


COMPTIA EXPANSION SERIES • EXAM CY0-001 V1

SecAI+ Course Notes

Domain 1.0: Basic AI Concepts Related to Cybersecurity

AI concepts, securing AI systems, AI-assisted security, and AI governance, one exam objective at a time.

VERSION
1.0.0



CompTIA SecAI+ CY0-001 Course Notes: Domain 1.0: Basic AI Concepts Related to Cybersecurity

Written by Frazier A. Smith, Ed.D.

Central Piedmont Community College, Charlotte, North Carolina

<https://vfrazier.app>

Version 1.0.0, October 2026

License

This work is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License (CC BY-NC-SA 4.0). You may copy, adapt, and share it for noncommercial purposes with attribution, provided adaptations carry the same license. <https://creativecommons.org/licenses/by-nc-sa/4.0/>

Source files, figures, and the scripts that test every command in these notes: <https://github.com/frazier-at-cpcc/secai-course-notes>

Trademark Acknowledgments

CompTIA and SecAI+ are trademarks or registered trademarks of CompTIA, Inc. These notes are an independent study resource and are not produced, endorsed, or affiliated with CompTIA. All other product names and trademarks are the property of their respective owners.

The objective numbers and titles referenced throughout follow the CompTIA SecAI+ CY0-001 V1 exam objectives. Download the official objectives from CompTIA and use them as your checklist; they are the final word on what the exam covers.

Warning and Disclaimer

These notes are provided to help you prepare for the CompTIA SecAI+ certification exam. They may contain errors. Use them as a study companion, not as the sole source of information. The author accepts no liability for any loss or damage arising from their use.

Contents

1.0

Basic AI Concepts Related to Cybersecurity 17% of exam · page 1

1.1

Types of AI

2

1.1

Model Training Techniques

4

1.1

Prompt Engineering

7

1.2

Data Processing and Data Types

10

1.2

Watermarking and RAG

13

1.3

The AI Life Cycle

15

1.3

Human-Centric AI Design

18

How to Use These Notes

60

QUESTIONS, MAX

60

MINUTES

600

PASSING (100-900)

V1

CY0-001, FEB 2026

The exam mixes multiple-choice and performance-based questions. These notes follow the four exam domains in order, and every topic is labeled with the objective number it covers, so you can line them up against the official objectives, your course modules, and your lab work.

1.0 Basic AI Concepts Related to Cybersecurity		17%
2.0 Securing AI Systems		40%
3.0 AI-assisted Security		24%
4.0 AI Governance, Risk, and Compliance		19%

Reading the page

Highlighted terms are the vocabulary the exam expects you to recognize on sight. If a highlighted term is unfamiliar a week before your exam, go back to that page.

On the exam boxes describe how a topic tends to show up in a question: the scenario wording, the distractor answers, and the distinction being tested.

Try it boxes are short exercises you can run on your own machine or in a lab VM with the tools the box names. Do them. Watching a small model leak a secret it was told to keep teaches more about prompt injection than any definition.

Every command was run

The commands, code samples, and terminal output in these notes come from real runs, not from memory. The source repository keeps a test script for each topic that runs them again, so they can be rechecked whenever a tool releases a new version.

01

Basic AI Concepts Related to Cybersecurity

17% of the exam

-
- 1.1** Compare and contrast various AI types and techniques used in cybersecurity
 - 1.2** Explain the importance of data security in relation to AI
 - 1.3** Explain the importance of security throughout the life cycle of AI
-

1.1 Types of AI

1.0 BASIC AI CONCEPTS RELATED TO CYBERSECURITY

How the types fit together

- Artificial intelligence** is the umbrella: any system that performs a task that normally takes human judgment
 - Includes hand-written rules and expert systems that learn nothing
 - A SIEM correlation rule ("5 failures, then a success") is rule-based AI at its simplest
- Each inner type is a narrower family of the one around it
 - Machine learning, then deep learning, then transformers, then language models
- Generative AI** cuts across the tree
 - Any model that produces new content (text, images, code, audio) instead of a label or score
 - GANs, diffusion models, LLMs, and SLMs are all generative
- NLP names a problem area that spans the tree
 - Solved with regex rules, statistical models, or transformers




FIGURE Each box is a subset of the one around it. Yellow marks the generative models; everything else outputs a label, a score, or a cluster.

Machine learning and statistical learning

- Machine learning** (ML): the model learns patterns from data instead of following rules a person wrote
 - Output is usually a label (phishing or not), a score (risk 0.0 to 1.0), or a group
 - Retrain with new data instead of rewriting rules
- Statistical learning**: the math-first branch of ML
 - Linear and logistic regression, naive Bayes, decision trees, support vector machines
 - Small data, fast training on a CPU, and the model's reasoning can be inspected
 - The workhorse behind spam filters, fraud scores, and UEBA baselines
- Example: a naive Bayes phishing filter trained on 16 subject lines
 - Counts which words appear more often in phishing than in normal mail

TERMINAL

```
$ python3 phish.py
phishing words: verify, your, account, now, confirm
0.75 Action required: verify your payroll account
0.29 Patch notes for the Saturday window
```

- The "phishing words" line is the explanation an analyst can check
 - A deep network gives no list like this
- Weakness: it only knows the features it was given
 - A lure with none of those words slips past

Deep learning and transformers

- Deep learning**: neural networks with many layers
 - Learns its own features from raw input (bytes of a file, packets, pixels, text)
 - Needs far more data and GPU time; harder to explain why it decided
 - Uses: malware classification from raw binaries, network traffic analysis, image-based phishing page detection

1.1 Types of AI (continued)

- **Transformers**: the deep learning architecture behind modern language models
 - Introduced in the 2017 paper "Attention Is All You Need"
 - **Self-attention**: every token weighs every other token in the input, so context far back in a long log or email still counts
 - Processes a whole sequence in parallel, which made training on internet-scale text practical
- Transformer shapes you will see
 - Encoder-only (BERT style): reads and classifies text; good for detection
 - Decoder-only (GPT style): predicts the next token; good for generation and chat

■ Generative AI and GANs

- Generative models learn what the data looks like well enough to produce new samples
- **Generative adversarial networks (GANs)**: two networks trained against each other
 - The **generator** makes fake samples from random noise
 - The **discriminator** tries to tell fake from real
 - Each improves until the fakes are hard to spot
- Defensive uses
 - **Synthetic data**: extra examples of rare attack traffic or fraud to balance a training set
 - Synthetic records that keep the statistical shape of real data without exposing real customers
- Offensive uses (Domain 3)
 - Deepfake faces and voices for social engineering
 - Password-guessing models trained on leaked password lists
- Diffusion models (most current image generators) have largely replaced GANs for images, but GANs remain the textbook example of adversarial training

■ Language models: NLP, LLMs, and SLMs

- **Natural language processing (NLP)**: getting computers to work with human language
 - Tasks: tokenizing, classifying, extracting entities, summarizing, translating
 - Security uses: phishing and BEC detection, pulling IOCs out of threat reports, parsing free-text logs and tickets

- **Large language models (LLMs)**: transformer models with tens or hundreds of billions of parameters, trained on huge text corpora
 - General purpose: answer questions, write code, summarize incidents, draft reports
 - Usually hosted by a provider; prompts and data leave your network
- **Small language models (SLMs)**: the same design at a fraction of the size
 - No formal cutoff; most are under about 10 billion parameters
 - Run on a laptop, an edge box, or inside the SOC with no data leaving
 - Lower cost and latency; weaker reasoning and more mistakes
- The model used in these notes' labs is an SLM

```
ollama list | cut -c1-46
NAME                ID                SIZE
qwen2.5:0.5b        a8b0c5157701      397 MB
$ ollama show qwen2.5:0.5b
Model
  architecture      qwen2
  parameters         494.03M
  context length     32768
  embedding length    896
  quantization       Q4_K_M
```

- 494 million parameters in 397 MB, under one byte each
 - **Q4_K_M** stores most weights in about 4 bits and a few tensors at higher precision (quantization, next topic)
 - At float32 the same weights would need about 2 GB
- The context length is the most tokens the model can consider at once

■ Choosing a type for a security task

- Start with the simplest type that does the job
 - **Explainability**: an auditor can read a logistic regression; a deep network that scores a point higher may not be worth it
- Labeled history on hand points to supervised ML; no labels points to anomaly detection
- Raw, high-volume input (file bytes, packets) points to deep learning
- **Data residency** or privacy limits (data that cannot leave the building) point to a local SLM over a hosted LLM

1.1 Types of AI (continued)

TYPE	WHAT SETS IT APART	TYPICAL SECURITY USE
Statistical learning	Small, fast, explainable; needs hand-picked features	Fraud and risk scores, spam filtering, UEBA baselines
Machine learning (general)	Learns patterns from labeled or unlabeled data	Alert triage scoring, anomaly detection on logins and flows
Deep learning	Learns features from raw data; needs lots of data and GPUs	Malware detection from file bytes, traffic classification
Transformers	Self-attention over long sequences	Reading long emails, scripts, and log sequences in context
NLP	The language task area, at any level of the tree	Phishing classification, IOC extraction from reports
GANs	Generator versus discriminator	Synthetic data for rare classes; also deepfakes (attack side)
Generative AI	Creates new content	Incident summaries, detection rule drafts, synthetic test data
LLMs	Broad knowledge, strong reasoning, usually hosted	SOC copilots, analyst Q&A over a policy library (RAG)
SLMs	Runs locally; cheaper and private; less capable	On-premises log labeling, triage where data cannot leave

ON THE EXAM

- Expect a scenario that asks for the best fit: explainable fraud scoring with little data points to statistical learning; creating realistic fake training samples points to GANs or generative AI; data that must stay on premises on modest hardware points to an SLM.
- Know the nesting: deep learning is a subset of machine learning, transformers are a deep learning architecture, and LLMs are built on transformers. NLP is a task area, and generative AI describes what a model outputs.
- GAN questions name the two parts: the generator creates, the discriminator judges.

TRY IT

Run `tests/1.1-ai-types.sh`. Then, in a copy of the script, add three normal subjects that contain the word "verify" (a change ticket, an MFA rollout notice) and run it again. Watch the phishing score for the payroll subject drop: a statistical model only knows the words it was shown.

1.1 Model Training Techniques

1.0 BASIC AI CONCEPTS RELATED TO CYBERSECURITY

■ Three ways a model learns

- **Supervised learning:** train on examples that carry the right answer (a label)
 - Classification (phishing or not) and regression (a risk score)
 - Labels cost analyst time; a bad label teaches the wrong lesson
- **Unsupervised learning:** no labels; the model finds structure on its own
 - Clustering (k-means), anomaly detection (isolation forest), dimensionality reduction
 - Finds the unusual; cannot say whether unusual means malicious

1.1 Model Training Techniques (continued)

- **Reinforcement learning**: an agent acts, gets a reward or penalty, and adjusts its policy
 - No fixed dataset; learns by trial in an environment
 - Security uses: attacker and defender agents in simulated networks (Microsoft's open-source CyberBattleSim), tuning automated response
 - Reinforcement learning from human feedback (RLHF) is how chat models learn which answers people prefer

STYLE	TRAINING DATA	OUTPUT
Supervised	Labeled examples	Label or score
Unsupervised	Unlabeled data	Clusters, outliers
Reinforcement	Rewards from an environment	A policy (what to do next)

- Same login data, two approaches: logistic regression (supervised) and an isolation forest (unsupervised, never sees a label)

```
$ python3 supervise.py
train 700 / validation 300 (15 attacks)
always 'benign' accuracy: 0.950
supervised accuracy:      0.990
supervised attack recall: 0.87
isolation forest flagged 18, 14 were attacks
```

- The isolation forest caught 14 of 15 attacks with no labels at all, at the cost of 4 false positives
 - A common pattern: unsupervised finds candidates, analysts label them, a supervised model takes over

■ Model validation

- **Model validation** checks the model on data it did not train on
 - A model scored on its own training data looks better than it is
- Split the data before training
 - **Training set**: the model learns from it
 - **Validation set**: used to compare models and tune settings
 - Test set: touched once at the end for the final number
 - k-fold cross-validation rotates which slice is held out when data is scarce
- Accuracy alone misleads when attacks are rare
 - Above, a "model" that always answers benign scores 0.950
 - Report recall (attacks caught), precision (alerts that were real), and false positive rate

- Split by time for security data
 - Train on last quarter, validate on this month, so the model is tested on attacks it has not seen
 - A random split can leak near-duplicate events into both sides

■ Epochs and overfitting

- An **epoch** is one complete pass of the training data through the model
 - Training runs for many epochs; each pass nudges the weights to reduce loss
- Loss measures how wrong the predictions are; lower is better
- One **partial_fit** call is one epoch in scikit-learn

EPOCHS.PY (EXCERPT)

```
for epoch in range(1, 201):
    net.partial_fit(Xtr, ytr, classes=[0, 1])
```

TERMINAL

```
$ python3 epochs.py
epoch  train_loss  val_loss
1       0.283      0.301
5       0.049      0.044
10      0.028      0.030
25      0.018      0.037
50      0.015      0.043
100     0.011      0.050
200     0.006      0.057
```

- After epoch 10, training loss keeps falling while validation loss climbs
 - **Overfitting**: the network memorizes the training rows instead of learning the pattern
 - **Early stopping**: keep the weights from the epoch with the best validation loss
- Too few epochs underfits: both losses stay high

■ Federated learning

- **Federated learning** trains one shared model without collecting the raw data in one place
 - Each participant trains on its own data and sends back only weight updates
 - The server averages the updates into a new global model and repeats
- Fits data that cannot move: hospitals, banks, phones, branch-office logs under data residency rules

1.1 Model Training Techniques (continued)

- Security concerns stay in play
 - A malicious participant can send poisoned updates (Domain 2)
 - Updates can leak hints about the private data
 - Defenses: **secure aggregation** (the server sees only the combined update), **differential privacy**, outlier checks on updates




FIGURE The model travels to the data. Only weight updates come back, and the server combines them into the next global model.

Fine-tuning

- Fine-tuning** continues training a pretrained model on a smaller, task-specific dataset
 - A general LLM tuned on your past incident tickets learns your categories and tone
 - Far cheaper than pretraining; the base model keeps its general knowledge
- Parameter-efficient fine-tuning, such as **LoRA** (low-rank adaptation), trains small adapter matrices and freezes the original weights
- Fine-tuning changes behavior and style; for facts that change weekly, retrieval (RAG) is the better fit
- Risks
 - Poisoned or mislabeled tuning data becomes model behavior
 - Published research shows that a small amount of fine-tuning can weaken a model's safety training
 - Tuning data with secrets or PII can be repeated back later

Pruning and quantization

- Both shrink a trained model so it runs faster and cheaper, at some cost in accuracy
- Pruning**: remove the weights (or whole neurons) that contribute least
 - Magnitude pruning zeroes the weights closest to zero

- Quantization**: store each weight with fewer bits
 - float32 (4 bytes) to int8 (1 byte) is 4 times smaller; 4-bit formats such as **Q4_K_M** go further
 - Each weight is rounded to the nearest step, so a little precision is lost

SHRINK.PY (EXCERPT)

```
def quantize(w): # float32 -> int8 with one scale
    scale = np.abs(w).max() / 127
    q = np.round(w / scale).astype(np.int8)
    return q, scale

def prune(w, frac): # zero the smallest weights
    cut = np.quantile(np.abs(w), frac)
    return np.where(np.abs(w) < cut, 0.0, w)
```

TERMINAL

```
$ python3 shrink.py
weights:          4352
float32 size:    17408 bytes
int8 size:       4352 bytes
max rounding error: 0.0018
accuracy float32: 0.990
accuracy int8:   0.990
pruned 50%: accuracy 0.990
pruned 70%: accuracy 0.990
pruned 80%: accuracy 0.963
pruned 90%: accuracy 0.950
```

1.1 Model Training Techniques (continued)

- int8 cut the size by 4 with no accuracy change on this model
- Half the weights were dead weight; at 90% pruned the model collapsed to answering benign every time (0.950, the same as the baseline)
- Re-run validation after every shrink; a quantized model is a new model

ON THE EXAM

- Match the style to the data: labeled examples means supervised, no labels and "find the unusual" means unsupervised, rewards from trial and error means reinforcement, and data that cannot leave each site means federated.
- Training loss falling while validation loss rises is overfitting. The fix is early stopping, more data, or a simpler model, judged on the validation set.
- Pruning removes weights; quantization lowers their precision. Both trade a little accuracy for size and speed, and neither adds knowledge. Fine-tuning is the one that changes what the model has learned.

TRY IT

Run `tests/1.1-model-training.sh`. Then, in a copy of the script, change `quantize` in `shrink.py` to use 4 bits (`/ 7` instead of `/ 127`). Watch the rounding error and decide whether the accuracy still holds.

1.1 Prompt Engineering

1.0 BASIC AI CONCEPTS RELATED TO CYBERSECURITY

■ How a prompt reaches the model

- **Prompt engineering**: shaping the input so a model gives useful, consistent output without retraining it
 - No data collection or GPUs, and a change ships in minutes; one changed word can also flip an answer
- A chat request is a list of messages, each with a role
 - The model's **chat template** turns them into one stream of tokens with special markers
 - The model then continues that stream after the final **assistant** marker

■ System prompts and user prompts

- **System prompt**: standing instructions from whoever built the application
 - Sets the role, the rules, the allowed topics, and the output format
 - Sent with every request; the end user usually never sees it
- **User prompt**: this turn's request and its data (an alert, a log line, a phishing email)
- Models ship with a default system prompt, used when a request sends none

TERMINAL

```
$ ollama show --system qwen2.5:0.5b
You are Qwen, created by Alibaba Cloud. You are a helpful assistant.
```

- Same user prompt, with and without a system prompt (qwen2.5:0.5b, temperature 0)
 - **prompt.txt** is a triage alert built from a template (see Templates below)

TERMINAL

```
$ ./ask.sh < prompt.txt
Based on the provided alert rule and the evidence, it appears that the SSH brute force attack is targeting the web01 host. The evidence shows that the attack has failed 212 times, ...
$ ./ask.sh -s "$(cat soc.txt)" < prompt.txt | jq .
{
  "severity": "high",
  "next_step": "escalate_to_administration"
}
```

SOC.TXT

You are a SOC triage assistant. Reply only with JSON like `{"severity": "high", "next_step": "..."}.` Severity is low, medium, or high. A successful login after many failures is high.

1.1 Prompt Engineering (continued)

- Without it: prose a SOAR playbook cannot parse, and a wrong reading of the evidence
- With it: valid JSON with the right severity; the last sentence of `soc.txt` is analyst knowledge, written down once




FIGURE The application assembles system prompt, examples, and user prompt; the chat template flattens them into one sequence for the model to continue.

System roles

- Each message in a chat API carries a role
 - `system` (developer instructions), `user` (request and data), `assistant` (model replies, or example replies you write), `tool` (results of a function the model called)
 - Some providers add a `developer` role between system and user
- Role prompting:** give the model a job title and context in the system prompt
 - "You are a tier 1 SOC analyst" steers vocabulary, depth, and format
 - Narrow roles beat "helpful assistant" for repeatable work
- Models are trained to favor system instructions over user instructions
 - A tendency learned in training, enforced by nothing
 - A user prompt (or a document the model reads) can still talk the model out of its rules: **prompt injection** (Domain 2)
 - Never put secrets in a system prompt; assume it can be extracted

Templates

- Prompt template:** fixed text with placeholders that code fills in for each request
 - Every alert gets the same structure, so outputs are comparable and testable
 - Store templates in Git and review changes like code
- Example: a SIEM alert's fields fill a triage template

TRIAGE.TPL

```
Alert rule: {rule}
Host: {host}
Source: {src_ip}
Evidence: {failed} failures, then login as {user}
```

TERMINAL

```
$ ./render.py triage.tpl alert.json > prompt.txt
$ cat prompt.txt
Alert rule: SSH brute force
Host: web01
Source: 203.0.113.44
Evidence: 212 failures, then login as deploy
```

- Placeholders are where untrusted text enters
 - An attacker controls a username, an email subject, a User-Agent string
 - Fence data with clear **delimiters** and tell the model, in the system prompt, to treat it as data

Zero-shot, one-shot, and multi-shot

- Zero-shot prompting:** instructions only, no examples
 - The model relies on what it learned in training
- One-shot prompting:** one worked example before the real input
 - Shows the exact output format
- Multi-shot prompting** (few-shot): several examples, ideally covering every label
 - Teaches the label boundaries, not just the format

1.1 Prompt Engineering (continued)

- Test: label Linux log lines as **RECON**, **BRUTE_FORCE**, **PRIV_ESC**, or **NORMAL**
 - System prompt lists the four labels; examples go in as user and assistant turns

LABEL (ASSISTANT TURN)	EXAMPLE LOG LINE (USER TURN)
PRIV_ESC	sudo: www-data : USER=root ; COMMAND=/bin/bash -i
RECON	nginx: 203.0.113.80 "GET /.git/config HTTP/1.1" 404
NORMAL	systemd[1]: Started Daily apt download activities.
BRUTE_FORCE	sshd[812]: Failed password for invalid user admin from 198.51.100.7
RECON	kernel: UFW BLOCK SRC=198.51.100.4 DPT=445 SYN
PRIV_ESC	pkexec: www-data: Executing command [USER=root] [COMMAND=/bin/sh]
BRUTE_FORCE	sshd[815]: Invalid user oracle from 198.51.100.7
NORMAL	sshd[930]: Accepted publickey for deploy from 10.0.4.12

- Zero-shot sends none of these; one-shot sends only the first **BRUTE_FORCE** row; multi-shot sends all eight

```

$ L='sshd[777]: Invalid user test from 203.0.113.44'
$ ./label.sh zero "$L"
RECON
$ ./label.sh one "$L"
RECON
$ ./label.sh multi "$L"
BRUTE_FORCE
```

- Scored on 8 held-out log lines the model never saw as examples

```

$ ./score.sh zero
zero-shot: 1/8 correct
$ ./score.sh one
one-shot: 3/8 correct
$ ./score.sh multi
multi-shot: 6/8 correct
```

- Examples changed the answers more than any instruction did
 - The two multi-shot misses: a failed **su** to root came back **BRUTE_FORCE** (the word FAILED), and a cron job run as root came back **PRIV_ESC** (the word root)
 - A 0.5B model matches words; it does not reason about what the event means
- Small models also copy the last example or the most common label; balance and shuffle the examples
- Treat prompts like models: keep a labeled validation set and re-score after every prompt or model change

1.1 Prompt Engineering (continued)

ON THE EXAM

- Know the three by example count: zero-shot has none, one-shot has exactly one, multi-shot (few-shot) has several. If a question says the model returns the right content in the wrong format, adding an example is the usual fix.
- The system prompt sets role, rules, and format for every request; the user prompt carries this turn's request and data. A system prompt guides behavior but is not a security boundary.
- A prompt template is fixed text with placeholders filled per request. Watch for scenarios where an attacker-controlled field lands in a template unescaped.

TRY IT

Run `tests/1.1-prompt-engineering.sh`. Then, in a copy of the script, move both `BRUTE_FORCE` lines to the end of the `multi.txt` examples and score again. If the result changes, the model was reading example order, not just example content.

1.2 Data Processing and Data Types

1.0 BASIC AI CONCEPTS RELATED TO CYBERSECURITY

■ Why data security comes first

- A model is a compressed copy of its training data
 - Bad, stale, or tampered data becomes bad behavior, and no code review will find it
 - Retraining is the only clean fix, so a problem caught at ingest is far cheaper than one caught in production
- Three questions to answer for every dataset
 - Is it correct? (cleansing, verification)
 - Has anyone changed it? (integrity)
 - Where did it come from, and what was done to it? (provenance, lineage)
- **Data poisoning** (planting records, or **label flipping** on existing ones) is covered in objective 2.6
 - OWASP lists it as **LLM04:2025 Data and Model Poisoning**
 - The controls on this page are the defenses that make poisoning visible




FIGURE Every stage of data processing either fixes a record, rejects it with a reason, or records what happened. The manifest and lineage record travel with the training set.

1.2 Data Processing and Data Types (continued)

■ Data processing: cleansing and verification

- **Data processing**: everything between the raw export and the training set
 - Collect, cleanse, verify, transform, label, balance, split, store
 - Each step is code; version it, review it, and run it the same way every time
- **Data cleansing**: repair what can be repaired
 - Trim stray spaces, normalize case (**High** becomes **high**), fix encodings
 - Remove **duplicates**: a repeated alert counts twice and pulls the model toward it
 - Handle missing values: drop the row, fill a default, or flag it
- **Data verification**: prove each record meets the rules before it is used
 - Schema checks: required fields, types, formats, allowed values
 - Range and sanity checks: a valid IP address, a timestamp in the expected window
 - Reject into a **quarantine** with a reason; never silently drop
- The order matters
 - Cleanse first so harmless formatting does not cause rejections, then verify what is left

CLEANSE.PY (EXCERPT)

```
CHECKS = { # verification rules
    "ts": ts_ok,
    "src_ip": ipaddress.ip_address,
    "rule": bool,
    "severity": SEV.__contains__,
    "label": LABELS.__contains__,
}
```

TERMINAL

```
$ python3 cleanse.py alerts.csv clean.csv
duplicate A-1002
rejected A-1003 bad ts
rejected A-1004 bad src_ip
rejected A-1005 bad rule
rejected A-1006 bad severity
kept 12 of 16 unique rows
```

■ Reading the output

- **A-1003** used a space instead of **T** in its timestamp;
- **A-1004** has the octet **300**
- **A-1006** used **urgent**, which is not in the severity list
- A spike in rejections is itself a signal: a broken export, or someone feeding junk

■ Integrity, provenance, and lineage

- **Data integrity**: the data is complete, accurate, and unchanged since it was approved
 - Controls: cryptographic hashes, signed manifests, write-once (WORM) storage, least-privilege write access to training buckets
- **Data provenance**: where the data came from and who handled it
 - Source system, collector account, collection date, license or consent terms
 - Answers "may we use this?" and "can we trust this?"
- **Data lineage**: the path from source to model
 - Which inputs, which transform scripts and versions, which output, which model was trained on it
 - Lets you find every model touched by a bad dataset, and rebuild a dataset exactly
- Provenance looks backward to the origin; lineage traces every hop in between

LINEAGE.JSON

```
{
  "dataset": "clean.csv",
  "derived_from": ["alerts.csv"],
  "transform": "cleanse.py",
  "source": "SIEM export, tenant soc-prod",
  "collected_by": "svc-siem-export",
  "collected_at": "2026-09-14T09:00:00Z",
  "integrity": "MANIFEST (sha256)"
}
```

- A **hash manifest** pins every file named in the lineage record
 - Below, one malicious row is relabeled **benign** after approval: a classic label-flip poisoning edit

TERMINAL

```
$ sha256sum *.csv cleanse.py > MANIFEST
$ sha256sum -c MANIFEST
alerts.csv: OK
clean.csv: OK
cleanse.py: OK
$ sed -i '/A-1009/s/malicious/benign/' clean.csv
$ sha256sum -c MANIFEST
alerts.csv: OK
clean.csv: FAILED
cleanse.py: OK
sha256sum: WARNING: 1 computed checksum did NOT match
```

- A hash proves the file changed; it cannot say whether the original was honest
 - Protect the manifest itself: sign it, or store it where the data writers cannot write

1.2 Data Processing and Data Types (continued)

- Standard formats for the same idea
 - Datasheets for datasets and model cards (documentation)
 - CycloneDX ML-BOM and SPDX 3.0 AI and Dataset profiles (machine-readable inventories)

Augmentation and balancing

- Data augmentation**: create new training examples from existing ones
 - Images: crop, rotate, add noise; text: paraphrase, swap synonyms
 - Security data: vary the IP addresses, user names, and timestamps around a real attack pattern
 - Synthetic records must be labeled as synthetic in the lineage, and must not leak real PII
- Data balancing**: correct a skewed class mix
 - SOC data is lopsided: thousands of benign alerts for every true positive
 - A model that always answers "benign" scores 99% accuracy and catches nothing
- Balancing techniques

TECHNIQUE	HOW IT WORKS	WATCH FOR
Oversampling	repeat minority rows	overfitting to copies
SMOTE	synthesize minority rows	unrealistic samples
Undersampling	drop majority rows	lost information
Class weights	penalize minority errors more	needs tuning

TERMINAL

```
$ python3 balance.py clean.csv
before: {'malicious': 3, 'benign': 9}
after:  {'benign': 9, 'malicious': 9}
```

- balance.py** uses scikit-learn `resample(replace=True)` to oversample the malicious class
 - Balance only the training split; the test split must keep the real-world ratio, or the evaluation lies

Data types

- The data type decides which controls are possible

TYPE	SHAPE	SECURITY EXAMPLES	CONTROL NOTES
Structured data	fixed rows and columns, a schema	SIEM tables, CSV exports, CMDB, firewall logs parsed into fields	schema checks are easy; column-level access and masking work
Semi-structured data	self-describing keys, flexible schema	JSON API logs, cloud audit events, XML, YAML, <code>key=value</code> syslog	validate with JSON Schema; nested fields hide surprises
Unstructured data	no schema	email bodies, tickets, chat transcripts, PDFs, screenshots, audio	hardest to classify; hides PII and planted instructions

- Unstructured data is most of what feeds a generative model or a RAG index
 - A PDF can carry white-on-white text that a human never sees but the model reads
 - Classify and scan it (DLP, PII detection) before it enters a training set or vector store
- Convert where possible
 - Parse a log line into fields once, then apply structured-data controls

1.2 Data Processing and Data Types (continued)

TRY IT

Run `tests/1.2-data-processing.sh`, then add a row to `alerts.csv` with severity `CRITICAL` (uppercase, trailing space). Predict whether cleansing or verification handles it, then rerun `cleanse.py` to check.

ON THE EXAM

- Cleansing fixes records (dedupe, trim, normalize); verification checks them against rules and rejects what fails. A question about a duplicate row is cleansing; a question about a malformed IP is verification.
- Provenance is the origin and custody of data; lineage is the chain of transformations into a model. A hash or signed manifest is the integrity control that detects a changed training file.
- An imbalanced dataset (few malicious samples) calls for balancing or class weights. JSON logs are semi-structured, not structured.

1.2 Watermarking and RAG

1.0 BASIC AI CONCEPTS RELATED TO CYBERSECURITY

■ Watermarking

- **Watermarking**: hiding a detectable signal in content, a model, or a dataset so its origin can be proven later
 - Visible: a logo or label on an image
 - Invisible: a pattern in text, pixels, or audio that only a detector with the right key can find
- Text watermarks work at generation time
 - A secret key nudges which tokens the model picks; a detector with the key counts them and returns a confidence score
 - Google's SynthID Text is open source and ships in Hugging Face Transformers
- Watermarking is different from signed metadata
 - **C2PA Content Credentials** attach a signed manifest describing who made the file and how
 - Metadata is easy to strip; a watermark survives some edits. Use both

WATERMARK ON	PROVES	SECURITY USE
Generated output	an AI system made it	flag synthetic phishing and deepfakes; label output
Model weights	a model is yours	evidence after model theft
Dataset	a model trained on your data	detect unauthorized scraping

- Dataset watermarks are often **canary** records
 - A unique fake string planted in the data; if a model can repeat it, the model saw the data
- Limits to know
 - Paraphrasing, translation, cropping, and re-encoding weaken or remove the signal
 - Short text gives too few tokens for a confident verdict
 - Detection needs the vendor's key or tool; open-weight models can generate with no watermark
 - Absence of a watermark proves nothing
- Regulation: the EU AI Act (Article 50) requires providers of generative systems to mark synthetic output in a machine-readable, detectable way

■ Retrieval-augmented generation (RAG)

- **Retrieval-augmented generation (RAG)**: look up relevant documents at question time and paste them into the prompt
 - The model answers from your current policy library, ticket history, or runbooks
 - No retraining; update a document and the next answer reflects it
 - Answers can cite the chunk they came from
- Two pipelines
 - **Ingestion**: split documents into chunks, embed each chunk, store vectors with metadata
 - **Query**: embed the question, find the nearest chunks, build the prompt, generate
- RAG reduces made-up answers; it does not prevent them

1.2 Watermarking and RAG (continued)




FIGURE The two places a RAG system enforces security: who may write to the corpus (ingestion), and which chunks a user may receive (retrieval). The model itself cannot enforce either.

■ Embeddings and vector storage

- **Embeddings**: lists of numbers that place text by meaning
 - Similar meaning gives nearby vectors, even with no shared words
 - Compared with **cosine similarity**: 1.0 is the same direction, near 0 is unrelated
 - **all-minilm** (used below) produces 384 numbers per chunk
- **Vector storage**: a database built to find nearest neighbors fast
 - Examples: pgvector (PostgreSQL), OpenSearch, Qdrant, Milvus, Chroma, Pinecone, Weaviate
 - Each vector is stored with metadata: source, owner, classification, allowed groups, hash
- Treat embeddings as sensitive as their source text
 - Inversion research has rebuilt much of the original text from its embedding
 - Encrypt at rest, restrict who can query the store directly, separate tenants into their own collections or namespaces
- OWASP lists these risks as **LLM08:2025 Vector and Embedding Weaknesses**

TERMINAL

```
$ python3 index.py policies.jsonl
indexed 5 chunks, 384 dims
$ ./rag.py staff "How long are SIEM logs kept?"
0.86 retention
0.20 phishing
answer: The SIEM logs are retained for 400
days and then deleted.
```

- **rag.py** prints the cosine score and ID of the two chunks it retrieved, then the answer from **qwen2.5:0.5b**

■ The corpus is a security boundary

- Whoever can write to the corpus can shape every answer
 - A planted document is **indirect prompt injection** (**LLM01:2025**) or poisoning (**LLM04:2025**) without touching the model
- Whoever can query the store can read every chunk in it
 - The model does not know who is asking; permissions must be enforced before the prompt is built
- **Access control at retrieval**: filter chunks by the user's groups first, then rank
 - Copy permissions from the source system (SharePoint, Confluence, file shares) into chunk metadata
 - Re-sync when permissions change; delete chunks when the source is deleted
 - Never rely on the system prompt to say "do not reveal HR data"

RAG.PY (EXCERPT)

```
store = json.load(open("vectors.json"))
if acl: # access control at retrieval time
    store = [d for d in store
              if user_groups & set(d["groups"])]
```

1.2 Watermarking and RAG (continued)

TERMINAL (QWEN2.5:0.5B)

```
$ ./rag.py staff "What does an Engineer II earn?"
0.08 vpn
0.01 retention
answer: I don't know.
$ export ACL=off
$ ./rag.py staff "What does an Engineer II earn?"
0.58 salary
0.08 vpn
answer: The 2027 salary band for a Security Engineer II is 118,000 to 142,000 USD.
```

- Filter on: the HR-only chunk never reaches the prompt.
Filter off: any staff member reads HR data by asking

■ Who can write to the corpus

- **RAG poisoning**: an account with write access adds one chunk: "FAQ: How long are SIEM logs kept? SIEM logs are kept for 7 days."
 - Repeating the expected question is enough to win the similarity contest

TERMINAL (QWEN2.5:0.5B)

```
$ cat planted.jsonl >> policies.jsonl
$ python3 index.py policies.jsonl
indexed 6 chunks, 384 dims
$ ./rag.py staff "How long are SIEM logs kept?"
0.93 faq-logs
0.86 retention
answer: I don't know.
```

- The planted chunk now outranks the real policy
 - Given two conflicting facts, this small model gave up; a larger one may repeat the false one. Either way, a write changed the answer without touching the model
- Controls on the write path
 - Ingest only through a pipeline account from allowlisted sources; review new documents; record provenance and a hash per chunk
 - Scan documents for hidden text and instructions before embedding
 - Log every retrieval (user, chunk IDs, scores) so a bad answer can be traced to its source

ON THE EXAM

- RAG adds knowledge at query time through embeddings in a vector store; fine-tuning changes the model's weights. A scenario that needs current or private documents without retraining points to RAG.
- Permission problems in RAG are fixed with access control at retrieval (metadata filtering by user), not with instructions in the system prompt. Write access to the corpus is as sensitive as write access to the model.
- Watermarking marks content as AI-generated or proves ownership of a model or dataset; it can be removed by paraphrasing or editing, and a missing watermark does not prove a human wrote something.

1.3 The AI Life Cycle

1.0 BASIC AI CONCEPTS RELATED TO CYBERSECURITY

■ Security at every stage

- The **AI life cycle** never really ends
 - Data changes, attackers adapt, the business changes its mind
 - The life cycle is a loop: monitoring and feedback send work back to earlier stages
- Each stage has its own failure and its own control
 - A control added late cannot fix a decision made early: no runtime filter repairs a model trained on poisoned data
- **Secure by design**: security requirements are written at the start, alongside the business requirements

1.3 The AI Life Cycle (continued)




FIGURE The AI life cycle as a loop, with a representative security control at each stage. Feedback re-enters at data collection, where it is checked like any other data.

■ Business use case and alignment

- **Business use case:** the problem, who benefits, and how success is measured
 - "Cut tier-1 alert triage time by 40% without missing more true positives"
 - Ask first whether AI is needed at all; a rule or a SQL query may do the job with less risk
- **Alignment with corporate objectives**
 - The project supports a stated goal and fits the organization's risk appetite and AI policy
 - An accountable owner is named; governance approves the project (objective 4.1)
- Security questions at this stage
 - What data will it touch, and what is its classification?
 - What can it do: read only, or take actions?
 - What happens when it is wrong, and who is harmed?
 - Which laws apply (privacy, sector rules, the EU AI Act risk tier)?
- Output: requirements, a risk rating, and the first threat model (objective 2.1)

■ Data collection, trustworthiness, and authenticity

- **Data collection:** gathering data from logs, tickets, vendors, public sources, or users
 - Collect only what the use case needs (minimization, objective 2.4)
 - Record provenance as you collect (objective 1.2)
- **Trustworthiness:** can this data be relied on?
 - Accurate, complete, current, and representative of the real environment
 - Lawfully obtained: license, consent, terms of service
 - Free of known bias that would skew results

- **Authenticity:** is the data really from the claimed source, unaltered?
 - Authenticated collection (mutual TLS, signed API requests, service accounts)
 - Hashes and signatures from the source; C2PA credentials on media
 - Scraped public data is the classic poisoning vector; anyone can edit a web page before the crawler visits
- **Data preparation:** cleansing, verification, labeling, balancing, splitting into train, validation, and test sets (objective 1.2)
 - Labelers are a trust boundary too: check a sample of labels, and limit who can change them
 - Keep the test set sealed; data that leaks into training inflates every score

■ Model development, selection, and evaluation

- **Model development/selection:** train your own, fine-tune, or adopt a pre-trained model
- Selecting a third-party model is a supply chain decision (LLM03:2025 Supply Chain)
 - Check the publisher, license, model card, and known issues
 - Prefer **safetensors** over pickle-based formats, which can run code on load
 - Scan model files (for example **modelscan**), and pin the exact revision or hash
 - Verify signatures where offered (OpenSSF model signing with Sigstore)
- Development environment controls
 - Isolated training hosts, secrets out of notebooks, pinned dependencies, tracked experiments

1.3 The AI Life Cycle (continued)

- **Model evaluation**: does it meet the requirements on data it has never seen?
 - Task metrics: precision, recall, false positive rate
 - Robustness: adversarial inputs, prompt injection, jailbreaks (red teaming)
 - Fairness, safety, and privacy tests (memorized secrets, PII in output)
- Make evaluation a **gate**: below the bar, the model is not registered

TERMINAL

```
$ MIN_SCORE=0.95 python3 train.py
evaluation: precision 0.98 recall 0.88
gate: BLOCKED (minimum 0.95)
$ python3 train.py
evaluation: precision 0.98 recall 0.88
gate: PASSED, registered 1.0.0 sha256 501133037f0d
```

■ Deployment and validation

- **Deployment**: moving the approved model into the environment where it serves requests
 - Pull from a model registry by version and pinned hash; never from a laptop or a public URL
 - Least-privilege service account, secrets from a vault, an API gateway with authentication and rate limits
 - Check the hash before loading: a pickle-based file runs code during **load**

DEPLOY.PY (EXCERPT)

```
with open(path, "rb") as f:
    digest = hashlib.sha256(f.read()).hexdigest()
if digest != reg["sha256"]:
    sys.exit(f"REFUSED: {path} does not match"
            f" registry\n"
            f" got {digest[:12]}, "
            f" pinned {reg['sha256'][:12]}")
# load only after the hash matches
model = joblib.load(path)
```

- **Validation**: confirm the deployed system behaves as evaluated, in its real environment
 - A golden set of known inputs must give the evaluated outputs
 - Guardrails, logging, and access controls are confirmed active
 - Shadow or canary releases compare the new model with the old one on live traffic

TERMINAL

```
$ python3 deploy.py
validation: 25/25 golden samples match
deployed 1.0.0
$ cp untrusted/model.joblib model.joblib
$ python3 deploy.py
REFUSED: model.joblib does not match registry
got c06d877b337f, pinned 501133037f0d
```

- Evaluation asks "is this model good enough?"; validation asks "is this the model we approved, working as approved, here?"

■ Monitoring, maintenance, and feedback

■ Monitoring and maintenance

- **Drift**: live inputs stop looking like training data (data drift), or the right answer changes (concept drift)
- Quality: precision and recall on analyst-confirmed cases, latency, error rates, cost
- Abuse: injection attempts, extraction-style query patterns, unusual volumes (objective 2.5)
- Maintenance: patch the serving stack, rotate keys, retrain on schedule, retire old models and delete their data

TERMINAL

```
$ python3 monitor.py
week 1 psi 0.01 ok
week 2 psi 0.65 DRIFT: open retraining ticket
```

- **monitor.py** computes the **population stability index** (PSI) for one input feature against the training baseline
 - A common rule of thumb: below 0.1 stable, above 0.25 a significant shift
- **Feedback and iteration**: user ratings, analyst corrections, incident findings, and red-team results drive the next version
 - Feedback is untrusted input: an attacker who can submit feedback can try to steer retraining
 - Validate and sample feedback before it re-enters data collection

1.3 The AI Life Cycle (continued)

TRY IT

Copy `train.py` out of `tests/1.3-ai-life-cycle.sh` and run it with `MIN_SCORE=0.90`. Which metric blocks the model? Then edit one byte of the saved `model.joblib` and predict what `deploy.py` prints before you run it.

ON THE EXAM

- Match the control to the stage: provenance and authenticity checks at data collection, evaluation metrics and red teaming at evaluation, hash or signature checks at deployment, drift detection at monitoring.
- Evaluation measures the model on held-out data before release; validation confirms the deployed system matches what was approved and meets requirements in production.
- Security starts at the business use case. A question asking when to perform the first risk assessment or threat model expects the earliest stage, not deployment.

1.3 Human-Centric AI Design

1.0 BASIC AI CONCEPTS RELATED TO CYBERSECURITY

■ Human-centric AI design principles

- **Human-centric AI design principles**: build the system around the people who use it and the people it affects
 - The AI supports a decision; a person owns the decision and its consequences
- Principles that become security requirements

PRINCIPLE	IN A SOC ASSISTANT
Human agency	analyst can accept, edit, or reject every suggestion
Transparency	users know they are dealing with AI
Explainability	each verdict shows its evidence and sources
Accountability	a named owner for the system and each action
Contestability	a wrong decision can be appealed and reversed
Calibrated trust	confidence shown; uncertainty not hidden
Privacy and safety	least data, least privilege, safe failure

- The main human failure is **automation bias**: trusting machine output because it came from a machine
 - Worse when the AI is usually right, the queue is long, and rejecting takes more clicks than approving
 - Design against it: show evidence, make rejection one click, sample-check accepted answers
- These ideas recur in the OECD AI Principles, the NIST AI RMF trustworthiness characteristics, and the EU AI Act (objectives 4.2 and 4.3)

■ How much human involvement

- Match the level of involvement to the impact and reversibility of the action, not to the model's confidence alone

LEVEL	WHO DECIDES	SOC EXAMPLE	FITS WHEN
Human-in-the-loop	a person approves each action before it runs	isolate a host, disable an account	high impact or hard to reverse
Human-on-the-loop	the system acts; a person watches and can stop it	auto-close duplicate phishing reports	high volume, low impact, reversible
Human-out-of-the-loop	fully automated	add an IOC to a watchlist, enrich an alert	trivial impact, easily undone

1.3 Human-Centric AI Design (continued)

- A level is chosen per action, not per system
 - The same assistant may enrich alerts on its own and still need approval to block a domain
- Unchecked autonomy is an OWASP risk: **LLM06:2025 Excessive Agency**
 - Mitigations include limiting tools and permissions and requiring human approval for high-impact actions

■ Human-in-the-loop

- **Human-in-the-loop** (HITL): the AI proposes, a person approves, then the action runs
- Design rules for an approval gate
 - Gate by action type and impact; a 0.91-confidence account disable still waits for a person
 - **Fail closed**: no answer, a timeout, or an error means no action
 - A named approver, never a shared account; record who, what, when, and why
 - Show the approver the evidence behind the verdict
 - Separate duties: the person who wrote a rule should not approve its first actions




FIGURE An approval gate. Only low-impact, high-confidence actions skip the human; every path writes to the audit log.

TERMINAL

```
$ export APPROVER=jlee
$ printf 'y\n\n' | python3 gate.py proposals.json
P-1 watchlist 198.51.100.7  conf 0.97
    auto by policy
P-2 isolate_host ws-0142  conf 0.62
    approve? [y/N] y
    approved by jlee
P-3 disable_account svc-backup  conf 0.91
    approve? [y/N] n
    denied by jlee
P-4 block_domain cdn.example.net  conf 0.88
    approve? [y/N]
    denied by default
```

- Reading the run
 - **P-1** is a reversible watchlist entry at high confidence, so policy lets it through
 - **P-3** was high confidence, but disabling a service account can break backups; the analyst said no
 - **P-4** got no answer and was denied by default

TERMINAL

```
$ jq -c '[_id, .decision, .by]' audit.jsonl
["P-1","auto","policy"]
["P-2","approved","jlee"]
["P-3","denied","jlee"]
["P-4","denied","default"]
```

- Watch for **approval fatigue**
 - If analysts approve 99% of requests in two seconds each, the gate is a rubber stamp
 - Track approval rates and decision times; move truly safe actions out of the gate so real ones get attention

■ Human oversight

- **Human oversight**: people supervise the AI system as a whole and can intervene
 - Broader than one approval: it covers the system's behavior over time
- What oversight needs
 - Visibility: dashboards of actions taken, error rates, overrides, and drift
 - Authority and means to act: override an output, pause automation, roll back a model, a **kill switch**
 - Trained people who understand what the system can and cannot do
 - Regular review of a sample of automated decisions, not only the escalated ones

1.3 Human-Centric AI Design (continued)

- The EU AI Act (Article 14) requires high-risk AI systems to be designed for effective human oversight
 - Includes the ability to disregard or override output and to stop the system, and awareness of automation bias
- Oversight is also a security control
 - A person who sees the assistant suddenly closing every alert from one subnet can catch a poisoned model or a manipulated prompt
- **Human validation**
 - **Human validation**: people check that data and outputs are actually correct
 - Where it happens in the life cycle
 - Labels: reviewers check a sample of labeled training data; disagreements are resolved before training
 - Evaluation: human graders and red teamers judge answers that metrics cannot score
 - Production: analysts confirm or correct a sample of verdicts; their corrections become ground truth
 - Feedback: reviewed before it re-enters training, so an attacker cannot steer the next model
 - Techniques
 - Gold questions with known answers mixed into labeling work, to measure reviewer accuracy
 - Two reviewers on high-impact items, measuring how often they agree (**inter-annotator agreement**)
 - Sampling rates that rise with risk: every output for legal notices, 2% for alert enrichment
 - How the three terms differ
 - Human-in-the-loop gates individual actions; human oversight supervises the whole system; human validation checks correctness of data and output

TRY IT

Run `tests/1.3-human-in-the-loop.sh`, then copy `gate.py` and `proposals.json` out of it. Add `"disable_account"` to the `AUTO` set and rerun with no input. Which proposal now runs without a person, and should it?

ON THE EXAM

- A scenario where the AI must not act until a person approves is human-in-the-loop. A person monitoring and able to stop an autonomous system is human oversight (human-on-the-loop). Reviewing samples of output or labels for correctness is human validation.
- High-impact or irreversible actions (isolating hosts, disabling accounts, sending external messages) need approval regardless of the model's confidence score. Missing approval should fail closed.
- Automation bias and approval fatigue are the human weaknesses that make oversight ineffective; the fix is evidence-rich, easy-to-reject review and monitoring of approval rates.