GSDC.
Global Skill Development Council

PATHWAY ROADMAP | CGAIC CERTIFICATION | PRINTABLE PDF

The full pathway roadmap brief.

All 4 persona roadmaps as 90-day printable plans. Per-pathway capstone picks, target employers, salary band, and the matching CGAIC modules. Plus the 9-module syllabus and sample exam.

4	9	2.5L+
PERSONAS	MODULES	CERTIFIED PROS

Inside the toolkit:

- **4 persona 90-day roadmaps**
- **Per-pathway capstone picks (3 artifacts)**
- **Per-pathway target employers**
- **Per-pathway salary band**
- **9 module syllabi · verbatim**

Program: Certified Generative AI in Cybersecurity (CGAIC)

Format: CGAIC Certification | **Duration:** 90 days

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

The 4 AI Cybersecurity Pathways · At a Glance

Four distinct career destinations from the same CGAIC certification. Pick one before you start the 90-day plan. The pathways differ in starting profile, depth of code, and end-state employer mix.

Pathway 1 · AI SOC Analyst

Best for: SOC analyst, threat intel analyst, junior IR.

Depth of code: Light — SIEM queries, Python scripting basics.

End state: Detection & triage of AI-powered attacks at L2/L3.

Pathway 2 · GenAI Red Teamer

Best for: Pentester, red teamer, offensive security engineer.

Depth of code: Heavy — prompt injection, jailbreak, LLM exploitation.

End state: Lead GenAI red-team engagements end-to-end.

Pathway 3 · AI Security Engineer

Best for: Security engineer, AppSec / DevSecOps.

Depth of code: Heavy — MLOps, secure model deployment, guardrails.

End state: Build & harden production AI systems.

Pathway 4 · AI Governance Lead

Best for: GRC, risk, compliance, audit, policy.

Depth of code: Light — frameworks, controls, regulator-facing.

End state: Own AI risk & compliance for the enterprise.

How to pick your pathway in 60 seconds

- **Currently in a SOC?** → Pathway 1. Your detection muscles transfer directly.
- **Currently breaking things for a living?** → Pathway 2. Red-team mindset is the prerequisite, not the goal.
- **Currently shipping code or platforms?** → Pathway 3. You'll harden what you already build.
- **Currently writing policy, managing risk, or in audit?** → Pathway 4. The board-facing pathway, lowest code requirement.

> *Pick one pathway up front. Trying to do all four at once is the single biggest reason candidates miss the 90-day window.*

Per-Pathway Salary Bands & Employer Mix

Mid-career total comp at USA tier-1 metro baseline. Use regional adjusters at the bottom for local market. All bands reflect 2026 placements with the CGAIC credential held.

Pathway	Mid (4–7 yrs)	Senior (8–12 yrs)	Director+ (12+ yrs)	Top Hiring Sector
AI SOC Analyst	$125K – $165K	$170K – $215K	$235K – $300K	MSSP, banking, healthcare
GenAI Red Teamer	$165K – $210K	$220K – $290K	$320K – $440K	Big Tech, defence, consulting
AI Security Engineer	$170K – $220K	$235K – $310K	$340K – $470K	Big Tech, scale-ups, fintech
AI Governance Lead	$135K – $180K	$195K – $260K	$285K – $395K	Banking, pharma, regulated EU

Which pathway pays the most?

AI Security Engineer and GenAI Red Teamer top the table at senior+ bands because both require working knowledge of model internals and active offensive/defensive tooling. AI Governance Lead climbs fastest in regulated geographies (EU, Singapore, UAE). AI SOC Analyst has the largest open-headcount, so volume of opportunities is highest even if absolute pay is lowest of the four.

Regional adjusters (multiply tier-1 USA midpoint)

- **USA tier-1 (NY/SF/Seattle):** 1.00× baseline · **USA tier-2:** 0.80×
- **UAE / KSA:** 0.88× tax-free · **Singapore:** 0.75× · **UK · London:** 0.65×
- **Germany / Netherlands:** 0.68× · **Australia:** 0.74× · **Canada:** 0.70×
- **India · metro:** 0.26× · **India · tier-2:** 0.16×

Worked example: AI Security Engineer · Senior · USA tier-1 midpoint $270K. In Singapore: $270K × 0.75 = ~$203K. In Dubai (tax-free): $270K × 0.88 = ~$238K with effective take-home premium ~25% over Singapore due to zero income tax.

Pathway 1 · AI SOC Analyst — 90-Day Roadmap

1

AI SOC Analyst

Detection & triage of AI-powered attacks at L2/L3. Maps SIEM, EDR, and AI-generated phishing/malware patterns.

Mid-career total comp · $125K – $165K · USA tier-1 baseline

WEEKS 1–3 · FOUNDATIONS

Modules 1–2 · LLMs, threat surface, AI attack taxonomy

Land the vocabulary. How LLMs work end-to-end as an analyst needs to know. The MITRE ATLAS taxonomy of AI attacks. Map every entry to a detection your SOC would (or wouldn't) catch today.

WEEKS 4–6 · GEN-AI PHISHING & SOCIAL ENGINEERING

Module 3 · Hands-on detection & triage labs

Build detections for AI-generated phishing, deepfake voice, BEC variants. Practical SIEM queries (Splunk / Sentinel / Elastic). Triage runbooks for each pattern. **Capstone Artefact 1** drafted here.

WEEKS 7–9 · AI-AUGMENTED MALWARE & LATERAL MOVEMENT

Module 4 + Module 6 (SOC slice)

AI-generated malware indicators, polymorphic payloads, prompt-injection-based C2. Lateral movement when attackers use LLM tooling. Update your detection-as-code pipeline.

WEEKS 10–12 · CAPSTONE BUILD & DEFENCE

Module 9 · Submit three artefacts + defend

Detection ruleset, incident-playbook one-pager, post-incident report on a simulated AI-driven breach. Defend in front of an evaluator and earn CGAIC.

Pathway 1 · Capstone Picks & Target Employers

The 3 capstone artefacts

1 AI Threat Detection Ruleset

A SIEM-portable detection pack covering AI-generated phishing, deepfake voice, AI-crafted malware patterns, and prompt-injection-based C2. **Deliverable:** 12–15 detections with MITRE ATLAS mapping, test cases, and false-positive notes. **Used in interviews:** walk through one detection end-to-end.

2 Incident Response Playbook · AI-Driven Breach

A one-pager IR playbook for an AI-augmented breach scenario — covering triage, containment, eradication, recovery, and lessons learned. **Deliverable:** single visual page plus a 3-page runbook. **Used in interviews:** defend one decision point under questioning.

3 Post-Incident Report · Simulated AI Phishing Campaign

A full forensic write-up of a simulated AI-driven phishing wave targeting your sector. **Deliverable:** 6–8 page report with timeline, IoCs, root-cause, and three recommendations. **Used in interviews:** the recommendations section is what hiring managers ask about.

Target employers · AI SOC Analyst

- **Managed Security Services (MSSP)**
- **Banking · global & regional**
- **Healthcare & payor networks**
- **SaaS & cloud platforms**
- **Big-4 cyber consulting**
- **Federal & state government**
- **Telco & carrier networks**
- **Retail & payments processors**
- **Insurance & reinsurance**
- **Critical infrastructure (utilities)**

Hiring signals to look for in postings

- "AI-powered threat detection," "MITRE ATLAS," or "GenAI attack patterns" mentioned in the JD body.
- SIEM platforms named — Splunk, Sentinel, Chronicle, Elastic — with reference to GenAI use-cases.
- Preferred or required: vendor-neutral AI security certification (CGAIC qualifies).

Pathway 2 · GenAI Red Teamer — 90-Day Roadmap

2

GenAI Red Teamer

Lead end-to-end red-team engagements against LLMs and AI-augmented stacks. Prompt injection, jailbreak chains, data exfiltration, model extraction.

Mid-career total comp · $165K – $210K · USA tier-1 baseline

WEEKS 1–3 · ATTACK SURFACE

Modules 1–2 · LLM internals from the attacker's seat

Tokenisers, system prompts, RAG architectures, agent loops, tool calls. Where each layer breaks. Build a personal attack-surface map for a target archetype.

WEEKS 4–6 · PROMPT INJECTION & JAILBREAK CHAINS

Module 3 + Module 5 (offensive slice)

Direct + indirect prompt injection. Multi-turn jailbreaks, encoded payloads, persona attacks, tool-call hijack. Practical labs on open and commercial models. **Capstone Artefact 1** drafted.

WEEKS 7–9 · MODEL EXTRACTION & DATA EXFIL

Module 5 + Module 7

Membership inference, training-data extraction, embedding leak, RAG poisoning. Building reproducible PoCs. Reporting in the OWASP LLM Top-10 format.

WEEKS 10–12 · CAPSTONE BUILD & DEFENCE

Module 9 · Submit three artefacts + defend

Red-team report on a target system, exploit chain demo, and remediation memo. Defend in front of an evaluator and earn CGAIC.

Code-depth note: Pathway 2 is the most code-heavy of the four. Expect to be writing exploit scripts in Python and chaining them with public LLM APIs from Week 4 onward. Comfortable with at least one scripting language before you start.

Pathway 2 · Capstone Picks & Target Employers

The 3 capstone artefacts

1 Red-Team Engagement Report (full)

A complete engagement report against a target AI system — scope, methodology, findings, evidence, remediation. **Deliverable:** 15–25 page report in OWASP LLM Top-10 format. **Used in interviews:** walk through one finding from discovery to remediation.

2 Exploit Chain Demo · Reproducible PoC

A reproducible exploit chain that goes from initial access (via prompt injection) to data exfiltration. **Deliverable:** Git repo with scripts, screen recording, and a written narrative. **Used in interviews:** live demo, then a defence of the impact narrative.

3 Remediation Memo for the Engineering Team

A short, builder-friendly memo translating red-team findings into concrete fixes (input filters, output filters, scope-limiting, monitoring). **Deliverable:** 3-page memo + a 1-page fix-priority matrix. **Used in interviews:** proves you can collaborate with builders, not just break things.

Target employers · GenAI Red Teamer

- **Big Tech AI labs & product orgs**
- **Defence & national-security contractors**
- **Specialist red-team consultancies**
- **Big-4 offensive security practice**
- **Frontier AI labs · safety teams**
- **Banking · adversarial testing**
- **Bug-bounty platforms · staff roles**
- **Fintech & crypto exchanges**
- **Independent contractor / boutique**
- **Critical national infrastructure**

Hiring signals to look for in postings

- "AI red team," "LLM penetration testing," "prompt injection" in the JD body — not just buried in keywords.
- OWASP LLM Top-10 or MITRE ATLAS named explicitly.
- Job ladder includes "AI Security Researcher" or "Adversarial ML Engineer" — these are the senior versions of the role.

Pathway 3 · AI Security Engineer — 90-Day Roadmap

3

AI Security Engineer

Build and harden production AI systems — secure model deployment, guardrails, MLOps security, RAG-pipeline integrity.

Mid-career total comp · $170K – $220K · USA tier-1 baseline

WEEKS 1–3 · SECURE-BY-DESIGN FOUNDATIONS

Modules 1–2 + Module 6 (eng. slice)

LLM stack from the engineer's seat. Where each component fits in your existing AppSec and DevSecOps controls. Threat-model an AI feature end-to-end.

WEEKS 4–6 · GUARDRAILS & INPUT/OUTPUT FILTERING

Module 5 + Module 6

Build prompt filters, output validators, scope-limiting agents, RAG-context sanitisation. Bench-test against the OWASP LLM Top-10. **Capstone Artefact 1** drafted.

WEEKS 7–9 · MLOPS SECURITY & SUPPLY CHAIN

Module 6 + Module 7

Model registry hardening, signing & provenance, supply-chain attack patterns (poisoned weights, malicious LoRA), monitoring & rollback, secret scanning in prompts.

WEEKS 10–12 · CAPSTONE BUILD & DEFENCE

Module 9 · Submit three artefacts + defend

A secure-AI reference architecture, working guardrail kit (as code), and a hardened RAG pipeline. Defend in front of an evaluator and earn CGAIC.

Pathway 3 · Capstone Picks & Target Employers

The 3 capstone artefacts

1 Secure AI Reference Architecture

A reference architecture diagram and one-page narrative for a production-grade AI feature — covering identity, network segmentation, guardrails, observability, and incident response hooks. **Deliverable:** diagram + 3 pages of narrative. **Used in interviews:** system-design rounds.

2 Guardrail Kit · Working Code

An installable guardrail library that wraps an LLM call with input/output filters, scope limits, and audit logging. **Deliverable:** Git repo, README, tests against the OWASP LLM Top-10. **Used in interviews:** live walkthrough of one filter's bypass-resistance.

3 Hardened RAG Pipeline

A demo RAG pipeline with end-to-end controls — source vetting, content sanitisation, retrieval-result filtering, response validation, and tamper-evident logging. **Deliverable:** repo + architecture note. **Used in interviews:** the threat-model walk-through.

Target employers · AI Security Engineer

- **Big Tech (hyperscalers, AI labs)**
- **AI-native scale-ups · Series C+**
- **Fintech & banking platforms**
- **SaaS & B2B platforms shipping AI**
- **Healthcare AI vendors**
- **Cloud security & platform vendors**
- **Defence-adjacent commercial firms**
- **Crypto / Web3 infrastructure**
- **E-commerce at scale**
- **Independent security boutiques**

Hiring signals to look for in postings

- "AI security engineer," "ML security," or "GenAI platform security" in the title — not just "security engineer with AI exposure."
- Stack mentions: vector DBs (Pinecone, Weaviate), model registries (MLflow, SageMaker), guardrail frameworks.
- Ladders into AI Security Architect or Principal AI Security Engineer at the senior tier.

Pathway 4 · AI Governance Lead — 90-Day Roadmap

4

AI Governance Lead

Own AI risk, policy, and compliance for the enterprise. Board-facing, regulator-facing, framework-driven. Lowest code requirement of the four pathways.

Mid-career total comp · $135K – $180K · USA tier-1 baseline

WEEKS 1–3 · FRAMEWORKS & LANDSCAPE

Modules 1–2 + Module 8 (intro)

NIST AI RMF, ISO/IEC 42001, EU AI Act, NYC Local Law 144, OWASP LLM Top-10, MITRE ATLAS. Map each to your enterprise risk taxonomy.

WEEKS 4–6 · CONTROL DESIGN & ASSURANCE

Module 8 (deep) + Module 5 (slice)

Translate frameworks into controls. Build the AI control library, the model-card template, the model-risk register. Adverse-impact testing & bias audits. **Capstone Artefact 1** drafted.

WEEKS 7–9 · REGULATOR & BOARD READINESS

Module 8 + Module 9 (governance slice)

Vendor governance, third-party AI assessments, regulator-facing artefacts, board-pack design for AI risk reporting. The Q&A you'll face from an auditor.

WEEKS 10–12 · CAPSTONE BUILD & DEFENCE

Module 9 · Submit three artefacts + defend

AI policy & control library, an EU AI Act readiness assessment, and a board-pack one-pager. Defend in front of an evaluator and earn CGAIC.

Pathway 4 · Capstone Picks & Target Employers

The 3 capstone artefacts

1 AI Policy & Control Library

An enterprise-grade AI policy plus the underlying control library, mapped to NIST AI RMF and ISO/IEC 42001. **Deliverable:** 12–15 page policy + control matrix in a spreadsheet. **Used in interviews:** defend one control's design choice under regulator-style questioning.

2 EU AI Act Readiness Assessment

A scoped readiness assessment for an enterprise's AI portfolio against the EU AI Act — risk classification, gap analysis, remediation plan, timeline. **Deliverable:** 8–10 page assessment + 1-page exec summary. **Used in interviews:** the risk-classification reasoning.

3 Board-Pack One-Pager · AI Risk Reporting

A single page suitable for monthly AI risk reporting to a board risk committee. **Deliverable:** one visual page with metrics, trend, top-3 risks, and proposed actions. **Used in interviews:** walk the page top-to-bottom in 90 seconds.

Target employers · AI Governance Lead

- **Banking · global & regional**
- **Pharma & life sciences**
- **Insurance & reinsurance**
- **Big-4 risk & consulting**
- **EU-regulated enterprises**
- **Government & federal agencies**
- **Sovereign-linked (UAE, SG, KSA)**
- **Large F500 with internal audit**
- **AI-tech vendors (governance customer-facing)**
- **Law firms with tech practice**

Hiring signals to look for in postings

- "AI Governance," "AI Risk Officer," "Responsible AI Lead" in the title — these are the modern titles.
- NIST AI RMF, ISO/IEC 42001, or EU AI Act named in the JD body.
- Reporting line into Chief Risk Officer, Chief Compliance Officer, or General Counsel — not into CISO alone.

9-Module Syllabus · CGAIC (Verbatim)

All 9 modules of the Certified Generative AI in Cybersecurity program. Every pathway uses the same 9 modules; the depth and ordering shift per persona.

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, prompt-injection-based C2, 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, 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 your pathway, build the three artefacts, defend in front of an evaluator, earn the CGAIC credential.

Each module ships with hands-on labs, a case-study reading set, and a graded assessment. Total program time-investment is 6–8 hours per week across the 90-day window.

🚀 NEXT COHORT STARTING SOON

Join the next CGAIC cohort with this brief in hand

You've now seen the syllabus. The next cohort uses this exact roadmap and is open to certification candidates — applying now earns the launch window discount.

Join The Next Cohort →

Module-to-Pathway Mapping

Every pathway covers every module — but the order, depth, and time-allocation shift. Use this matrix to know what to focus on inside each module when you're on a specific pathway.

Module	SOC Analyst	Red Teamer	Sec Engineer	Governance
M01 · Foundations	Core	Core	Core	Core
M02 · Threat Landscape	Core	Core	Core	Core
M03 · GenAI Phishing	Heavy	Medium	Medium	Awareness
M04 · AI-Augmented Malware	Heavy	Heavy	Medium	Awareness
M05 · Prompt Injection & Exploitation	Medium	Heavy	Heavy	Awareness
M06 · Secure-by-Design	Awareness	Medium	Heavy	Medium
M07 · MLOps Security	Awareness	Medium	Heavy	Medium
M08 · Governance & Compliance	Awareness	Awareness	Medium	Heavy
M09 · Capstone	Core	Core	Core	Core

How to read the matrix

- **Core** — every pathway must complete this module fully. No exceptions.
- **Heavy** — primary depth area for that pathway. Spend the most hands-on hours here.
- **Medium** — solid working knowledge expected. You can be questioned on it in interviews.
- **Awareness** — understand the concepts and vocabulary. You won't be expected to build from scratch.

The capstone in Module 9 is where the pathway diverges most. Same evaluator format, but the three artefacts are pathway-specific — the ones laid out earlier in this brief.

Sample Exam — Part 1 of 2

Six representative questions across the four pathways. The real CGAIC exam is 60 multiple-choice + the capstone defence. Answers at the end of part 2.

Q1 · MODULE 03 · SOC PATHWAY

A user reports a phishing email. The text is grammatically perfect, the sender domain is a near-look-alike, and the link goes through three redirects to a phishing page. The strongest single SIEM signal you should pivot on is:

(a) The Unicode homoglyph ratio in the sender domain.

(b) The behavioural pattern of redirects matching known AI-generated phishing infrastructure.

(c) The email body length only.

(d) The user's prior reporting history.

Q2 · MODULE 05 · RED TEAM PATHWAY

During a red-team engagement, a customer-support chatbot ignores its system prompt when asked in base64-encoded form. The correct 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.

Q3 · MODULE 06 · SECURITY ENGINEER PATHWAY

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.

Sample Exam — Part 2 of 2

Q4 · MODULE 08 · GOVERNANCE PATHWAY

Under the EU AI Act, an HR system that uses an LLM to score job-candidate interview transcripts is 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.

Q5 · MODULE 07 · SECURITY ENGINEER PATHWAY

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 02 · ALL PATHWAYS

An adversary uses an LLM to generate variants of a known PowerShell loader, then weaponises one variant that AV doesn't recognise. In MITRE ATLAS, this maps most directly to:

(a) Initial Access via Drive-by Compromise.

(b) Defense Evasion · AI-Generated Polymorphic Variants.

(c) Discovery · Account Discovery.

(d) Persistence via Scheduled Task.

Answer key

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

If you scored 5–6 out of 6

You already have the working vocabulary for one or more pathways. The certification's value is mostly the capstone defence and the credential signal — both of which the labs and Module 9 deliver.

If you scored 3–4 out of 6

You have foundations but gaps. The 90-day program is well-paced for you — most scores in this range land at 90%+ on the real exam after the cohort.

FAQs · Honest Answers Before You Enrol

Can I switch pathways after I start?

Yes — pathway selection is a soft commitment until Week 4. After Week 4, switching costs you 2–3 weeks of rework on the capstone direction, so most candidates stick. The cohort lets you switch up to twice without re-enrolling.

Do I need to know how to code?

It depends on the pathway. **SOC and Governance:** no real coding required — light scripting helps but isn't blocking. **Red Teamer and Security Engineer:** assume comfort with Python and one cloud platform; you'll be writing exploit scripts or guardrail code from Week 4 onward.

How is CGAIC different from CISSP or CEH?

CISSP and CEH cover the classical security body of knowledge; neither focuses on AI-specific threats, OWASP LLM Top-10, MITRE ATLAS, or LLM-stack defence. CGAIC is purpose-built for the AI security work — designed to complement your existing certifications, not replace them.

Will my employer recognise this?

CGAIC is a vendor-neutral GSDC credential, recognised across global enterprise security and Big-4 consulting. The capstone defence is what hiring managers ask about most — more than the badge itself.

What's the time commitment per week?

Plan for 6–8 hours per week — about an hour a day on weekdays plus a Saturday cohort session. The cadence is designed to be sustainable around a full-time role.

Is the capstone defence remote or in-person?

Remote, via video. You walk the evaluator through your three artefacts and answer ~30 minutes of questions. Recordings are kept for credential audit; defenders own their evaluation criteria.

Pre-Enrolment Checklist · Printable

Tear this page out (or print it). Run through this before you enrol. Every box you can tick raises your odds of finishing the 90 days on time and landing the role you want after.

Pathway decision

- ✓ You've picked **one** of the four pathways (page 2). You'll re-confirm in Week 4.
- ✓ You've checked the **code-depth requirement** matches your current skills.
- ✓ You've read the **three capstone artefacts** for your pathway and they sound interesting (not torturous).
- ✓ You've checked the **target employer list** for your pathway and at least three companies are realistic targets.

Calendar & cadence

- ✓ You can put aside **6–8 hours per week** for 90 days without significant work disruption.
- ✓ You can put aside one **Saturday block (3–4 hrs)** for the live cohort session.
- ✓ You've blocked **two weeks at the end** for capstone build and defence.
- ✓ Your **line manager knows** you're doing this — or you've decided it stays off the radar.

Workspace & tooling

- ✓ You have a personal workspace with internet you control (not just locked-down employer device).
- ✓ You have access to at least one commercial LLM API (OpenAI / Anthropic / Google) — even a paid personal account.
- ✓ For Pathway 2 / 3: Python, Git, and a cloud account (AWS / Azure / GCP free tier is fine).
- ✓ For Pathway 4: a working knowledge of one risk-management framework (NIST CSF, ISO 27001, COSO) helps.

Career positioning

- ✓ Your resume mentions **"AI security"** or a pathway-specific term in the top three lines.
- ✓ Your LinkedIn headline matches the **pathway title** you've picked.
- ✓ You have at least **one analytics or technical artefact** you can show today, even before CGAIC.
- ✓ You've identified the **3–5 employers** you'll target in the 90 days *after* the credential.

Glossary & About This Brief

Glossary

- **CGAIC:** Certified Generative AI in Cybersecurity — the GSDC vendor-neutral AI security credential covered in this brief.
- **OWASP LLM Top-10:** The current OWASP top-10 application-security risks specific to LLM applications.
- **MITRE ATLAS:** The Adversarial Threat Landscape for AI Systems — MITRE's tactics-and-techniques framework for AI attacks.
- **Prompt injection:** An attack where the model is induced to ignore its system instructions via crafted input — direct (in the user message) or indirect (in retrieved or tool-returned content).
- **RAG:** Retrieval-Augmented Generation — an architecture where an LLM is grounded with retrieved documents at query time.
- **Red team (AI):** A structured offensive engagement against an AI system to find security and safety failure modes before adversaries do.
- **EU AI Act:** The European Union's risk-tiered regulation of AI systems, with high-risk categories carrying conformity-assessment and registration obligations.
- **NIST AI RMF:** The U.S. NIST AI Risk Management Framework — a voluntary, broadly-adopted approach to governing AI risk across the lifecycle.

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 capstone defence and the assessment, 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.

The Full Pathway Roadmap · On One Page

The 4 pathways (pages 2, 4, 6, 8, 10)

AI SOC Analyst · GenAI Red Teamer · AI Security Engineer · AI Governance Lead. Same CGAIC certification, four distinct career destinations. Pick one before you start — switching costs time.

Per-pathway salary bands (page 3)

Mid-career USA tier-1 midpoint by pathway: SOC $125–165K · Red Team $165–210K · Sec Engineer $170–220K · Governance $135–180K. Senior bands clear $300K for the top three pathways.

Per-pathway capstone artefacts (pages 5, 7, 9, 11)

Three deliverables per pathway. SOC: detection ruleset + IR playbook + post-incident report. Red Team: engagement report + exploit chain + remediation memo. Sec Engineer: reference architecture + guardrail kit + hardened RAG. Governance: AI policy + EU AI Act assessment + board-pack one-pager.

Per-pathway target employers (pages 5, 7, 9, 11)

~10 employer categories per pathway. Total ~40 distinct employer categories across the four pathways. Pick 3–5 to target in the 90 days after the credential.

The 9-module syllabus (page 12)

Foundations · Threat Landscape · GenAI Phishing · AI Malware · Prompt Injection · Secure-by-Design · MLOps Security · Governance · Capstone. Heavy vs medium vs awareness depth shifts by pathway (matrix on page 13).

How CGAIC fits in your career

Vendor-neutral credential. Recognised across enterprise security, Big-4, and global hiring. The capstone defence is the recurring artefact you'll point to in interviews — the credential opens the door, the artefacts close the offer.