

Before You Start

Read this after "What Is an AI Personal Operating System" — and before you open any guide or template.

This is not a form-filling exercise

The files in this training package look like templates. They have sections and headings and placeholders. It's tempting to open them, type your answers, and move on.

Don't.

The best AI Personal Operating Systems aren't built by filling in forms — they're built through conversation. When you answer questions in a live AI chat, the AI can probe, reflect back, challenge your assumptions, and surface things you wouldn't have thought to write down. When you fill in a template alone, you get what you already know. When you talk it through, you get more.

A quick word on the files — what's a .md?

Every file in this package (and in the system you're about to build) ends in `.md` — a **markdown file**. If you've never seen one, don't be put off: it's just a plain text file with a few simple symbols for formatting. `#` makes a heading, `**bold**` makes bold, `-` makes a bullet. That's essentially all of it.

Why not Word or PDF? Because these files aren't documents for people — they're instructions for AI. Every AI tool reads and writes markdown natively. Plain text means your system works in Claude today, any other tool tomorrow, and is never locked inside an app. Small, portable, future-proof.

Here's the part that matters: **you'll rarely open these files by hand**. In a working AI POS, the AI reads them at the start of every session and updates them when things change. When you want to see one, just ask — "show me my identity file" — and when you want to change one, say so — "update my rules." The AI is your editor.

If you do want to browse them yourself, they open in any text editor (Notepad included — the formatting symbols are visible but harmless). Free apps like Obsidian display them beautifully, but that's entirely optional — nothing in this training requires it.

What you need — read this honestly before investing the time

The full experience — where the AI reads your files automatically at the start of every session — requires an AI tool that can **connect to a folder on your computer**, and that capability generally sits on **paid plans** (Claude's desktop app is what this package is written and tested with; ChatGPT's desktop app has equivalent capability). If that's you, everything in this package works as described.

If it's not you, don't stop — the system has three levels, and **the files you build are identical at every level**:

1. **Connected folder** (paid desktop tool) — the full system. Files load automatically every session.
2. **Uploaded files** (most tools, including cloud-only ones like Gemini) — upload your Layer 1 files to a project or workspace; the AI reads them there. Nearly the full experience, minus live folder access.
3. **Personal instructions only** (any tool, free plans) — the compressed summary in your settings carries the essentials into every chat.

Connecting your folder (level 1). "Connected" means the AI can see and work in your AI Access folder directly. In Claude's desktop app, look at the message box on a new chat: switch the **Chat / Cowork toggle to Cowork**, then use the **folder selector beneath the message box** to choose your AI Access folder. Once selected, Claude remembers it for future sessions. App layouts change often — if what you see doesn't match this, just ask the AI: *"How do I give you access to a folder on my computer?"* It will walk you through its own current interface. Other desktop tools have equivalent steps.

Start at whatever level your tool and plan allow. Build the same files. If you upgrade later, your system is ready and waiting — nothing gets rebuilt.

We deliberately don't maintain a tool-by-tool capability list — AI tools change monthly, and a stale list is worse than none. The rule of thumb: if your AI can connect to a local folder, you're at level 1; if it can hold uploaded files in a project, level 2; otherwise level 3.

How to use these materials

Each file in this package is a guide — it tells you what to think about and what questions to explore. But the actual building happens in a chat session with AI.

The process for each file:

1. Open the relevant guide or template
2. Start a chat session with your AI tool

3. Copy the **kick-off prompt** at the top of the relevant template *together with everything below it*, and paste the lot into your AI as one message — every template starts with one, written for you. (To see what the conversation should feel like, read [2 - Build Guides/what_a_build_conversation_looks_like.pdf](#).)
 4. Let the conversation go where it goes — answer questions, push back, refine
 5. When you're happy, ask AI to produce the finished file
 6. Save it in the right place
-

For voice standards and reference files — upload your work first

For your voice standards, presentation standards, and document standards, the best starting point isn't answering questions — it's uploading examples of your actual work.

Attach at least 2–3 pieces of writing, a couple of presentations, or documents you're proud of to your chat — use the paperclip or + button in the AI's message box, the same place you type. Ask AI to analyse them. It will surface patterns you weren't consciously aware of — your natural sentence rhythm, structural habits, phrases you reach for, what you consistently do well. You validate what's right, correct what's off, and add what's missing.

This produces a far more accurate result than self-reporting ever will.

What to expect

Building a complete AI POS typically takes many hours across a few sessions — not because it's complicated, but because the thinking takes time. The identity and rules files usually flow quickly once you get into conversation. Voice standards take longer because they require examples and reflection. Reference files are the most individual — what "good" looks like is personal, and it takes a real conversation to surface it properly.

Don't try to do everything in one sitting. Build Layer 1 first — identity, voice, rules. Once those are in place, your AI already knows you well enough to be genuinely useful. Layer 2 can follow as you go.

The goal

By the end of this process you'll have a set of files that load into every AI session you open — so instead of starting from zero every time, your AI starts from context. It knows who you are, how you think, what you're working on, and how to work with you.

That's when AI stops being a tool you use occasionally and starts being something you genuinely rely on.

Let's build it.