administrator needs a security mechanism that can validate input.

Which of the following controls should the administrator use?

A. Prompt firewall
B. Rate limits
C. Token limits
D. Input quantity

Answer: A. Prompt firewall

Explanation

This question tests the ability to identify an AI security control that examines and validates incoming prompts or other input before processing. A prompt firewall is designed for this purpose, whereas rate, token, and quantity limits constrain usage rather than validate content.

✓ **A** Prompt firewall
Correct. A prompt firewall can inspect, validate, and filter input before it reaches an AI model, helping prevent malicious or malformed prompts.
https://www.comptia.org/certifications/secai

✗ **B** Rate limits
Incorrect. Rate limits control how frequently requests can be submitted; they do not validate the content of the input.
https://www.comptia.org/certifications/secai

✗ **C** Token limits
Incorrect. Token limits restrict the size of model input or output, but they do not validate whether the input is safe or appropriate.
https://www.comptia.org/certifications/secai

✗ **D** Input quantity
Incorrect. Limiting input quantity controls volume or size rather than validating the content of the input.
https://www.comptia.org/certifications/secai

Free demo · full exam: 126 questions

Question 18

Single choice · Domain: D3 AI-assisted security › Use an AI agent to automate SOC investigation and containment through SOAR playbooks

A SOC team has an AI agent that performs web searches and calls to the SOAR solution. The team is concerned about enterprise uptime and case resolution time.

Which of the following is the most appropriate use of the AI agent?

A. To analyze and contain offending users or hosts using SOAR playbooks
B. To perform research using open-source intelligence to enrich the alerts
C. To aggregate SOC metrics and generate reports for the leadership team
D. To create tabletop exercises so the team can increase its incident response speed

Answer: A. To analyze and contain offending users or hosts using SOAR playbooks

Explanation

The question tests selecting an AI-agent use case that directly improves enterprise uptime and case resolution time. Automating analysis and containment through SOAR playbooks enables faster incident response and reduces the impact and duration of security incidents.

✓ **A** To analyze and contain offending users or hosts using SOAR playbooks
Correct. An AI agent that can call SOAR can analyze alerts and initiate approved containment playbooks, directly reducing incident impact, improving uptime, and shortening case resolution time.

https://www.comptia.org/certifications/secai

✗ **B** To perform research using open-source intelligence to enrich the alerts
Incorrect. OSINT enrichment can improve alert context, but it does not directly contain threats or restore service, so it is less appropriate for the stated uptime and case-resolution objectives.

https://www.comptia.org/certifications/secai

✗ **C** To aggregate SOC metrics and generate reports for the leadership team
Incorrect. Metrics and reporting support management visibility but do not directly improve operational uptime or accelerate containment and case resolution.

https://www.comptia.org/certifications/secai

✗ **D** To create tabletop exercises so the team can increase its incident response speed
Incorrect. Tabletop exercises can improve preparedness over time, but creating them does not provide the immediate automated investigation and containment capability needed to improve uptime and resolution time.

https://www.comptia.org/certifications/secai

Free demo · full exam: 126 questions

Question 19

Single choice · Domain: D3 AI-assisted security › Use prompt templates to provide domain context and reduce semantic ambiguity in AI chatbot responses

An organization implements a domain-specific AI chatbot. After operating normally for weeks, the model returns contextually incorrect responses — treating 'worm' as a biological pest rather than a computer worm when answering a cybersecurity question.

Which of the following should the organization do to address the issue?

A. Configure guardrails.
B. Encrypt the weights at rest.
C. Apply model access controls.
D. Deploy prompt templates.

Answer: D. Deploy prompt templates.

Explanation

This question tests selecting a prompt-engineering technique that helps a domain-specific model interpret ambiguous terms according to the cybersecurity context. Prompt templates can supply consistent context and instructions, whereas encryption and access controls address protection and authorization rather than response accuracy.

✗ **A** Configure guardrails.
Incorrect. Guardrails constrain or filter model behavior, such as unsafe content or prohibited actions, but they do not primarily resolve an ambiguous term being interpreted in the wrong domain.
https://www.comptia.org/certifications/secai

✗ **B** Encrypt the weights at rest.
Incorrect. Encrypting model weights at rest protects stored model data from unauthorized disclosure or tampering; it does not improve the model's contextual interpretation of the word 'worm.'.
https://www.comptia.org/certifications/secai

✗ **C** Apply model access controls.
Incorrect. Model access controls restrict who can use or administer the model and what actions they can perform. They do not correct semantic ambiguity in chatbot responses.
https://www.comptia.org/certifications/secai

✓ **D** Deploy prompt templates.
Correct. A prompt template can consistently identify the cybersecurity context and define relevant terminology, helping the chatbot interpret 'worm' as malware rather than a biological pest.
https://www.comptia.org/certifications/secai

Free demo · full exam: 126 questions

Question 20

Single choice · Domain: D3 AI-assisted security › Identify AI techniques used to detect malicious websites and phishing links

An organization deploys a browser-based AI plug-in to detect malicious websites and phishing links in corporate email.

Which of the following techniques is used in this AI plug-in?

A. Code quality testing
B. Pattern recognition and signature matching
C. Automated penetration testing
D. Automated incident response

Answer: B. Pattern recognition and signature matching

Explanation

This question tests recognition of how AI-assisted security tools use pattern recognition and signature matching to identify suspicious websites and phishing URLs in email and browser activity.

✗ **A** Code quality testing
Incorrect. Code quality testing evaluates software code for defects, maintainability, or vulnerabilities; it is not the primary technique used to identify malicious websites or phishing links.

https://www.comptia.org/certifications/security/secai

✓ **B** Pattern recognition and signature matching
Correct. An AI plug-in that detects malicious websites and phishing links can analyze patterns in URLs, page content, and known indicators, while signature matching compares activity against recognized malicious signatures.

https://www.comptia.org/certifications/security/secai

✗ **C** Automated penetration testing
Incorrect. Automated penetration testing actively evaluates systems by simulating attacks; it is not the normal detection technique used by a browser plug-in to identify malicious links during email or web use.

https://www.comptia.org/certifications/security/secai

✗ **D** Automated incident response
Incorrect. Automated incident response performs containment or remediation actions after a security event is identified, whereas this plug-in is focused on detecting malicious websites and phishing links.

https://www.comptia.org/certifications/security/secai

Free demo · full exam: 126 questions

Question 21

Single choice · Domain: D3 AI-assisted security › Apply human-in-the-loop oversight to AI-generated recommendations

During an update, an AI system flags some potential compatibility issues and provides recommendations. An administrator reviews the recommendations before addressing the issues.

Which of the following processes describes this scenario?

A. Data validation
B. Data preparation
C. Human-in-the-loop
D. Model evaluation

Answer: C. Human-in-the-loop

Explanation

The scenario describes an AI system generating compatibility recommendations while an administrator reviews them before action is taken. This is human-in-the-loop operation, where human judgment supervises or validates AI output.

✗ **A** Data validation
Incorrect. Data validation checks whether input data meets defined quality, format, or accuracy requirements; it does not describe a human reviewing AI recommendations before implementation.
https://www.comptia.org/certifications/secai

✗ **B** Data preparation
Incorrect. Data preparation involves collecting, cleaning, transforming, and organizing data for analysis or model use. The scenario focuses on human oversight of AI output, not preparing data.
https://www.comptia.org/certifications/secai

✓ **C** Human-in-the-loop
Correct. Human-in-the-loop describes an AI process in which a person reviews, validates, or approves AI-generated recommendations before action is taken.
https://www.comptia.org/certifications/secai

✗ **D** Model evaluation
Incorrect. Model evaluation measures a model's performance against defined criteria, such as accuracy or precision. The administrator is reviewing recommendations for operational approval rather than formally evaluating model performance.
https://www.comptia.org/certifications/secai

Dumpsgate

Question 22

Single choice · Domain: D4 AI governance, risk, and compliance › Calculate generative AI usage costs using storage retrieval and prompt-processing activity

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?

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. Storage retrieval and prompt processing

Explanation

This question tests the ability to identify the principal usage factors for monitoring generative AI costs in a regulated healthcare environment. Storage retrieval and prompt processing are measurable consumption activities, unlike the other choices, which describe unrelated or incomplete infrastructure functions.

✗ **A** Connection access and exchange gateway
Incorrect. Connection access and exchange gateways describe connectivity or integration components, but they are not the primary measures used to calculate generative AI consumption costs.
https://www.comptia.org/certifications/secai

✗ **B** Encryption and decryption processing
Incorrect. Encryption and decryption may contribute to infrastructure overhead, but they are not the main generative AI usage metrics for cost calculation.
https://www.comptia.org/certifications/secai

✓ **C** Storage retrieval and prompt processing
Correct. Generative AI costs are commonly monitored through resource consumption such as storage retrieval and prompt processing, including the processing of input and output tokens.
https://www.comptia.org/certifications/secai

✗ **D** Catalog servicing and exchange processing
Incorrect. Catalog servicing and exchange processing are not standard primary measures for calculating generative AI usage costs.
https://www.comptia.org/certifications/secai

Question 23

Multiple choice · Domain: D4 AI governance, risk, and compliance › Identify organizational roles that support AI governance, risk, and compliance implementation

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

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

Answer

- D. **Security architect**
- E. **Security operations center (SOC) analyst**

Explanation

This question tests the ability to distinguish security roles that implement, enforce, and monitor AI governance and compliance controls from technical roles primarily focused on development, data analysis, networking, or endpoint support.

✗ **A** Desktop specialist
Incorrect. A desktop specialist primarily supports end-user systems and endpoint troubleshooting, rather than implementing organization-wide AI governance, risk, and compliance controls.

https://www.comptia.org/certifications/secai

✗ **B** Data scientist
Incorrect. A data scientist may contribute technical insight into data and model behavior, but the role is not primarily responsible for implementing security governance, risk, and compliance controls.

https://www.comptia.org/certifications/secai

✗ **C** Software developer
Incorrect. A software developer may implement secure coding practices and AI functionality, but software development is not the role most directly responsible for organizational AI GRC implementation.

https://www.comptia.org/certifications/secai

✓ **D** Security architect
Correct. A security architect designs security architectures, controls, and risk treatments that can operationalize AI governance and compliance requirements across systems.

https://www.comptia.org/certifications/secai

✓ **E** Security operations center (SOC) analyst
Correct. A SOC analyst supports ongoing monitoring, detection, investigation, and evidence collection, which helps enforce and demonstrate compliance with AI security governance requirements.

https://www.comptia.org/certifications/secai

✗ **F** Network engineer
Incorrect. A network engineer supports network availability and security, but does not generally have primary responsibility for implementing AI governance, risk, and compliance programs.

https://www.comptia.org/certifications/secai

Free demo · full exam: 126 questions

Question 24

Single choice · Domain: D4 AI governance, risk, and compliance › Apply AI-related policies and procedures to guide compliant end-user behavior

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

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

Answer: B. Policies and procedures

Explanation

Policies and procedures provide end users with explicit rules and required practices for using AI in accordance with organizational, legal, and regulatory requirements. The question focuses on preventing non-compliant behavior rather than controlling access or preventing data exfiltration.

✗ **A** AI center of excellence
Incorrect. An AI center of excellence can provide centralized expertise, standards, and oversight, but it does not directly give all end users the operational rules they must follow. Policies and procedures are the primary mechanism for communicating compliant behavior.

https://www.comptia.org/content/articles/what-is-artificial-intelligence-ai

✓ **B** Policies and procedures
Correct. Clearly documented policies and procedures tell end users which AI uses are permitted, what data may be entered into AI systems, and which legal, regulatory, and organizational requirements apply. This directly helps reduce AI-related non-compliance.

https://www.comptia.org/content/articles/what-is-artificial-intelligence-ai

✗ **C** Implementing data loss prevention
Incorrect. Data loss prevention can detect or block unauthorized disclosure of sensitive information, including through AI tools, but it addresses data exfiltration rather than the broader set of compliance requirements that end users must understand and follow.

https://www.comptia.org/content/articles/what-is-artificial-intelligence-ai

✗ **D** Enabling multifactor authentication (MFA) for access
Incorrect. MFA strengthens authentication and reduces the risk of unauthorized account access, but it does not instruct end users on acceptable AI use or ensure compliance with AI-related policies and regulations.

https://www.comptia.org/content/articles/what-is-artificial-intelligence-ai

Free demo · full exam: 126 questions

Question 25

Single choice · Domain: D4 AI governance, risk, and compliance › Implement token usage limits to control AI service costs

An AI security administrator receives an inquiry about an unusually high monthly bill from the AI solution provider. The administrator thinks the majority of staff might be using the most powerful model available.

Which of the following AI measures should the administrator implement to lower costs?

A. Storage monitoring
B. Modality types
C. Prompt firewalls
D. Token limits

Answer: D. Token limits

Explanation

This question tests the ability to apply governance controls that limit AI resource consumption and prevent excessive use of costly models. Token limits constrain the amount of text processed or generated, helping control usage-based billing.

✗ **A** Storage monitoring
Incorrect. Storage monitoring can identify storage consumption, but it does not directly limit staff use of an expensive AI model or control token-based processing charges.

https://www.comptia.org/certifications/secai

✗ **B** Modality types
Incorrect. Modality types describe the forms of data an AI system handles, such as text, images, or audio; selecting or monitoring modalities is not the most direct control for excessive model usage and associated costs.

https://www.comptia.org/certifications/secai

✗ **C** Prompt firewalls
Incorrect. Prompt firewalls help inspect, filter, or restrict prompts for security and safety purposes, but they are not primarily designed to cap the volume of model processing that drives the monthly bill.

https://www.comptia.org/certifications/secai

✓ **D** Token limits
Correct. Token limits restrict the amount of input and output processed by the AI service, reducing excessive consumption and helping prevent users from generating unexpectedly high usage charges.

https://www.comptia.org/certifications/secai

Dumpsgate

Question 26

Single choice · Domain: D4 AI governance, risk, and compliance › Establish an AI center of excellence to govern, evaluate, publish, and approve AI usage patterns

A company is adopting AI and wants to create policies and procedures that include a structure for evaluating, publishing, and approving patterns for AI usage.

Which of the following should the company establish to meet this goal?

A. AI center of excellence
B. AI legal affairs office
C. AI audit department
D. AI data science division

Answer: A. AI center of excellence

Explanation

The question tests the governance structure used to create consistent policies, procedures, and approval processes for organizational AI adoption. An AI center of excellence commonly coordinates standards, reusable patterns, evaluation, publication, and approvals across the organization.

✓ **A** AI center of excellence
Correct. An AI center of excellence provides centralized governance for AI adoption, including evaluating, publishing, and approving standardized AI usage patterns.
https://www.comptia.org/certifications/secai

✗ **B** AI legal affairs office
Incorrect. An AI legal affairs office may advise on legal and regulatory issues, but it is not the primary cross-functional structure for evaluating, publishing, and approving reusable AI usage patterns.
https://www.comptia.org/certifications/secai

✗ **C** AI audit department
Incorrect. An AI audit department would assess compliance and controls after or during implementation, but it would not normally serve as the central body for developing and approving organizational AI usage patterns.
https://www.comptia.org/certifications/secai

✗ **D** AI data science division
Incorrect. A data science division focuses on developing and applying models and analytics; it is not necessarily responsible for enterprise-wide AI governance, policy, pattern publication, and approval.
https://www.comptia.org/certifications/secai

Dumpsgate