ORIGIN ACADEMY · COURSE SYLLABUS

Agentic AI Engineer

Design, build and ship AI agents: LLMs, context engineering, tool use and MCP, retrieval, memory, evaluation and safe production deployment.

16 weeks · about 12 hours a week · live cohort of up to 10
Curriculum last updated 19 November 2026

Availability

This course opens soon. Join the waitlist on the course page to hear first about dates and prices.

Before you start

- Some programming experience (Python is taught in Phase 0)
- Comfort with the command line and web APIs helps
- English good enough to follow technical documentation and live sessions: about a credit (C6 or better) in WASSCE or NECO English Language, or an equivalent such as KCSE or IGCSE English (roughly CEFR level B1). You don't need the certificate: Phase 0 is free, so you can check for yourself
- A laptop running Linux, macOS, or Windows with WSL

How the course works

- A live cohort of up to 10 students, with a mentor.
- At least half your time is hands-on: every module ends with a lab.
- Live sessions are recorded, so you can rewatch them or catch up.
- Questions between sessions are answered within 24 hours on weekdays.
- A mentor reviews every phase project, like a pull request at work.

Phase 0 is free for everyone, before you enroll.

Curriculum

PHASE 0 · WEEKS 1–2 · FREE

Python for AI Engineering

Python, HTTP and REST APIs, async code, typing and working with data.

- Python essentials for engineers
- HTTP and REST APIs
- Async Python
- Typing and data validation
- Working with JSON and data

You build: API client toolkit

PHASE 1 · WEEKS 3–4

How LLMs Work

Transformers and tokens, context windows, sampling, prompting, structured outputs, model choice and cost.

- Transformers and tokens
- Context windows and pricing
- Sampling: temperature and top-p
- Prompting techniques
- Structured outputs
- Choosing a model: open vs closed weights

You build: Structured data-extraction service

PHASE 2 · WEEKS 5–6

Tool Use & MCP

Function calling, designing good tools, the Model Context Protocol, and running tools safely with sandboxing and permissions.

- Function calling
- Designing good tools
- The Model Context Protocol: hosts, clients and servers
- Building an MCP server
- Sandboxing and permissions

You build: Assistant that uses real APIs

PHASE 3 · WEEKS 7–8

Retrieval (RAG)

Embeddings, vector databases, chunking, hybrid search and citations.

- Embeddings
- Vector databases
- Chunking strategies
- Hybrid search and reranking
- Citations and grounding

You build: Document Q&A with citations

PHASE 4 · WEEKS 9–10

Agent Loops, Memory & Context

The agent loop, ReAct and planning, short- and long-term memory, context engineering and multi-agent patterns.

- The agent loop
- ReAct and planner-executor patterns
- Short-term and long-term memory
- Context engineering and compaction
- Multi-agent systems
- Frameworks vs building from scratch

You build: Multi-step research agent

PHASE 5 · WEEKS 11–12

Evals, Safety & Guardrails

Eval sets, deterministic and model-graded evals, prompt injection, PII and human-in-the-loop review.

- Building eval sets
- Deterministic and model-graded evals
- Prompt injection and jailbreaks
- Privacy and PII redaction
- Human in the loop

You build: Eval suite and red-team report

PHASE 6 · WEEKS 13–14

Production Agents

Observability and tracing, latency, cost control, streaming, deployment, monitoring, and licences for models, data and code.

- Tracing and observability
- Latency, streaming and caching
- Cost control
- Deployment and monitoring
- Licences for models, data and code

You build: Deployed agent with dashboards

PHASE 7 · WEEKS 15–16

Capstone

Build an agent that automates a real business workflow, to an enterprise brief.

- Scoping with the business owner
- Build and evaluate
- Red-team review
- Demo

You build: End-to-end workflow agent

Your certificate

You earn it by doing two things:

- Attend every live session. If you miss one, make it up by watching the recording and submitting that week's lab within 7 days. You can do this up to 2 times per course.
- Score more than 50% on the final evaluation at the end of the course.

Your Enterprise Readiness Rating is separate: it scores your projects, capstone and assessments so employers can see how job-ready you are.

Questions

Email academy@originclouds.io, or see the course online: https://academy.originclouds.io/courses/agentic-ai-engineer