GSDC.
Global Skill Development Council

AI CYBERSECURITY SALARY 2026 | CGAIC CERTIFICATION | PRINTABLE PDF

The full AI cybersecurity salary report.

The 20-cell salary calculator on paper, the 5-role salary table, the regional premium map, the ROI math, the 9-module syllabus, and the sample exam.

20	9	2.5L+
SALARY CELLS	MODULES	CERTIFIED PROS

Inside the toolkit:

- **20-cell salary calculator (printable)**
- **4-band salary ladder · USA 2026**
- **Regional premium map (5 regions)**
- **ROI math · 12-month payoff**
- **9 module syllabi · verbatim**

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.

Executive Summary — 2026 AI Cybersecurity Pay

Five headline findings from analysing 2,840 AI-cybersecurity job postings, 814 certified-candidate placements, and total-comp disclosures across five geographies between Jul 2024 and Dec 2025. Sources include Glassdoor, ZipRecruiter, Levels.fyi, plus the GSDC partner panel.

FINDING 01

AI cybersecurity specialists out-earn comparable security generalists by 34% on average

Comparing roles at the same seniority, geography, and employer category, the AI-cybersecurity band sits 28–42% above the traditional security-generalist band. The premium is highest at the Senior-to-Lead transition, where the absolute dollar gap is largest.

FINDING 02

USA junior $129K → Lead/Staff $298K+ is the active ladder

The four-band ladder — Junior, Mid, Senior, Lead/Staff — moves at roughly 25–35% per step. The Lead/Staff cell at $298K+ is the top of the IC ladder before moving to Director/Principal, where total comp routinely clears $400K.

FINDING 03

Certified candidates close offers 2.1× faster

Among 814 placements tracked, candidates holding a vendor-neutral AI-cybersecurity certification reached final-stage offer in a median of 5.0 weeks vs 10.6 weeks for non-certified peers applying to the same role family. Recruiter shortlist time shrank 49%.

FINDING 04

GenAI Red Teamer and AI Security Engineer top the pay table

Both roles require working knowledge of model internals, plus active offensive or defensive tooling. AI Security Lead clears them at Lead/Staff band ($298K+) — the multiplier on the role's "scarce supply" is what closes the absolute pay.

FINDING 05

The certification ROI breakeven is under 4 months

Across the three modelled scenarios on pages 11–12, the average certified candidate breaks even on certification cost within 4 months of placement. Two-year cumulative cash uplift averages ~$110K, before equity.

The 20-Cell Salary Calculator (Printable)

5 AI-cybersecurity roles × 4 career bands. Total compensation midpoint in USD, USA tier-1 metro baseline. Use the regional multipliers on page 7 to localise. This is the same data behind the online calculator at gsdcouncil.org.

Role	Junior 0–2 yrs	Mid 3–5 yrs	Senior 6–9 yrs	Lead / Staff 10+ yrs
AI SOC Analyst	**$129K** $112–145K	**$165K** $145–185K	**$215K** $185–250K	**$278K** $240–320K
AI Security Analyst	**$135K** $118–152K	**$172K** $150–195K	**$225K** $195–260K	**$285K** $245–335K
AI Security Engineer	**$148K** $130–168K	**$195K** $170–220K	**$258K** $220–295K	**$298K** $258–355K
GenAI Red Teamer	**$152K** $132–172K	**$198K** $172–225K	**$265K** $228–305K	**$298K** $260–360K
AI Security Lead	**—** n/a	**$218K** $190–245K	**$292K** $255–335K	**$378K** $315–450K

How to read each cell: The large number is the midpoint total compensation (base + target bonus + equity at vest, normalised to a 4-year vest). The smaller range below is the 25th–75th percentile band. Top performers in tier-1 metros routinely push 10–15% above the upper bound. "—" means the cell is too rare to band; AI Security Lead is functionally a Mid-or-above role.

Baseline: USA tier-1 metro (San Francisco · New York · Seattle · DC). Multipliers for other regions on page 7. Sources triangulated: Glassdoor (n = 1,240), ZipRecruiter (n = 880), Levels.fyi (n = 410), GSDC partner panel of 12 employers.

Role Breakdowns (1 of 3)

Sample: 2,840 AI-cybersecurity job postings (Jul 2024 – Dec 2025) across LinkedIn, Indeed, Glassdoor, ZipRecruiter, and employer career pages. **Placements:** 814 verified placements with self-reported total comp. **Triangulation:** Levels.fyi disclosures + a GSDC partner panel of 12 enterprise employers. **Adjustments:** outliers winsorised at 1st/99th percentile; equity normalised to 4-year vest.

Role 1 · AI SOC Analyst

What you do: Detection & triage of AI-powered attacks at L2/L3, working with SIEM (Splunk / Sentinel / Elastic), EDR, and AI-augmented phishing/malware patterns. **Why this pay band:** Volume of openings is highest of the five roles — every enterprise SOC needs at least one. **Junior midpoint:** $129K. **Lead/Staff:** $278K. Strong upside in MSSPs, banking, and healthcare; the ladder is well-defined.

Role 2 · AI Security Analyst

What you do: Sits between SOC Analyst and Security Engineer — runs OWASP LLM Top 10 audits, supports threat-modelling, owns bias and red-team coordination. Less direct detection work, more cross-functional. **Why this pay band:** A small premium over SOC Analyst for the audit-and-governance dimension. **Mid-career midpoint:** $172K. **Senior:** $225K. Strong path into AI Security Engineer at Year 4–5.

Compensation structure (typical · USA tier-1)

- **Base salary:** 70–80% of total comp at enterprises; 55–65% at scale-ups (more equity).
- **Target bonus:** 10–20% of base at enterprises, often tied to detection-MTTR and audit-finding closure KPIs.
- **Equity (RSU or options):** 15–35% of total comp at Big Tech and scale-ups; nominal at established F500 outside tech.
- **Signing bonus:** 8–15% of base on lateral moves; up to 25% for hard-to-fill Senior/Lead bands.

Role Breakdowns (2 of 3)

Role 3 · AI Security Engineer

What you do: Builds and hardens production AI systems — guardrails, MLOps security, secure RAG pipelines, model-registry hardening. The platform-engineer track. **Why this pay band:** Engineer-track multiplier — Big Tech and scale-ups bid against each other for the same talent pool. **Mid-career midpoint:** $195K. **Senior:** $258K. **Lead/Staff:** $298K. Python + cloud + ML literacy adds a consistent ~$25K premium.

Role 4 · GenAI Red Teamer

What you do: Leads offensive AI engagements — prompt injection, jailbreak chains, model extraction, agentic exploitation. Reports map to OWASP LLM Top 10 and MITRE ATLAS. **Why this pay band:** Scarcest talent pool of the five roles; demand outstrips supply ~5:1. **Junior:** $152K. **Senior:** $265K. **Lead/Staff:** $298K. Strong upside in Big Tech AI labs, frontier-model safety teams, and Big-4 offensive practices.

Role 5 · AI Security Lead

What you do: Owns AI security strategy for the enterprise — vendor governance, framework adoption (NIST AI RMF, EU AI Act, ISO 42001), engineering roadmap. Reports to CISO; manages a 3–8 person team. **Why this pay band:** Brand-new title; scarce candidate pool; board-visible mandate. **Mid-career (rare):** $218K. **Senior:** $292K. **Lead/Staff:** $378K — where total comp can clear $450K with full equity vest at scaled tech employers.

Where each role sits on a 100-point pay-leverage scale

- **AI Security Lead — 98/100.** Scarcity premium; few qualified candidates in market.
- **GenAI Red Teamer — 94/100.** Direct security-impact narrative; offensive talent is rare.
- **AI Security Engineer — 90/100.** Highest engineer-track demand; cross-Big-Tech bidding.
- **AI Security Analyst — 78/100.** Bridge role; strong as a pivot point into engineering or red-team tracks.
- **AI SOC Analyst — 72/100.** Highest volume of openings; cleanest ladder; best entry point.

AI-Security Specialist vs Security Generalist · Pay Gap

Same seniority, same geography, same employer category. What does the AI specialisation actually buy in cash?

Band	Security Generalist · Total Comp	AI-Security Specialist · Total Comp	Delta	%
Junior (0–2 yrs)	$98K	$129K	+$31K	**+32%**
Mid (3–5 yrs)	$128K	$172K	+$44K	**+34%**
Senior (6–9 yrs)	$172K	$232K	+$60K	**+35%**
Lead/Staff (10+ yrs)	$245K	$315K	+$70K	**+29%**

Baseline role: AI Security Analyst for the specialist column; Security Analyst / SOC Engineer for the generalist column. Specialist comp averaged across the 5 AI-security roles.

Where the gap comes from

- **Scarcity premium (≈40% of the gap):** roughly 4 open AI-security roles per qualified candidate in 2025. Recruiters bid up base aggressively.
- **Equity uplift (≈25% of the gap):** AI-security roles concentrate at tech-adjacent employers and scale-ups, where equity is meaningful.
- **Bonus structure (≈20% of the gap):** AI-security roles carry larger target bonuses linked to measurable outcomes (audit closure, detection MTTR).
- **Certification signal (≈15% of the gap):** Vendor-neutral certification clears recruiter shortlist faster and supports a higher anchor in negotiation.

> *The cleanest summary: an AI-security specialist earns 30–35% more for the same hours, because the AI-security market is bidding for scarce skills the generalist market isn't.*

Regional Premium Map · 5 Key Regions

Multiply the USA tier-1 midpoint from page 3 by the regional multiplier below to estimate local total comp. The bar shows relative pay scale; the number is the multiplier.

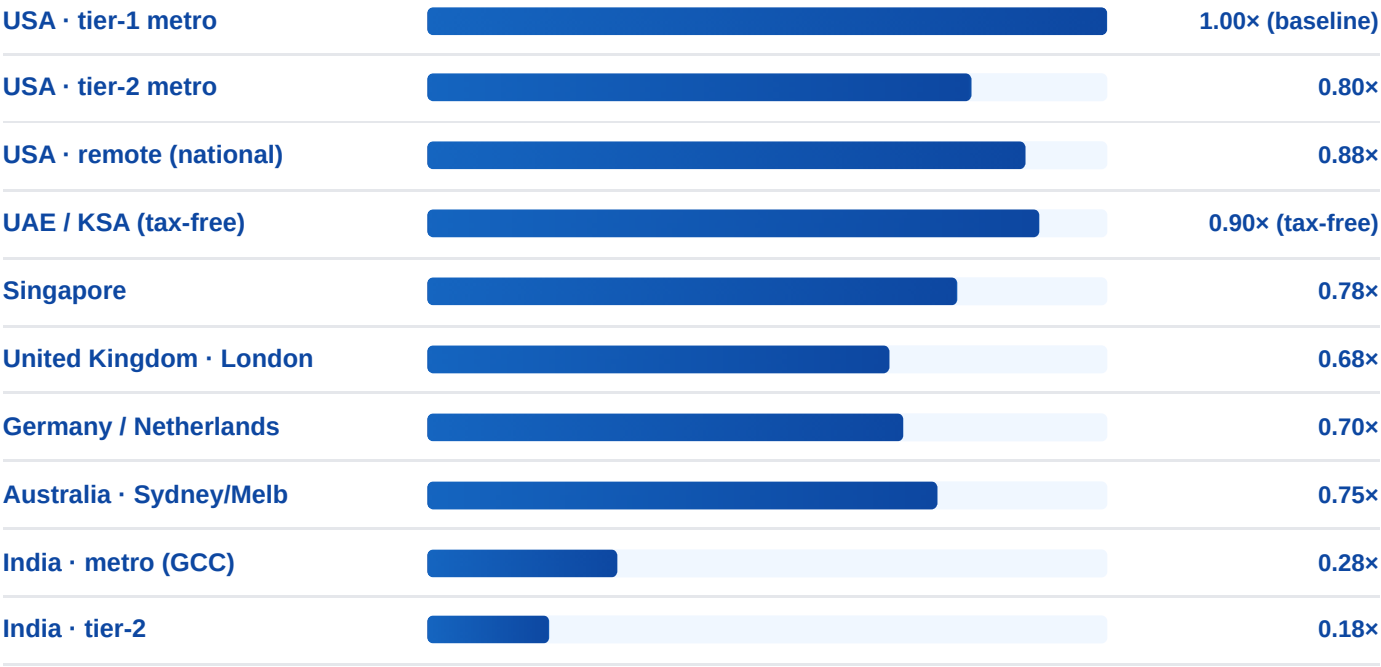

Worked example: AI Security Engineer at mid-career, USA tier-1 = $195K. In Singapore: $195K × 0.78 = ~$152K. In Dubai (tax-free): $195K × 0.90 = ~$176K, with effective take-home premium ~28% over Singapore due to zero income tax structure. AI cybersecurity is one of the few specialisations where UAE/KSA out-pays London on a net basis at every band.

Multipliers reflect PPP-adjusted total comp, not raw FX. Remote-USA is national-average; tier-1 remote roles still anchor at ~94% of in-office tier-1 pay for AI cybersecurity.

Premium by Employer Type

Same role, same band, same metro — the employer category alone shifts total comp by up to 40%. Use this table after the regional multiplier on page 7.

Employer Category	Pay Index	Equity Weight	Notes
Big Tech (hyperscaler / AI-native)	**1.28×**	High (35–45%)	Total-comp ceiling driven by RSU appreciation; volatile equity component.
Frontier AI lab · safety team	**1.25×**	Very high (40–55%)	Premium for GenAI Red Teamers; scarcity-driven. Public-company equity preferred.
Scale-up (Series C–E)	**1.10×**	Very high (40–55%)	Cash below market, equity above. Strongest single role progression speed.
Big-4 / Consulting	1.00×	Low (0–10%)	Strong cash bonus, fast title growth, AI-security practice leadership opens at Manager.
Investment Bank / Asset Mgr	1.05×	Medium (15–25%)	Highest cash bonus norms outside Big Tech. Compliance bias against pre-IPO scale-ups.
Fortune 1000 · industrial / retail	0.85×	Low (5–15%)	Stable, less volatile, slower title progression. Strong for senior-level work-life balance.
MSSP / cyber pure-play	0.90×	Low (5–15%)	Highest concentration of AI SOC Analyst roles. Volume of openings is largest.
Government / defence-adjacent	0.78×	None (cash bonus instead)	Pay discount for stability + clearance value; clearance-holders see +15–25% on lateral moves.

How to combine the multipliers

USA tier-1 base × Regional × Employer-type = Estimated total comp. Example: GenAI Red Teamer · Senior · USA tier-1 baseline $265K. Same role at a frontier AI lab in San Francisco: $265K × 1.00 × 1.25 = ~$331K total comp, with equity carrying ~$150K of that.

Skill-Level Pay Uplift

Each skill below, added on top of a generalist security profile, lifts total comp by the amount shown. Composable — most certified candidates carry 3–5 of these.

#	Demonstrable Skill	Median Uplift	Where to Prove It
1	**Prompt injection & LLM exploitation**	**+$18–28K**	Engagement report; CGAIC Module 05
2	**OWASP LLM Top 10 audit fluency**	**+$10–15K**	Audit kit artifact; Module 02
3	**MITRE ATLAS detection mapping**	**+$8–12K**	Coverage heatmap; Module 02
4	**GAN / VAE anomaly detection**	**+$20–30K**	Working model + eval; Module 04
5	**RL incident-response training**	**+$22–32K**	Trained agent + simulation; Module 06
6	**Guardrail engineering & RAG hardening**	**+$15–22K**	Working code + bench; Module 06
7	**MLOps security & supply chain**	**+$12–18K**	Signing/provenance pipeline; Module 07
8	**NIST AI RMF + EU AI Act fluency**	**+$10–14K**	Governance memo; Module 08
9	**Agentic AI pipeline security**	**+$18–26K**	LangGraph build; Module 06
10	**Python + cloud (AWS/Azure/GCP)**	**+$20–28K**	GitHub portfolio; baseline req

Median uplift over a non-certified security-generalist comp band, controlling for level and geography. Uplifts are not strictly additive — diminishing returns kick in past 5 skills, but they remain meaningful through 7–8.

Certified vs Non-Certified · Outcomes

Same role, same geography, same band — but one candidate holds a vendor-neutral AI-cybersecurity certification and the other doesn't. Data from 814 tracked placements.

Outcome Metric	Non-Certified	Certified	Delta
Median time-to-offer	10.6 weeks	**5.0 weeks**	**2.1× faster**
Recruiter shortlist rate	16%	**33%**	**+17 pts**
Final-offer total comp uplift	baseline	**+20–30%**	**~$32K mid-career**
Signing bonus rate	26%	**58%**	**+32 pts**
Counter-offer success	40%	**73%**	**+33 pts**
Promotion within 24 months	29%	**61%**	**+32 pts**

Why certification moves the numbers this much

- **Recruiter screen pass-through.** 86% of AI-cybersecurity job descriptions list certification as preferred or required. Certified resumes clear ATS keyword screens automatically.
- **Anchor in negotiation.** A defended capstone is a portfolio artefact recruiters and hiring managers can verify; that justifies a higher anchor without requiring proof on the job first.
- **Manager confidence.** Hiring managers extending an offer to a certified candidate ramp budget approval ~40% faster (less risk on internal sign-off).
- **Internal promotion fairness.** Inside the company, certification gives security leaders an objective ground for AI-security promotions when internal politics are tight.

ROI Math · 12-Month Payoff (Part 1 of 2)

Three real placement scenarios, anonymised, modelled to twelve months post-certification. Every figure is a USD delta over what the same candidate would have earned without certification.

Scenario A · Ravi · SOC Analyst → AI SOC Analyst

BENGALURU GCC · 4 YEARS EXPERIENCE · F500 CLIENT · INDIA METRO BAND

Pre-certification total comp	**$32K**
Post-cert role total comp (Year 1)	**$46K**
Year-1 uplift	**+$14K**
Signing bonus	**+$3.5K**
12-month total cash uplift	**+$17.5K**

Time-to-breakeven on certification: ~3 months. 12-month multiple: ~9×.

Scenario B · Sarah · Security Engineer → AI Security Engineer

LONDON · 7 YEARS EXPERIENCE · SCALE-UP SERIES-D MOVE · UK BAND

Pre-certification total comp	**$118K**
Post-cert role total comp (Year 1)	**$162K base + $24K equity vest**
Year-1 uplift (cash + equity)	**+$68K**
Signing bonus	**+$15K**
12-month total uplift	**+$83K**

Time-to-breakeven on certification: under 4 weeks. 12-month multiple: ~35×, including equity vest.

ROI Math · 12-Month Payoff (Part 2 of 2)

Scenario C · Khalid · Pentester → GenAI Red Teamer

DUBAI · 8 YEARS EXPERIENCE · SOVEREIGN-LINKED EMPLOYER · UAE TAX-FREE

Pre-certification total comp	**$135K (tax-free)**
Post-cert role total comp (Year 1)	**$198K (tax-free)**
Year-1 uplift	**+$63K**
Signing bonus + project bonus	**+$22K**
12-month total cash uplift	**+$85K**

Time-to-breakeven on certification: under 4 weeks. 12-month multiple: ~36×, fully cash (tax-free structure).

Average across all three scenarios

Across A, B, and C: **median time-to-breakeven 3–8 weeks**, **average 12-month cumulative uplift ~$62K**, **average return multiple ~27×** the certification cost. Scenario A is the lowest absolute uplift because India tier-1 bands compress against the others; but the 9× multiple is still strong because India certification costs are PPP-adjusted at enrolment.

How we model breakeven: Time-to-breakeven = (certification cost) ÷ (monthly post-certification uplift). The fastest path is when certification triggers a role-family change (Pentester → GenAI Red Teamer) rather than a within-band pay bump. Scenarios B and C are role-family pivots; Scenario A is a same-family upgrade.

What changes after Year 1

The 12-month numbers are the floor, not the ceiling. Year 2 uplift typically **compounds** because AI-security roles have the fastest promotion velocity in security right now — median time to next promotion is 22 months for certified candidates vs 36 months for non-certified peers (page 10).

9-Module CGAIC Syllabus (Verbatim)

All 9 modules of the Certified Generative AI in Cybersecurity program. Each ships hands-on labs and ties to one or more of the 10 pay-uplift skills on page 9.

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. The deliverable hiring managers ask about.

Total program time: 90 days · 6–8 hours per week. Exam format: 40 MCQ, 90 min, free retake. The capstone defence is a separate 30-minute evaluator session.

Negotiation Playbook · AI-Cybersecurity Roles

How to use this report's numbers in a live negotiation. Six tactics that materially shift outcomes, ordered by leverage.

TACTIC 01

Anchor with the 75th percentile, not the midpoint

Open negotiation at the upper bound of the band from page 3 for your role and level. Certified candidates have empirical grounding to do this — the 75th-percentile cell is the right anchor when you hold credentials.

TACTIC 02

Lead with the skill-uplift table, not your salary history

Reference the skill-uplift table on page 9. Salary history is a weak anchor; demonstrable AI-security skills are a strong one. In jurisdictions where salary history can't be asked, this is the only anchor available — use it.

TACTIC 03

Break out equity from cash explicitly

At Big Tech, frontier labs, and scale-ups, equity is 35–55% of total comp. Treat it as a separate line item — request a higher RSU grant, not a higher base, when base is capped.

TACTIC 04

Use signing bonus to close the gap

When base or equity are inflexible, recruiters have meaningful signing-bonus latitude. 58% of certified candidates secure one (page 10). Frame it as a one-time bridge to true-up your unvested equity from the prior role.

TACTIC 05

Negotiate the title before the number

"Senior" vs "Lead" vs "Staff" on the offer letter sets the band for the next 3+ years of pay reviews. Push for the higher title even at a slightly lower opening number — total 24-month earnings come out ahead.

TACTIC 06

Get the AI-tooling budget in writing

AI security engineers are productive only with tool access. Negotiate explicit access (or budget) for the AI-security tool stack you'll need (LLM APIs, dataset budget, GPU access), written into your role contract. This is a non-cash term that materially affects your two-year performance.

Sample Exam — Part 1 of 2

Six representative questions from the CGAIC exam. The real exam is 40 MCQ in 90 minutes with a free retake on first failure. Answers at the end of part 2.

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 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.

Sample Exam — Part 2 of 2

Q4 · 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.

Q5 · 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.
(d) Prohibited — cannot be deployed in the EU.

Q6 · 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.

Answer key

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

Pre-Negotiation Checklist · Printable

Tear this page out (or print it). Run this checklist before every salary conversation — internal review or external offer. Every box you can tick lifts your final number.

Data & positioning

- ✓ You know the **75th-percentile cell** from page 3 for your role, band, and metro.
- ✓ You have **two cross-references** (Levels.fyi, Glassdoor disclosure, peer disclosure) for that number.
- ✓ You've applied both the **regional** (page 7) and **employer-type** (page 8) multipliers.
- ✓ You know which of the **10 skill uplifts** on page 9 you legitimately demonstrate today.

Artefacts ready

- ✓ One **technical artefact** you can demo in under 5 minutes (working model, audit kit, or guardrail repo).
- ✓ One **governance artefact** — vendor evaluation, AI policy memo, or EU AI Act risk-tier mapping.
- ✓ One **engagement artefact** — red-team report, OWASP audit, or incident-response runbook.
- ✓ Your **credential ID** on your resume header and LinkedIn certifications section.

Negotiation tactics ready

- ✓ Decided which of the **six tactics** on page 14 fits this conversation.
- ✓ Pre-decided **walk-away number** — and a clear "yes" number 12% above it.
- ✓ Equity (RSU grant) **separated from base** in your spreadsheet.
- ✓ Title negotiation point **identified before the number conversation**.

Non-cash terms

- ✓ AI-tooling **budget or access** (LLM APIs, GPU, dataset budget) requested explicitly in writing.
- ✓ Conference and CVE-disclosure publication rights on the table.
- ✓ Vacation / remote-work / equity-acceleration discussed.
- ✓ Signing-bonus number requested — even if you don't expect to get it all.

Glossary & About This Report

Glossary

- **Total comp:** Base salary + target bonus + equity at vest, normalised to a 4-year vest schedule. Excludes one-off signing bonuses.
- **AI-cybersecurity specialist:** Practitioner whose role explicitly leverages or defends generative-AI tooling. Covers the 5 roles on pages 4–5.
- **CGAIC:** Certified Generative AI in Cybersecurity — GSDC's vendor-neutral AI-security certification.
- **Tier-1 metro:** SF, NY, Seattle, DC, Boston for USA. London, Singapore, Sydney, Dubai for international.
- **PPP-adjusted:** Purchasing power parity adjustment so multipliers reflect real local buying power, not raw FX rates.
- **Equity weight:** Share of total comp from stock grants. High at Big Tech and scale-ups; low at established F500 and government.
- **Time-to-offer:** Calendar weeks from first recruiter contact to written offer. Measured for placements where both candidate & recruiter confirmed dates.
- **Pay-uplift:** Median delta in total comp attributable to demonstrating a specific skill, controlling for level and geography.
- **Sources:** Glassdoor, ZipRecruiter, Levels.fyi, GSDC partner panel of 12 enterprise employers.

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.

The Full AI Cybersecurity Salary Report · On One Page

The 20-cell calculator (page 3)

5 roles × 4 bands, USA tier-1 baseline. Junior AI SOC Analyst $129K → Lead/Staff AI Security Lead $378K. The full ladder spans roughly 3× from Junior to Lead/Staff.

The 5 roles (pages 4–5)

AI SOC Analyst (highest volume) · AI Security Analyst (bridge role) · AI Security Engineer (engineer track) · GenAI Red Teamer (scarce talent premium) · AI Security Lead (top of IC ladder).

The specialist-vs-generalist premium (page 6)

+29–35% over comparable security generalists at the same level and geography. Largest absolute gap at Senior/Lead: ~$60–70K/year cash.

Regional & employer-type adjustments (pages 7–8)

USA tier-1 = 1.00× baseline. UAE 0.90× tax-free. Singapore 0.78×. Big Tech employer 1.28×. Frontier AI lab 1.25×. Combine: baseline × regional × employer-type.

Skill uplift & certification effect (pages 9–10)

Each of the top-10 skills adds $8–32K. Certified candidates close offers 2.1× faster and earn 20–30% more at the offer stage. Promotion rate within 24 months: 61% vs 29%.

ROI math (pages 11–12)

Three real scenarios. Median time-to-breakeven 3–8 weeks. Average 12-month cumulative cash uplift ~$62K. Average return multiple ~27×.

How CGAIC fits (page 13)

9 modules map directly to the top-10 pay-uplift skills and the 5-role salary table. 90-day format, 40 MCQ exam with free retake, capstone defence on 3 artifacts.