

Course Outline

Power Up with Agentic AI

"From doing the work to managing the agents that do it"

What This Course Is About

This course teaches you to hand work over to AI safely: to stop doing the tasks that soak up your week and start managing the agents that do them.

It is the hands-on follow-on to AI Technology for Everyday Use. That course taught you to use AI well. This one assumes you can, and goes further: what agentic AI actually is, which of your workflows deserve an agent, how to build one, how to keep it honest, and what it costs.

Over five weeks you will build one real agentic workflow from your own life or work, and leave with it running, plus the judgement to know when it should not. Each session is **hands-on**: taught material, a live demo, then a supported build block on your own project.

The course is tool-agnostic. You choose Copilot, ChatGPT, Claude or Gemini in week 1 and stay with it. The ideas are the same in every tool, and the metaphors carry them: the model is the chef, the agent is the new hire, the context window is the desk, the harness is the scaffolding.

Who This Course is For & How it Works

This course is for graduates of AI Technology for Everyday Use, and for anyone at that level: you write multi-part prompts, you have built a custom helper, you have connected a tool or uploaded files, and you use AI most weeks.

You'll work with:

- One real time soak from your own week, carried through all five sessions
- A fit test and a mapping habit for deciding what to hand over
- Reusable job descriptions, checklists, stop rules and a runbook template
- Deterministic controls built in from the first session, not bolted on at the end

You won't be asked to:

- Code or script (the optional advanced track is the one exception)
- Share sensitive personal or work data
- Hand over judgement to AI, or leave anything running unattended

AI is treated as a **capable new hire**, not an authority. You are the manager.

Your 5-Week Learning Journey

(5 core sessions plus optional orientation and wrap-up)

Orientation (optional): Getting set up

“Tools, logistics, and a first look at the furniture”

- Understand what the course will (and won't) cover, and how the build blocks work
- Choose your tool for the five weeks and get it set up with help on hand
- Take the short self-check and start week 1 at speed

Session 1: Meet the New Hire

“What agentic AI is, and what it is not”

- Tell probabilistic from deterministic, and see live why the difference matters
- Name the furniture: model, context window, goals, tools, skills, plugins, MCP, memory, loops
- Place any tool on the ladder from prompt to workflow to agent, and audit your week for time soaks

Session 2: Which Work Deserves an Agent

“Choosing the right work to hand over”

- Apply a five-question fit test to any task before you automate it
- Map one real workflow step by step and mark each step human, fixed, or handed over
- Write an agent's job description with a clear definition of done, then attach the tool and skill your map needs as homework

Session 3: Give It a First Run

“Building the workflow in your own tool”

- Turn a mapped workflow into a working sequence with branches and loops
- Attach the tool and skill from your homework to the real workflow so the agent can act, not just answer
- Add the first deterministic controls: fixed steps, templates, and a hard stop on every loop

Session 4: Check the Work, Mind the Meter

“Keeping the agent honest and affordable”

- Put a harness around the agent: checkpoints, approvals, logs and evaluation, in plain English
- Recognise the four ways agents fail and design against each one
- Read the meter: what tokens cost, what your plan hides, and when the workflow must stop and ask

Session 5: When One Agent Is Not Enough

“Sub-agents, hand-offs, and keeping it alive”

- Decide when to split work across agents and when not to, and design the contract between them
- See what the advanced runtimes add for those who want to go further

- Write the one-page runbook that keeps your workflow alive after the course ends

Wrap-Up (30 minutes): Still Running?

“What worked, what broke, and what to build next”

- Report back on what is still running and what has already needed fixing
- Optional show-and-tell of finished workflows
- Choose calm, practical next steps for the months ahead

Technology Toolkit

You need a paid plan on one of these for the five weeks; agentic features sit behind the paywall in every tool.

- Microsoft Copilot, ChatGPT, Claude or Gemini, your choice, fixed in week 1
- The connectors, skills and scheduling features inside your chosen tool
- Optional advanced track: an open agent runtime such as OpenClaw or Hermes, demo-led, small metered spend possible
- A runbook template, a metaphor glossary, and recordings of every session

Learning Outcomes

By the end of the course, you will have:

- **One working agentic workflow** that handles a real time soak from your own week, with a runbook to keep it alive
- **A repeatable method** for deciding what to hand over, briefing an agent, and checking its work
- **Deterministic controls as a habit:** definitions of done, stop rules, checkpoints and budgets, so nothing runs unattended

Format and Fees

- Five weekly sessions of 2 hours, online via Zoom, recorded. Cohort of 10 to 12 so every build block has live support
- Optional 45-minute orientation the week before, and a 30-minute wrap-up after session 5, both included
- €300 self-funded · €600 employer-funded

The Promise to You

Each week you'll leave with something that runs, and a clearer sense of what to trust it with.

The agent does the work. You stay in charge of what it does, how you check it, and when it stops. This course is designed to be calm, practical, and people-centred, building real capability without pressure or hype.