

Course 05 — AI Employee Mastery

A 6-modules CyberG7 system — research-backed, build-as-you-go.

THE COURSE AT A GLANCE

The operator's path: not another prompt-engineering course, but the stance shift and the concrete six-phase build that turns you from a customer of someone else's chatbot into the owner of a hardened, productizable AI employee for ~\$50/month.

Who This Is For

The operator's path: not another prompt-engineering course, but the stance shift and the concrete six-phase build that turns you from a customer of someone else's chatbot into the owner of a hardened, productizable AI employee for ~\$50/month.

This course is built for:

- Solo founders and early-stage operators who are fluent ChatGPT/Claude users but have hit the consumer-tool ceiling
- Technical-curious non-developers who can follow copy-paste instructions and want to own their AI infrastructure
- Indie hackers and consultants who want to productize AI deployments as a service for clients
- Lean ops leaders at small teams looking to reclaim 8-15 hours/week of shallow operational work

You'll feel right at home if any of these sound familiar:

- Paying for multiple AI subscriptions yet still drowning in inbox, calendar, and repetitive morning ops
- Consumer chatbots that forget context, don't run on a schedule, and live behind a tab
- Fear that self-hosting an agent is too technical or too expensive to run 24/7
- Anxiety about security and data exposure when giving an AI access to email and calendar

What You'll Build

By the end, you won't just understand the ideas — you'll have assembled a working system, module by module. Across the course you'll develop:

- The Stance Shift, Not a Tooling Shift
- The Four Pillars Test

- The Founder Economic Case
- The Shallow-Ops Filter for First-Wave Delegation
- The Three Failure Patterns
- Glossary
- The Six Steps in Strict Order
- The 30-Minute Hardening Drill
- Brain-vs-Orchestrator Decoupling

Course Outline

Module 1 · The AI Employee Method — From Chatbot User to Operator

Most founders meet AI as a chatbot in a browser tab: they prompt well, they pay \$20/month, they extract real value, and then they hit a hard ceiling. Consumer tools do not remember across sessions, do not run on a schedule, do not initiate, and do not have access to the inbox, the calendar, or Notion. This module reframes the entire problem: the next move 99% of founders never make is to stop being a customer of someone else's chatbot and become the operator of their own AI employee. An AI employee is a software agent you hire, onboard, and manage, running on infrastructure you own.

Key themes:

- The Stance Shift, Not a Tooling Shift
- The Four Pillars Test
- The Founder Economic Case

Module 2 · Build Self-Hosted Agentic Infrastructure in a Weekend

This module turns the stance from Module 1 into running infrastructure. You provision a VPS, harden it in 30 minutes, install an orchestrator, connect a subscription-billed LLM brain, and lock the public surface behind a mesh VPN. The order matters: six steps in strict sequence, where skipping any one causes the deployment to fail by week two. The module ends with the 60-minute recovery drill — the practice that turns a working build into actual infrastructure rather than a hobby that dies on the first outage.

Key themes:

- The Six Steps in Strict Order
- The 30-Minute Hardening Drill
- Brain-vs-Orchestrator Decoupling

Module 3 · Wire the Channels Where Your AI Employee Shows Up

An AI employee that lives behind a login never gets used; the difference between a tech demo and an indispensable teammate is whether it meets you where you already work. This module wires three channels — Telegram for fast mobile interactions, the WebUI for deep work, and WhatsApp for client-facing deployments — and gives you a single decision rule so you never spend energy choosing which channel to use. It also sets the security baseline for every channel before any of them goes live.

Key themes:

- The Three Channels and What Each Is For

- The Channel-Decision Rule
- The Security Baseline for Every Channel

Module 4 · Configure Identity, Memory, Skills, and Tools — The 9-File Persona

A raw LLM is a generalist; this module turns it into a particular person doing a particular job. The mechanism is the 9-file persona spec — AGENTS, BOOTSTRAP, IDENTITY, SOUL, USER, MEMORY, SKILLS, TOOLS, HEARTBEAT — organized by a three-tier mental model of governance, personhood, and capability. This is what "configuration as code" means in practice: each file owns one dimension, so you can change skills without breaking identity or rotate API keys without disturbing the agent's ethics. The whole config directory is version-controlled, so every behavior change is a commit and any mistake rolls back in seconds.

Key themes:

- The Three-Tier Mental Model
- SOUL.md — The Hard Limits That Make Delegation Safe
- MEMORY.md — Three Tiers That Deliver Persistence

Module 5 · Operate the Four Office Workflows That Reclaim Your Week

This module is where the promise becomes a measured number. You wire the agent into Google Workspace and Notion, give it its own identity and minimal OAuth scopes, and ship four canonical workflows that run on a schedule: a 7am morning brief, task ingestion every two minutes, pre-meeting summaries at T-15, and email triage every 15 minutes. Then you run the 2-week measurement protocol that converts "this feels good" into a defensible figure — typically 8-15 hours per week reclaimed.

Key themes:

- The Agent's Own Workplace Identity
- The Four Canonical Workflows
- The Draft-Only Rule and the Trust Curve

Module 6 · Harden, Secure, and Productize Your AI Employee

The capstone makes the deployment production-grade and then turns it into a business. You implement three-tier secrets management, automated backups, and a 12-point quarterly security audit, all framed by the 2026 Lethal Trifecta threat model. Then you learn the AI-Employee-as-a-Service play: real 2026 agency economics of roughly \$2-5K setup plus \$500-1.5K monthly retainer per client, with the margin and client-count math that turns one operator into a six-figure practice.

Key themes:

- Three Hardening Pillars
- The 12-Point Quarterly Audit
- The 2026 Lethal Trifecta Threat Model

Outcomes

Complete the course and you'll be able to:

- Score your current AI setup on the Four Pillars (0-2 each); a total below 5/8 means you have a tool, not an employee.

- Provision the cheapest viable VPS first (Hetzner CX22 or DO Basic 4GB) — do not over-buy hardware before you have a working agent.
- Wire Telegram first; it is the lowest-friction channel and the one that builds the daily habit.
- Fill the 9 files in one sitting (30-60 minutes for the first agent) rather than tweaking prompts piecemeal.
- Give the agent its own Google Workspace account with minimal OAuth scopes (read + compose, never send).
- Move every secret into a three-tier flow (vault to .env to process) and schedule quarterly rotation.

ENROLL

Enrol Now — \$497 Start *Course 05 — AI Employee Mastery* today — the full module-by-module system lives at <https://ai-employee-mastery-edu.cyberg7.com.sg>.