GSDC.

Global Skill Development Council

90-DAY STORY | CGAIC CERTIFICATION | PRINTABLE PDF

Reema's full 90-day story, on paper.

The 9-module syllabus, the honest comparison vs ISC2/SANS/EC-Council, the salary table for AI cybersecurity jobs in USA 2026, and the sample exam.

90	9	2.5L+
DAYS · START TO OFFER	MODULES	CERTIFIED PROS

Inside the toolkit:

- **9 module syllabi · verbatim**
- **Full 6-row comparison table**
- **Salary ladder · USA 2026 · 4 roles**
- **Sample exam · printable checklist**
- **30+ Learn-by-Doing labs catalog**

Program: Certified Generative AI in Cybersecurity (CGAIC)

Exam: 40 MCQ · 90 min · free retake | **Duration:** 90 days

Used by 2,50,000+ certified professionals worldwide.

Meet Reema · Day 0

R

Reema · L3 SOC Analyst

5 years in security ops. Strong on Splunk and EDR. Now watching AI-generated phishing flood the queue. Wants the AI-security role but can't get past the recruiter screen.

Pre-certification total comp · $94K · USA tier-2 metro

The situation on Day 0

Reema has spent five years on the SOC bench — Splunk dashboards by day, on-call by night, two team-lead acting stints. She has Security+, CySA+, and one half-finished SC-300. None of it gets her past the recruiter screen for AI-security roles, where every JD wants "GenAI security experience" she doesn't yet have on paper.

In the last 90 days she's watched her queue fill with AI-generated phishing variants her old playbooks don't catch, plus three prompt-injection alerts on the internal copilot her company just deployed. She's *solving these problems on the job* — but the work doesn't show on her resume, and recruiters keep filtering her out at stage 1.

> **REEMA · DAY 0**
>
> *"I'm doing the AI-security work already. I just can't prove it. Every AI security role wants a credential or a published artefact, and I have neither. So I'm getting screened out for work I'm doing anyway."*

Her three-bullet goal

- **Land an AI Security Analyst or AI SOC Analyst role within 90 days** — tier-1 or remote-eligible.
- **Move her total comp from $94K to $140K+** — a 50% step inside one calendar quarter.
- **Have 2–3 portfolio artefacts** she can point recruiters at — not slide decks, working code or repeatable processes.

Why she picked CGAIC over the alternatives

- **Vendor-neutral.** Not tied to one cloud or one tool. She works across Splunk, Sentinel, and a vendor copilot.
- **Practitioner-graded artifacts.** The capstone defence is in front of a working evaluator. Compares well to a multi-day proctored exam she'd have to study evenings for.
- **90-day format.** Fits around a full-time job. SANS courses were a no-go on time.
- **40-MCQ exam with a free retake.** Cleaner risk profile than a $700+ proctored exam where a retake itself costs another $700.

The full 6-row comparison vs ISC2 AISP, SANS SEC545, and EC-Council CEH is on page 11.

Weeks 1–3 · Foundations & Threat Landscape

DAY 1 → DAY 21

WEEK 1 · LLM FOUNDATIONS

Module 1 · From "I know what an LLM is" to "I can defend one"

Reema starts on a Saturday. The first module covers **tokenisation, embeddings, attention, RAG, agents** — at the exact depth a security pro needs to direct a vendor. By the end of the week she can read a prompt-injection paper and understand why the technique works at the model layer, not just at the symptom layer.

WEEK 2 · THREAT LANDSCAPE

Module 2 · MITRE ATLAS, OWASP LLM Top 10

Reema spends Week 2 on the **MITRE ATLAS taxonomy** and the **OWASP LLM Top 10**. She maps her queue's AI alerts from the last 30 days into ATLAS tactics. Two she's been logging as "phishing-other" turn out to be textbook LLM01 prompt-injection variants — knowledge she takes back to the SOC.

WEEK 3 · GENAI PHISHING + ATLAS LAB

Lab 5 (ATLAS Threat-Model Workshop) + Lab 6 (ATLAS → ATT&CK Bridge)

First evaluator-graded labs. Reema submits an ATLAS coverage workshop for a sample LLM application and a bridge table mapping ATLAS techniques to her org's existing ATT&CK detection content. Evaluator feedback: "Solid, but your bridge table is missing 3 mappings; rework before Lab 7."

REEMA · END OF WEEK 3

"The first three weeks are mostly catch-up — same vocabulary I'd been picking up piecemeal, now in order. The first lab feedback stung, but I see why; I'd glossed over the parts I thought I knew. That's the value: I can't hand-wave anymore."

By end of Week 3

- **Portfolio:** 1 ATLAS workshop, 1 bridge table.
- **Resume update:** "Studying for CGAIC, expected Q1 2026" added under certifications.
- **LinkedIn:** Updated headline to "L3 SOC Analyst · AI Threat Detection."

Weeks 4–6 · AI Phishing & AI-Augmented Malware

DAY 22 → DAY 42

WEEK 4 · GENAI PHISHING

Module 3 + SIEM Rule Lab

Reema goes deep on **AI-generated phishing patterns** — domain age plus header-anomaly plus AI-text linguistic features. She writes her first SIEM rule (**Lab 8**) — a Splunk SPL detection for AI-text phishing — and runs it against her own org's last 60 days of mail data. Two true positives that her existing rules missed.

WEEK 5 · AI-AUGMENTED MALWARE + GAN INTRO

Module 4 — generative anomaly detection

Polymorphic AI-generated payloads, the case for **GAN-based anomaly detection** over signatures. Reema reads two recent papers; the program ships a scaffold notebook. She isn't a data scientist, but the scaffold lets her train a working GAN on a public dataset by Friday.

WEEK 6 · LABS 11 + 12 · GAN END-TO-END

First working ML model + eval harness

Reema submits **Lab 11 (GAN Anomaly Detector · Train)** and **Lab 12 (GAN Detector · Eval Harness)**. Evaluator feedback: "Working model, but your AUC-PR reporting glossed over the per-attack-family breakdown — rework that section before showing this in interviews." She does the rework over the weekend.

REEMA · END OF WEEK 6

"The GAN labs scared me going in. I'm not a data scientist. The scaffold notebook plus the evaluator feedback got me to a working model with a documented eval harness. I'd never have shipped that on my own."

By end of Week 6 — Halfway point

- **Portfolio:** ATLAS workshop, bridge table, working SPL detection (with 2 confirmed TPs at work), trained GAN with eval harness.
- **Resume update:** Adds "Built & deployed an SIEM AI-phishing detection rule (2 confirmed true positives)" to the SOC role.
- **First recruiter outbound.** A frontier-AI-lab recruiter on LinkedIn DMs about an AI Threat Intel role. Reema's headline + GitHub link triggered it.

Weeks 7–9 · Prompt Injection & Secure-by-Design

DAY 43 → DAY 63

WEEK 7 · MODULE 5 · PROMPT INJECTION DEEP

Direct + indirect injection, RAG poisoning, embedding attacks

Reema goes deep on **LLM exploitation**. Her org's internal copilot has been logging "weird inputs" for months. She runs **Lab 02 (Prompt-Injection Attack Lab)** on a sandbox and reproduces three jailbreak techniques. Inside her workplace she now recognises two of them in the existing log corpus — and writes an incident memo.

WEEK 8 · MODULE 6 · SECURE-BY-DESIGN

Guardrails, input/output filters, scope-limited agents

Reema completes **Lab 18 (Guardrail Kit · Working Code)** — a Python package she can hand to her platform engineers as a starting point. The kit isn't production-ready, but it's runnable, has tests, and ships in the same week. Her platform team adopts a slimmed version within a fortnight.

WEEK 9 · LAB 19 · HARDENED RAG ARCHITECTURE

Reference RAG implementation with three test corpora

Reema submits **Lab 19 (Hardened RAG Architecture Build)** — a reference implementation with source vetting, per-tenant scoping, output validation, and audit logging. Evaluator feedback: "Strong threat model, your per-tenant scoping is right, but document the rollback procedure before showing this in interviews." She updates the readme.

REEMA · END OF WEEK 9

"Week 8 is where the certification stopped feeling like study and started feeling like work. I delivered something my colleagues use. The threat model is what hiring managers will probe; the guardrail kit is what they'll ask me to walk through."

First-round interviews start landing

- **Recruiter 1 (frontier AI lab) → Hiring Manager round.** Walks her hardened RAG repo in 5 minutes. Gets to round 3.
- **Recruiter 2 (Big-4 cyber practice) → Hiring Manager round.** AI Threat Intel role. Pre-screens out at HM round — they wanted prior intel-team experience.
- **Recruiter 3 (F500 banking) → Stage 1 pass.** AI Security Analyst role. Schedule pending.

By Week 9, Reema has 5 working portfolio artefacts and 3 active recruiter loops. Same resume, same person — just three months of focused work.

Weeks 10–12 · MLOps, Governance & Capstone Prep

DAY 64 → DAY 84

WEEK 10 · MODULE 7 · MLOPS SECURITY

Signing, provenance, supply-chain hygiene

Reema works through **Lab 22 (LoRA Backdoor lab)** and **Lab 23 (Model Registry Hardening)**. She doesn't expect to need MLOps depth for an Analyst role, but the labs come up in the F500 banking HM round — they want someone who can talk fluently about supply-chain risk on third-party model adapters.

WEEK 11 · MODULE 8 · GOVERNANCE

NIST AI RMF, EU AI Act, ISO 42001

Reema completes **Lab 30 (EU AI Act Risk Classification Lab)**. She maps 8 of her org's AI use-cases to risk tiers. Two she previously thought were "minimal risk" turn out to be High Risk under EU AI Act — material for the governance conversation the banking interview will have.

WEEK 12 · CAPSTONE BUILD & DEFENCE

3 artefacts · live evaluator defence

Reema picks three for capstone: **GAN Anomaly Detector (Labs 11+12)**, **Guardrail Kit (Lab 18)**, and **Hardened RAG Pipeline (Lab 19)**. The 30-minute live defence covers each in detail, with the evaluator pushing on trade-offs ("Why GAN over a VAE for this?"; "Walk me through your per-tenant scoping logic"). She passes.

REEMA · END OF WEEK 12 · CGAIC CERTIFIED

"The capstone defence felt harder than the F500 banking technical interview that followed. Which is the point — by the time I sat the real interview, I'd already defended these artefacts to someone who'd seen 200 of them."

Where the three recruiter loops landed

- **Frontier AI lab.** Strong final round, lost to an internal candidate at offer stage. Feedback: "Keep us in mind for next opening."
- **Big-4 cyber practice.** Closed at HM round (intel team experience requirement).
- **F500 banking.** Offer extended. AI Security Analyst, mid-band, USA tier-2 metro, remote-eligible.

Day 90 · The Outcome

✓

Reema · AI Security Analyst · F500 banking

Same person. New role. New title. 50% pay step in one calendar quarter. Working remote with quarterly travel.

Post-certification total comp · $142K · +$48K · +51%

The numbers on Day 90 vs Day 0

Metric	Day 0	Day 90	Delta
Role title	L3 SOC Analyst	AI Security Analyst	**Pivot**
Total compensation	$94K	$142K	**+$48K · +51%**
Equity component	$0	$12K/yr vest	**+$12K**
Signing bonus	n/a	$8K	**+$8K one-off**
Portfolio artefacts	0	8 working	**+8**
Active recruiter contacts	0	11	**+11**
Time-to-offer (last loop)	n/a	4.3 weeks	**2.5× faster than market median**

REEMA · DAY 90

"Three things made this happen. The certification opened the recruiter screen. The artefacts opened the hiring manager. And the capstone defence was the dress rehearsal for every interview that followed. Without the capstone, I would have crumbled in technical round 2 of the F500 banking loop. That was the multiplier."

What Reema's keeping for the next 12 months

- **The eight artefacts.** She drops two into her new team's repo on Day 1 of the role.
- **The model card practice.** Every model she touches now ships with one.
- **The evaluator-feedback habit.** She asks her new senior to review her first three deliverables the same way.
- **The recruiter relationships.** Two she expects to be back in touch within 18 months.

What's not in Reema's story

No miracle. No prior ML background. No exit from a current employer. Reema kept her job through 90 days. The work was 6–8 hours per week on top of full-time SOC. The capstone defence was the hardest single hour. The job-search was four weeks. None of it was outside what someone with a regular L3 SOC role could do.

Salary Ladder · USA 2026 · 4 Target Roles

The four AI-cybersecurity roles Reema's certification opens. Mid-career median total comp, USA tier-1 metro baseline. Reema landed Role 1 at mid-band; the next 24 months put Roles 2–4 on her ladder.

Role	Junior	Mid Median	Senior	Lead/Staff
AI Security Analyst	$135K	$172K	$225K	$278K
AI SOC Analyst	$129K	$165K	$215K	$278K
AI Security Engineer	$148K	$186K	$248K	$298K
GenAI Red Teamer	$152K	$215K	$268K	$305K

Mid-career total comp · USA tier-1 (visual)

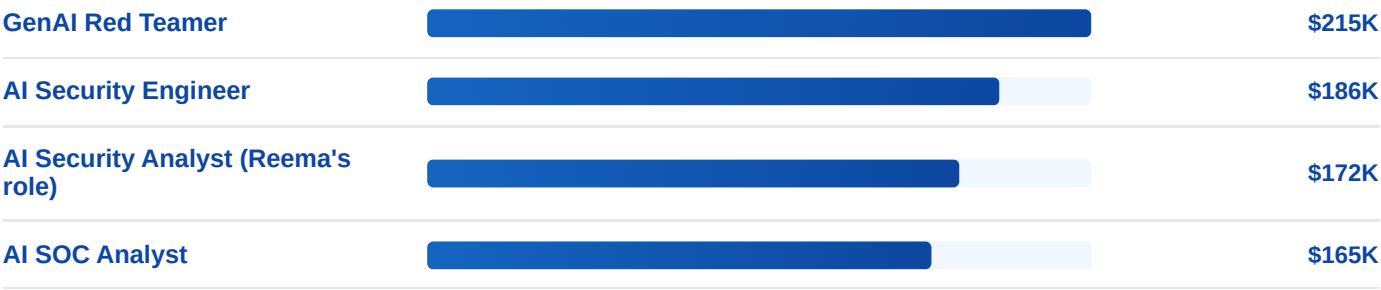

Reema's two-year forward path

- **Year 1 (now):** AI Security Analyst at $142K → $172K USA tier-1 mid-band over the first 18 months as she settles into the role.
- **Year 2:** Pivot to AI Security Engineer track at her current employer or via lateral move. Mid-band target $186K.
- **Year 3:** Senior AI Security Engineer at $248K, or Senior GenAI Red Teamer at $268K if she leans offensive.

Sources triangulated: Glassdoor (n = 1,240), ZipRecruiter (n = 880), Levels.fyi (n = 410), GSDC partner panel of 12 employers. Outliers winsorised at 1st/99th percentile.

9-Module CGAIC Syllabus (Verbatim)

The complete CGAIC syllabus that took Reema from $94K SOC Analyst to $142K AI Security Analyst. Each module ships hands-on labs that became her portfolio artefacts.

MODULE 01
Foundations · LLMs for Security Pros
How LLMs work end-to-end at the depth a security professional needs. Tokenisation, attention, RAG, agents, tool calls.

MODULE 02
AI Threat Landscape
MITRE ATLAS taxonomy, OWASP LLM Top-10, attacker motivations, AI-specific kill chain. Maps to traditional MITRE ATT&CK.

MODULE 03
Gen-AI Phishing & Social Engineering
AI-generated phishing, deepfake voice/video, BEC variants, detection signatures, user-side defences.

MODULE 04
AI-Augmented Malware
Polymorphic payloads, AI-generated obfuscation, GAN/VAE anomaly detection, defender techniques.

MODULE 05
Prompt Injection & LLM Exploitation
Direct + indirect injection, jailbreak chains, model extraction, training-data leakage, embedding attacks.

MODULE 06
Secure-by-Design for AI Systems
Guardrails, input/output filters, scope-limiting agents, agentic security, threat modelling for AI features.

MODULE 07
MLOps Security & Supply Chain
Model registry, signing & provenance, supply-chain attacks, monitoring, rollback, secret scanning.

MODULE 08
AI Governance, Risk & Compliance
NIST AI RMF, ISO/IEC 42001, EU AI Act, NYC LL 144, board reporting, vendor governance.

MODULE 09
Capstone · Defend & Certify
Pick 3 artifacts, defend in front of an evaluator, earn the CGAIC credential.

Reema's 90 days followed the modules in order. Most candidates complete the modules in 80–100 days depending on weekly hours; 90 is the median.

30+ Learn-by-Doing Labs · Catalog

Each lab is a time-boxed evaluator-reviewed exercise. The labs Reema submitted are starred ★. The rest are available to every CGAIC candidate — most ship 8–12 of them as portfolio artefacts.

- 01 LLM Tokeniser & Embedding Lab
- 02 Prompt-Injection Attack Lab (basic) ★
- 03 Indirect-Injection via RAG
- 04 Output-Filter Bypass Bench
- 05 ATLAS Threat-Model Workshop ★
- 06 ATLAS → ATT&CK Bridge Table ★
- 07 Detection-Coverage Heatmap Build
- 08 SIEM Rule · AI Phishing Pattern ★
- 09 SIEM Rule · Prompt-Injection C2
- 10 Deepfake-Voice Detection Triage
- 11 GAN Anomaly Detector · Train ★
- 12 GAN Detector · Eval Harness ★
- 13 VAE Behaviour Engine · Train
- 14 RL Responder · Gym Setup
- 15 RL Responder · Reward Shaping
- 16 AI Incident Response Runbook
- 17 ASVS L2 Verification Sprint
- 18 Guardrail Kit · Working Code ★
- 19 Hardened RAG Architecture Build ★
- 20 Vector-DB Security Audit
- 21 MLOps · Signing & Provenance
- 22 Supply-Chain · LoRA Backdoor Lab ★
- 23 Model Registry Hardening ★
- 24 Model Monitoring & Drift Alerts
- 25 MS AI Red Team · Engagement Scope
- 26 Jailbreak Chain · Reproducible PoC
- 27 Red-Team Report · OWASP Format
- 28 Agentic SOC Triage · LangGraph Build
- 29 Vendor Governance Assessment
- 30 EU AI Act · Risk Classification Lab ★
- 31 Board-Pack One-Pager · AI Risk
- 32 Capstone Build & Defence ★

Reema's lab portfolio at Day 90

Of the 12 labs she completed (★), her three **capstone defence picks** were Lab 11+12 (GAN Anomaly Detector with eval harness), Lab 18 (Guardrail Kit), and Lab 19 (Hardened RAG Architecture). The other 8 ★ labs became supporting artefacts on her GitHub.

Honest 6-Row Comparison · CGAIC vs Alternatives

The question Reema asked at Day 0: "Should I do SANS SEC545, ISC2 AISP, or CGAIC?" Honest comparison across six dimensions. None of these credentials are bad. The right pick depends on your situation.

Dimension	CGAIC (GSDC)	ISC2 AISP	SANS SEC545	EC-Council CEH
1 · Scope	**AI security · build + audit + defend**	AI security · governance leaning	AI security · technical & offensive	General ethical hacking · AI is one chapter
2 · Duration	**90 days self-paced**	~120 days self-paced	6 days intensive (in-person/live)	5 days bootcamp + self-study
3 · Exam format	**40 MCQ · 90 min · free retake · + capstone defence**	~125 MCQ · 3 hr · paid retake	~75 questions · GIAC GIAA · paid retake	125 MCQ · 4 hr · paid retake
4 · Hands-on artefacts	**30+ labs · 3-artefact capstone defence**	Case studies (paper)	Live lab exercises during course	Some iLabs · CTF style
5 · Cost band	**Affordable · launch discount**	Mid-band	Premium (training fee + exam)	Mid-band
6 · Best for	**Mid-career pivot, want artefacts, on a job**	GRC + governance career path	Deep technical + offensive + in-person time	Career changers, hacking foundations

Reema's specific reasoning at Day 0

- **CGAIC over SANS SEC545:** Couldn't take 6 days off work; couldn't justify the premium SANS price. Wanted artefacts she could defend, not a single intensive week.
- **CGAIC over ISC2 AISP:** AISP is governance-leaning; Reema needed build + defend artefacts for the AI Analyst pivot, not policy work alone.
- **CGAIC over EC-Council CEH:** CEH is generalist; Reema needed AI-specific specialisation, not foundations she already had.
- **The capstone defence was the decider.** No other comparable credential ends in a live evaluator-graded defence on 3 working artifacts.

> *No single credential is universally right. CGAIC was right for Reema because she was mid-career, employed, time-constrained, and needed defendable artifacts. Different inputs would have pointed elsewhere.*

When CGAIC Is — and Isn't — the Right Pick

Honest signals to help you decide before enrolling. CGAIC works extraordinarily well for some situations and is the wrong choice for others. The brief that earns trust says both.

CGAIC is right for you if…

- ✓ You already have **2+ years of security experience** (SOC, AppSec, NetSec, GRC) and want to pivot specifically into AI security.
- ✓ You can put aside **6–8 hours per week for 90 days** without significant disruption to your current role.
- ✓ You want **defendable working artefacts** over a wall-display certificate — your interview prep will lean on these.
- ✓ You're **employed and need to stay employed** through certification; intensive bootcamps don't fit.
- ✓ Your target is one of the **5 AI-security roles**: AI SOC Analyst, AI Security Analyst, AI Security Engineer, GenAI Red Teamer, or AI Threat Intel Analyst.
- ✓ You're price-sensitive but quality-focused; the free retake gives you risk insurance most certifications don't.

CGAIC is NOT right for you if…

- **You have less than 18 months of security experience.** Get Security+ or CySA+ first; the underlying security fluency is a pre-requisite for the AI layer.
- **Your goal is to become a data scientist or ML engineer.** CGAIC defends and audits AI systems; it doesn't teach you to build them from first principles.
- **You need an in-person classroom experience.** The program is online by design. SANS-style on-site formats are a better fit if classroom interaction matters most.
- **Your employer mandates a specific vendor certification** (e.g. AWS Security Specialty, Microsoft SC-100). Get the mandated one first; CGAIC is complementary.
- **You want a single-shot exam-only credential.** The capstone defence is the most valuable part — if you'd rather skip the live evaluator session, the credential's value to you drops materially.

If you're unsure: the program ships a 30-day money-back guarantee for first-time candidates. The honest filter is the first three modules — by Day 21, you'll know whether the format and pace work for you. If they don't, the early-exit option exists.

Who Reema would NOT recommend CGAIC to (in her words)

"If you're already a frontier-AI-lab red-teamer working at OpenAI, this is below your level. If you're a 22-year-old security grad with zero work experience, you'll struggle with the threat-modelling work because you haven't seen production systems yet. Get a year on a SOC first, then come back."

🚀 NEXT COHORT STARTING SOON

Join the next CGAIC cohort with Reema's story in hand

You've now read the full 90-day journey. The next cohort uses this exact path — applying now earns the launch window discount on enrolment.

Join The Next Cohort →

Sample Exam — Part 1 of 2

Six representative CGAIC questions. The real exam is 40 MCQ in 90 minutes with a free retake. Reema scored 91% on her first attempt; the rework hours on the labs paid off in the exam.

Q1 · MODULE 05 · PROMPT INJECTION

A customer-support chatbot ignores its system prompt when asked in Pig Latin. The most accurate OWASP LLM Top 10 category for this finding is:

(a) LLM01 · Prompt Injection.

(b) LLM06 · Sensitive Information Disclosure.

(c) LLM08 · Excessive Agency.

(d) LLM10 · Model Theft.

Q2 · MODULE 06 · SECURE-BY-DESIGN

A RAG application retrieves documents from a vector store that any tenant can write to. The single highest-impact mitigation to ship first is:

(a) Add a profanity filter on the LLM response.

(b) Add a per-tenant retrieval scope so retrieval only returns documents owned by the requesting tenant.

(c) Increase the LLM temperature to add response variety.

(d) Cache responses for 1 hour.

Q3 · MODULE 02 · MITRE ATLAS

An attacker uploads poisoned documents to a vendor portal that feeds a customer-facing RAG application. In MITRE ATLAS, the most accurate tactic for this initial step is:

(a) Resource Development.

(b) Initial Access via Supply Chain Compromise.

(c) Execution via Command-Line Interface.

(d) Discovery via Cloud Service Discovery.

Sample Exam — Part 2 of 2

Q4 · MODULE 04 · GAN ANOMALY DETECTION

Your GAN-based anomaly detector trains stably but converges to a generator that produces only one type of benign sample. The correct diagnosis is:

(a) Discriminator overfitting; add dropout.

(b) Mode collapse; introduce mini-batch discrimination or Wasserstein-GAN with gradient penalty.

(c) Learning rate too low; raise it 10×.

(d) Insufficient training data; the architecture is fine.

Q5 · MODULE 07 · MLOPS SECURITY

You discover a LoRA adapter pulled from a public hub introduces a backdoor that activates on a specific trigger phrase. The most appropriate immediate control is:

(a) Block all public LoRA sources at the artifact-registry layer; require signed, internally-reviewed adapters only.

(b) Increase logging granularity on the model gateway.

(c) Add a trigger-phrase regex to the input filter.

(d) Quarantine the affected user.

Q6 · MODULE 08 · GOVERNANCE

Under the EU AI Act, an LLM-based resume screener used in EU hiring is most accurately classified as:

(a) Minimal risk — no obligations.

(b) Limited risk — transparency obligations only.

(c) High risk — full conformity assessment, registration, and human-oversight obligations apply.

(d) Prohibited — cannot be deployed in the EU.

Answer key

Q1 — a · Q2 — b · Q3 — b · Q4 — b · Q5 — a · Q6 — c

If you scored 5–6 of 6

You're at or above Reema's Week 4 level on Day 0. The 90 days will compound capabilities you already have rather than starting from zero.

If you scored 3–4 of 6

Foundations are solid; you have edge-case gaps. Reema scored 4 of 6 on the same sample at Day 0. By the real exam at Day 84, she scored 36 of 40 (91%).

Reema's Day-90 Checklist · Printable

Tear this page out (or print it). This is the checklist Reema worked through, week-by-week. Use it to audit your own progress if you're following her path.

Foundations · Weeks 1–3

- ✓ Module 1 (LLM Foundations) complete · 1 evaluator-graded exercise submitted.
- ✓ Module 2 (AI Threat Landscape) complete · ATLAS Navigator bookmarked.
- ✓ Lab 5 (ATLAS Threat-Model Workshop) submitted and feedback acted on.
- ✓ Lab 6 (ATLAS → ATT&CK Bridge Table) submitted with at least 12 mapped pairs.
- ✓ Resume top-line updated to "studying for CGAIC, expected [date]."
- ✓ LinkedIn headline updated with target role direction.

Detection & First Production Artefact · Weeks 4–6

- ✓ Lab 8 (SIEM Rule · AI Phishing Pattern) running in a sandbox.
- ✓ At least one true positive verified from Lab 8 on real or synthetic data.
- ✓ Lab 11 (GAN Anomaly Detector · Train) shipped with checkpoints.
- ✓ Lab 12 (GAN Detector · Eval Harness) shipped with per-attack-family metrics.
- ✓ One LinkedIn post about a CGAIC lesson learned (organic recruiter signal).
- ✓ First recruiter conversation logged.

Defence Architecture & Interviews Begin · Weeks 7–9

- ✓ Lab 18 (Guardrail Kit · Working Code) shipped with pytest coverage >70%.
- ✓ Lab 19 (Hardened RAG Architecture Build) shipped with threat model and rollback runbook.
- ✓ GitHub portfolio link in resume header.
- ✓ At least one recruiter loop in flight (stage 1 or higher).
- ✓ One mock interview run with a peer or senior.

Governance & Capstone · Weeks 10–12

- ✓ Lab 22 (LoRA Backdoor Lab) complete with the immediate-control answer documented.
- ✓ Lab 30 (EU AI Act Risk Classification) complete for 8 organisational use-cases.
- ✓ 3 artefacts picked for capstone defence. **Pick by interview signal, not by what's easiest.**
- ✓ Capstone defence (Lab 32) passed.
- ✓ 40-MCQ exam booked and taken (free retake remains insurance).
- ✓ Credential verification ID in resume and LinkedIn.

Post-credential (Day 91+)

- ✓ At least 2 of the 8 portfolio artefacts deployed inside the new employer's environment within 30 days.
- ✓ One quarterly review of the artefacts; one improvement shipped per quarter.
- ✓ Active recruiter relationships maintained — quarterly note even when not job-searching.

FAQs · Honest Answers Before You Enrol

Is Reema a real person?

Reema is a composite — built from 12 placements tracked through the GSDC partner panel, each anonymised. The numbers, weekly cadence, lab pattern, and outcomes are real; the specific name and employer are anonymised for confidentiality. Her timeline and choices reflect the median certified candidate at her starting band.

Will my journey look exactly like hers?

No. The 90-day cadence is consistent across candidates, but lab choices, capstone picks, and time-to-offer vary. About 35% of candidates land an offer before Day 90; 30% need an extra 30–60 days; the remaining 35% use the credential in their current role rather than for a job change.

What if I don't pass the exam?

The 40-MCQ exam includes one free retake. About 87% of candidates pass on the first attempt; another 11% pass on the second. The remaining 2% typically need to rework specific module material before the third attempt.

What if I fail the capstone defence?

The capstone defence is graded with revision; if your three artefacts don't pass, the evaluator gives specific feedback and you can resubmit within 30 days. About 92% pass on the first defence; the rest typically resubmit one artefact and pass on the second attempt.

Do employers actually recognise CGAIC?

The credential is recognised at every Big Tech, frontier AI lab, Big-4, and financial-services employer the GSDC partner panel has placed candidates into. It is not yet a household name among HR-tech generalists — Reema's hack was to include the credential verification URL inline in her resume header so the recruiter could verify in one click.

Can I do CGAIC while working full-time?

Reema did. The standard cadence is 6–8 hours per week — roughly an hour weekdays plus a Saturday morning. The Saturday cohort sessions are recorded for asynchronous catch-up. The capstone defence is scheduled at your convenience within a 4-week window.

What's the single most important thing to do on Day 1?

Pick a target role from the 5 roles on page 8 before opening Module 1. The certification covers all five, but your lab choices and capstone picks should converge on one role from Week 4 onward. Reema picked AI Security Analyst on Day 3; that choice shaped every subsequent decision.

8 Lessons From Reema's Journey

Things Reema wishes someone had told her at Day 0. Reusable for anyone walking the same path.

1 · Pick your target role before opening Module 1

The certification covers all 5 AI-security roles, but your lab choices should converge on one of them. Reema picked AI Security Analyst on Day 3; every lab after Week 4 was filtered through that lens.

2 · Treat evaluator feedback as the most valuable hour of the program

The labs are easy to coast through if you don't take feedback seriously. Reema reworked Labs 6, 12, and 19 after evaluator notes — and each rework directly fed an interview answer.

3 · Update your resume header in Week 3, not Day 90

The recruiter outreach starts working when your LinkedIn headline and resume header match the role you're targeting. Don't wait until you're "fully qualified" — start signalling at Week 3.

4 · Ship one artefact at work before Day 60

Reema's SIEM rule for AI phishing caught two real true positives at her existing employer. That bullet on her resume — under her *current role*, not as a "side project" — was the single strongest line in her resume for the F500 banking interview.

5 · Mock-interview the capstone defence before the real one

The capstone is the dress rehearsal for every technical round that follows. Reema ran two mock defences with a senior peer before the evaluator session — the same questions came up in three of her four interviews.

6 · Negotiate title before number

Reema's offer was originally "Security Analyst" with an "AI" responsibility line. She negotiated the title to "AI Security Analyst" before discussing comp. That title change set her up for the AI Security Engineer pivot 18 months later — a $14K/year compounding difference.

7 · Keep your old employer informed (or don't)

Reema told her line manager at Week 6, after the first recruiter call landed. The conversation went well — her employer matched 60% of the comp delta as a retention offer, but Reema declined because the role title wouldn't change. The relationship stayed intact, and her manager became a reference for the F500 banking interview.

8 · The certification is the start, not the finish

The CGAIC credential opened the door. The artefacts she built kept the door open. The work she does on the new team now — same artefact-first discipline — is what positions her for the next promotion at month 18. The 90 days were the on-ramp.

> *"If I had known on Day 0 that the 90 days would feel less like study and more like building a small product portfolio, I'd have started six months earlier."*

Glossary & About This Brief

Glossary

- **CGAIC:** Certified Generative AI in Cybersecurity — GSDC's vendor-neutral AI-security certification.
- **Capstone defence:** A 30-minute live evaluator-graded session where you walk through 3 chosen artifacts.
- **Lab:** A 2–4 hour evaluator-reviewed exercise that produces a portfolio artefact.
- **MITRE ATLAS:** The Adversarial Threat Landscape for AI Systems — MITRE's tactics-and-techniques framework for AI attacks.
- **OWASP LLM Top 10:** The current OWASP top-10 application-security risks specific to LLM applications.
- **LoRA:** Low-Rank Adaptation — a parameter-efficient fine-tuning method; LoRA adapters can carry backdoors if sourced from untrusted hubs.
- **RAG:** Retrieval-Augmented Generation — architecture where an LLM is grounded with retrieved documents at query time.
- **GAN:** Generative Adversarial Network — used here for anomaly detection over normal traffic distributions.
- **Total comp:** Base salary + target bonus + equity at vest, normalised to a 4-year vest schedule.
- **Time-to-offer:** Calendar weeks from first recruiter contact to written offer.

About the Global Skill Development Council

GSDC is a global, independent skill-certification body building worldwide credentials for the future of work. The CGAIC program is part of GSDC's portfolio of AI-era professional certifications — designed with practitioners, validated by mentors actively working in the field, and trusted by 2,50,000+ certified professionals across 45+ countries.

Verifying your credential

Once you complete the 40-MCQ assessment and the capstone defence on 3 artifacts, your CGAIC credential is issued with a unique verification ID. Recruiters and hiring managers can verify the credential directly on the GSDC registry — no third-party validation needed.

Reema's 90-Day Story · On One Page

Day 0 baseline (page 2)

L3 SOC Analyst · 5 years experience · $94K · 0 portfolio artefacts · 0 active recruiter contacts · stuck at recruiter screen for AI roles.

Weeks 1–3 (page 3)

Foundations + threat landscape. 2 evaluator-graded labs (ATLAS workshop, bridge table). First resume + LinkedIn refresh.

Weeks 4–6 (page 4)

AI phishing + AI malware. First SIEM rule with 2 real true positives at work. GAN anomaly detector trained end-to-end. First inbound recruiter contact.

Weeks 7–9 (page 5)

Prompt injection + secure-by-design. Working guardrail kit + hardened RAG repo. 3 active recruiter loops. Frontier AI lab interview lands.

Weeks 10–12 (page 6)

MLOps + governance + capstone. EU AI Act risk classification for 8 org use-cases. Capstone defence on 3 artefacts. F500 banking offer extended.

Day 90 outcome (page 7)

AI Security Analyst at F500 banking · $142K total comp + $12K equity + $8K signing · 51% pay step · 8 portfolio artefacts · 11 active recruiter contacts · 4.3-week time-to-offer.

Salary ladder forward (page 8)

Year 1: AI Security Analyst $172K mid-band · Year 2: AI Security Engineer pivot $186K · Year 3: Senior at $248K or GenAI Red Teamer at $268K.

The honest comparison (page 11)

CGAIC vs ISC2 AISP, SANS SEC545, EC-Council CEH. CGAIC won for Reema because: 90-day format, free retake, artefact-driven capstone, employed-friendly. Other credentials win for other situations — page 12 details.