InfoQ | AI SECURITY & PRIVACY ENGINEERING PROGRAM

InfoQ Certified AI Security & Privacy Engineering Program syllabus

Five-week online certification cohort. Four hours a week.

Week-by-week syllabus

Each week pairs a curated QCon talk with a 4-hour live working session. You apply the week's frameworks to the security and privacy decisions in your own work, alongside senior engineers from different companies.

Week 1: Working with Sensitive Data + AI

- **Focus**: How do you work with large AI systems and still protect personal and confidential data? This week covers the strategies and tools to identify sensitive data in AI workflows and protect it. You will work through how to spot sensitive data flows and the basic privacy controls you can put around personal data before it reaches a model.
- **Sample discussion task**: Name the concrete privacy and security problems you want to address during the cohort, and the learning goals that serve both your organization's needs and your own career growth. Then work through how to identify sensitive data flows in a real workflow, and which basic privacy controls you can offer for personal data in AI workflows.
- **Weekly homework**: Test some of the open-source libraries and tools recommended in the session against your actual workflows. Build a few examples that show what works, and find edge cases or places where the tooling does not work well, so you have a real comparison rather than a vendor claim.

Week 2: Threat Modeling and Red Teaming

- **Focus**: To find the problems in your system, you have to model threats and red team real systems. This week is hands-on practice in thinking like an attacker: recognizing threats, prioritizing the ones that matter, and red teaming an LLM yourself.
- **Sample discussion task**: Work through which privacy and security threats apply when you use large AI models in data workflows, and how to assess and prioritize risk for those workflows. Decide which types of red teaming attacks and threat modeling are relevant to your own work and your learning goals.
- **Weekly homework**: Continue working through threat modeling methods that may apply to your work, such as LINDDUN and Plot4AI. Try red teaming automation, and keep working through notebooks on AI security.

Week 3: Necessary Controls: Guardrails, Data Flow Controls and Sandboxes

- **Focus**: There are many protections you can build into AI workflows. This week focuses on the priorities: guardrails, flow sanitization and controls, and sandboxes. You leave able to analyze where each control fits, which ones suit your product or use case, and where to start implementing them.
- **Sample discussion task**: Work through how guardrails actually work and which ones are relevant to your real-world use cases. Decide when data flow controls apply and how to add them to your architecture, and how sandboxes work and when you need them for self-running AI workflows such as agents.
- **Weekly homework**: Try at least one open-weight guardrail model or sanitization control against synthetic data you have built to resemble the real data you work with. Document what works and what doesn't. Alternative: change and set up your agent's sandbox, then see whether you can still get the agent to do something harmful or unexpected, and document the result either way.

Week 4: Observability, Testing, and Evaluations

- **Focus**: Stepping back: how do you know your controls are working, and that you prioritized the right threats and protections in the first place? This week covers how observability, testing, and evaluations give you visibility and reassurance that you are making systems safer.
- **Sample discussion task**: Define what privacy and security observability means for AI systems and workloads. Work through what data is useful and how to collect it while still protecting privacy. Discuss how to use machine learning evaluations and MLOps/AIOps testing to assess models, and the software around them, against privacy and security requirements.
- **Weekly homework**: Build a small evaluation suite based on your homework from previous weeks. Alternative: try Arize Phoenix or similar observability software and test whether you can log flows and add privacy engineering into the trace collection.

Week 5: Building out Governance and Auditing

- **Focus**: Who is responsible for which safety, privacy, and security efforts, and how do you manage that engineering work? This week looks at governance models and engineering leadership roles for growing privacy and security engineering in AI, while covering compliance and auditing duties. Most of the session is given to the capstone presentations.
- **Sample discussion task**: Work through which governance models, structures, and processes can support growing AI privacy and security at your organization. Discuss how to build trust and learning into your compliance and auditing functions, and what helps with organizational buy-in and communication.
- **Capstone Project**: The capstone is introduced in Week 1 and built throughout the program. The deliverable is a 20-minute group presentation followed by a peer discussion. End on a privacy or security question your group is still working out.