Building, Securing and Hacking Intelligent Agentic Systems v2026

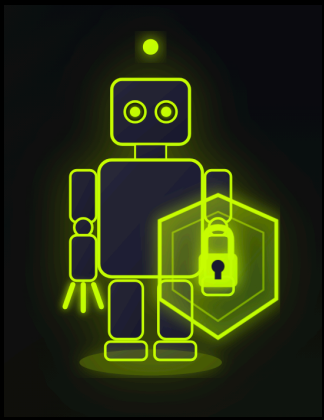

Target Audience

Beginner, Intermediate and anyone interested in learning more about LLMs, GenAI, Agentic AI, and the offensive and defensive security aspects of this ecosystem.

Introduction

This is an 3-day in-depth hands-on course designed for developers, engineers and security professionals who want to master the core principles behind Agentic systems aka AI agents.
We won't shy away from the required, sometimes theoretical concepts to grasp the technology, but also the security dynamics in play.

This course goes far beyond basic prompt engineering. It explores the low-level mechanics of LLM integration, AI Agents, MCP (Model Context Protocol) and the architecture behind all this and of course the security implications.

You'll start by interacting with LLMs using direct API calls, gradually progressing to SDKs, low-code interfaces, and full-fledged agent frameworks. We emphasise and focus on agentic design patterns, RAG and tool use, planning and decision-making to build agents that can reason, coordinate, and act in complex environments.

Our labs will notably focus on agent use in applications security, automation and DevOps operations but are applicable in any context.

Although OpenAI is used throughout the course for its accessibility and broad compatibility to explain and practice the concepts, we also cover emerging frameworks such as Google's ADK (Agent Developer Kit), A2A (Agent-to-Agent protocols), MCP (Model Context Protocol) and other open-source projects promoting interoperability across different models and providers. Notably MCP is a game-changing concept requiring detailed attention to avoid security pitfalls in favour of simplicity.

This course is continuously updated to reflect the rapid evolution of agentic AI, ensuring learners stay at the forefront of real-world, production-ready implementations.

While this is not a die-hard hacking course, it is designed to spark the mindset of a true hacker, someone who breaks things to understand them and questions defaults and thinks out-of-the-box.

This course will equip you to better understand how agentic systems work under the hood, justify and apply secure design patterns, and confidently engage with the next wave of AI-driven automation and lay a solid foundation for building your own agents (for fun and profit). This is your launchpad into the world of agentic AI with a hacking twist.

Agenda

Agenda subject to change.

Day 1

- Getting familiar with the concepts
- LLM Transformer model
- Chat completions API
- Responses API (and Assistants API)
- Python frameworks (OpenAI API and SDK, and LangChain)
- RAG (Retrieval Augmented Generation)

By the end of this day, you should be equipped to build simple but powerful agents that can solve real-world problems. An exciting lab will help you understand Linux system calls, and program execution flows in a way you never imagined.

Day 2

- Tool deep dive (the scary part)
- Model Context Protocol (MCP)
- Agentic Frameworks
- Agent Orchestration frameworks
- Guardrails

By the end of day 2, you should be able to build agents that interact with the `real-world` using data sources and tooling. The vast amount of tooling available will inspire every true hacker to start building agents and lay the foundation to integrate almost anything.

Day 3

- Agents in an enterprise setting
- An introduction to Model and Agent Evaluation
- Tracing and observability
- RAG pitfalls
- AI red teaming

During this last day, we'll explore how agents are used in the enterprise and point out the consequences of ignoring basic security hygiene.

Pre-requisites

- Willingness to explore new concepts
- Notebook with access to GitHub and use of SSH client
- Signup to Agentic-labs (Slack)
- Visual Studio Code or other code editor

- OpenAI API key required
  - Signup at https://platform.openai.com/docs/overview
  - Add a credit balance (10$)
  - Create an API key (sensitive, store it safely)

- Sign-up
  - Hugging Face https://huggingface.co/ (free)
  - Google AI Studio https://aistudio.google.com/apikey (free)
  - Serper https://serper.dev/signup (free)
  - MCPCloud https://www.mcp-cloud.ai (free)
  - GitHub Account (free)
  - Azure / AWS (optional)

- Coding skills? There is an AI agent for this.
- GitHub Copilot (if you have it available)
- ChatGPT (free), Claude Desktop, Gemini, etc.
- Cline / Roo Vscode extension (you can use your OpenAI key)
- Anthropic API key (PAYG, but relatively costly)

You will be provided with access to a GitHub repo, where you can clone the labs for you references and local reading. Labs will run on provided Ubuntu virtual machines hosted in AWS (and/or Digital Oceans). Access will be via SSH and public/private key.

Labs are tested on macOS, but this is at your own risk and might require that you install additional tooling like Python, Node, etc.

Protect you API keys and accounts as you do always!
Do not process any confidential, business or privacy related information data through the labs or the VMs. This is at your own risk and responsibility. Please do not misuse or use any adversarial tactics outside your own VM and lab environment.

Trainer – Speaker Bio and contact information

Philippe Bogaerts has over 9 years of hands-on experience in containerisation and Kubernetes, and more than 20 years in security and application delivery. He has built a solid foundation in designing secure, scalable, and future-ready architectures for cloud-native applications.

Today, his focus is on cloud-native security and AI security, areas he is deeply passionate about and actively investing in through continuous learning and hands-on development. He thrives in environments that challenge him to evolve, experiment with new technologies, and push the boundaries of what's possible.

As a technology advocate and leader, he enjoys guiding teams, growing businesses, and bridging the gap between deep technical topics and business outcomes. He is equally passionate about sharing his knowledge as a trainer, helping professionals and organisations strengthen their skills and prepare for the future. He believes that a strong mix of practical experience, curiosity, and a commitment to learning is essential for tackling the challenges of today and tomorrow.

You can reach me via the Slack channel after signup
Philippe Bogaerts
https://www.linkedin.com/in/philippebogaerts/
Mobile: +32473654689
Mailto: philippe.bogaerts@kubiosec.tech