DATAMARCOS

LIVE ONLINE · MENTOR-LED PROGRAM

Cybersecurity with AI

8 weeks · 6–8 hours a week · Part-time, live online

An eight-week, part-time live program for security and engineering professionals who want to use AI carefully in security operations, and to secure LLM and agent applications against prompt injection, data leakage and related risks. Designed with US learners in mind, including weekend-morning live batches; join from anywhere.

Fees and EMI options shared on a counsellor call.

Program guide · Curriculum, projects and format. No prices are listed in this brochure.

Noida One, 303-304, Rasoolpur Nawada, Industrial Area, Sector 62, Noida, Uttar Pradesh 201309, India

www.datamarcos.com · hello@datamarcos.com · +91 88519 30644

About this program

Who it is for

- Security analysts, SOC and detection engineers who want a structured way to use AI in investigations
- Software and platform engineers who are shipping chat, RAG or agent features and need to secure them
- IT and security professionals in the US and other regions who can join a weekend-morning or other live batch
- People who already work around logs, alerts or applications, not complete beginners in cybersecurity

Who it is not for

- People looking for an official CompTIA, SANS, (ISC)² or government certificate. This program issues a Datamarcos certificate of completion
- Complete beginners who need a first course in cybersecurity fundamentals
- People seeking placement or salary guarantees — we do not offer those

Starting point

- Some exposure to IT, networks, application logs or security operations is helpful
- No prior machine-learning or LLM-engineering experience is required
- You should be willing to inspect logs, prompts or application behaviour and write down what you actually checked

What you will be able to do

These are skills and work products from the program — not placement, salary or promotion promises.

- 1. Use an AI assistant on a detection or investigation task and verify, with notes, what you accepted, edited or rejected
- 2. Write a short investigation or triage brief that a colleague can follow, including residual uncertainty
- 3. Describe, in practical terms, the OWASP Top 10 for LLM applications and how it applies to a chat, RAG or agent feature
- 4. Spot common prompt-injection and data-leakage patterns and recommend a human review or control step
- 5. Plan and document a short red-team pass on an LLM feature, including findings and residual risk
- 6. Leave with a capstone pack: a security-ops brief and/or an LLM threat model, plus recommended controls — not a vendor certification

Projects

Project 1: Verified security-ops brief

An investigation or detection note on a realistic sample, with an explicit list of what an AI assistant suggested and what you confirmed against the source.

Project 2: LLM application threat model

A worksheet that maps one chat, RAG or agent feature to the OWASP Top 10 for LLM applications and names realistic controls.

Project 3: Capstone: red-team note and recommended controls

A short, documented test, residual risk and a runbook. This is coursework, not an official penetration-testing or vendor certification.

DATAMARCOS

Curriculum · 8 weeks

6–8 hours a week, live and mentor-led. Open each week in the live sessions for labs and feedback.

Week 1: Where AI helps in security work — and where it fails

- Assisting versus automating: drafting, summarising and clustering, not unsighted close-out
- Typical failure modes: invented indicators, missed context, over-confident wording
- What must stay a human decision in a security workflow
- Choosing a practice task from detection, triage or investigation

Lab: Compare an assistant's summary of a sample alert with the source evidence and list errors.

Week 2: Threat detection and log analysis with AI assistance

- Writing a precise question over logs or alerts so the answer can be checked
- Using an assistant to suggest queries, clusters or next questions
- Verifying indicators, timestamps and asset names against the source
- A short method note: what the model suggested and what you confirmed

Lab: Run an assisted analysis on a sample log or alert set and tick off every claim.

Week 3: Triage, investigation and human review

- A triage workflow that includes a review step before any action is taken
- Writing an investigation brief another analyst can pick up
- What not to paste into a public assistant: secrets, personal data, unreleased findings
- When a finding should be escalated rather than auto-closed

Lab: Draft a triage brief from a sample case and mark the human review points.

Week 4: Phase checkpoint: a verified security-ops brief

- Assembling the analysis, the checks you ran and the residual uncertainty
- Peer and mentor review
- What the second half of the program will add: securing the AI systems themselves

Lab: Present your security-ops brief and the evidence behind it.

Week 5: OWASP Top 10 for LLM applications, in practice

- The attack surface of a chat, retrieval-augmented or tool-using agent application
- The OWASP Top 10 for LLM applications as a working checklist, not as a certificate
- How this differs from a traditional web or network review
- Picking one feature to threat-model for the rest of the program

Lab: Walk a sample LLM feature and mark relevant OWASP LLM items.
You ship: A first-pass threat-model worksheet.

Week 6: Prompt injection, data leakage and output handling

- Direct and indirect prompt injection, in practical terms
- Sensitive data in prompts, context windows and logs
- Insecure output handling: when a model reply is treated as a trusted instruction
- Basic mitigations: input and output review, least privilege for tools, and fallbacks

Lab: Attempt a small, authorised prompt injection test on a practice application and record what happened.
You ship: A findings note with suggested controls.

Curriculum (continued)

Week 7: Red-teaming an LLM or agent feature

- Scoping a short red-team pass: goals, rules of engagement and what is out of scope
- Recording findings so an engineer can reproduce them
- Residual risk: what you did not test
- This is structured testing, not a professional penetration-testing or official red-team certification

Lab: Run a short, documented test plan against your chosen feature.
You ship: A red-team note with findings and residual risk.

Week 8: Capstone: brief, threat model and recommended controls

- Assembling a capstone pack from the security-ops brief and/or the LLM threat model and red-team note
- Recommended controls and a short runbook a team could actually follow
- Live walk-throughs and mentor feedback

Lab: Present the capstone pack and the checks you ran.

Format, fees and refunds

Part-time and live online, in a small mentor-led cohort. Weekend-morning batches are offered for learners in the US. Evening batches (India time) are offered for learners in the Gulf and Africa. Europe-oriented groups typically meet after work and on Saturdays. Recordings are shared if you miss a live session. Exact days, times and the next cohort date are confirmed on a counsellor call.

Designed with US learners in mind, including weekend-morning live batches. Join live from anywhere.

A Datamarcos certificate of completion is awarded when you finish the program work. Cohort start dates are shared on a counsellor call — we do not publish a date until then.

Fees and EMI options shared on a counsellor call.

7-day refund policy: if the program is not right for you, ask within 7 days of enrollment.

Mentor

Live sessions are mentor-led by Deepesh Mishra, founder of Datamarcos, a BITS Pilani alumnus, and a solution architect and corporate trainer in data and AI who has trained teams at 200+ organizations.

Questions, answered

Is this an official CompTIA, SANS, (ISC)² or government certificate?

No. Completing the program leads to a Datamarcos certificate of completion. We use public frameworks such as the OWASP Top 10 for LLM applications as working checklists. We do not award, and do not prepare you for, a third-party certification exam.

How long is the program and how much time does it take?

Eight weeks, part-time and live online, at about 6–8 hours a week including the live session and practice work.

Who is it for?

Security, SOC and detection professionals, and engineers who ship chat, RAG or agent features and need to secure them. Some IT, logging or security exposure is helpful. It is not a first course in cybersecurity fundamentals.

Are the sessions live, and when do US batches meet?

Yes. Sessions are live and mentor-led, with recordings if you miss a class. Weekend-morning batches are offered for learners in the US; other regional times are available. Exact days and times are confirmed on a counsellor call.

Do I need to be an AI engineer already?

No. You should be comfortable inspecting logs, prompts or application behaviour. Machine-learning or LLM-engineering experience is not required. Engineers who mainly want to deploy and evaluate production LLM services may also look at LLMOps & AI Evaluation.

What about fees?

Fees and EMI options shared on a counsellor call.

Is there a refund policy?

Yes. Datamarcos has a 7-day refund policy. If the program is not right for you, ask within 7 days of enrollment.

DATAMARCOS

Talk to a counsellor

Share a few details on datamarcos.com and a counsellor will talk through program fit, the next cohort, fees and enrollment.

Fees and EMI options shared on a counsellor call.

7-day refund policy: if the program is not right for you, ask within 7 days of enrollment.

Noida One, 303-304, Rasoolpur Nawada, Industrial Area, Sector 62, Noida, Uttar Pradesh 201309, India

www.datamarcos.com · hello@datamarcos.com · +91 88519 30644

This brochure does not promise placement, salary or promotion outcomes. It is not legal, immigration or official-government certification advice.