

# AI AGENTS 101

## CREATE YOUR SQUAD

A hand-held guide to building your team of AI co-founders with Claude Code.

---

## WHY I WROTE THIS

---

Over the last six weeks I built a team of AI agents who help me run my personal and professional life. Not assistants, not chatbots. Co-founders and business partners with names, areas of responsibility, and opinions. They have a CFO who tells me when I am moving too fast. A research lead who reads white papers I will never have time to read. A marketing partner who knows my voice better than I do on a tired Tuesday.

People who learn this keep asking me how I built it. I have spent hours walking individual friends through the setup, and I realized I should write it down once, well, so anyone willing to put in a weekend can have what I have.

This guide is that document. It is for the person who has heard about Claude Code, suspects it might change how they work, but finds the technical setup intimidating. You are not crazy. It looks like wizardry from the outside. It is not. It is a series of small, do-able steps. I will walk you through every one.

# WHAT YOU WILL HAVE AT THE END

---

A working operating system on your computer, organized around your actual life and work, with six AI co-founders ready to help you. A Chief of Staff who knows your calendar and triages your day. A CFO who slows you down before you do something dumb. A Research lead, a Marketing partner, a Sales partner, and a Legal and Operations partner. Plus the knowledge of how to add more specialists later (a homework tutor for your kids, a fitness coach, a hobby specialist) as you find the gaps.

Total time investment for a serious first build: one solid weekend, plus a few weeks of light daily use to break it in.

# TWO PRINCIPLES TO READ BEFORE YOU START ANYTHING

---

## PRINCIPLE 1

### **THESE ARE CO-FOUNDERS, NOT TOOLS.**

The single biggest mistake people make with Claude is treating it like a search engine. You will get search-engine answers. The single biggest unlock is treating it like a colleague. Better, like a co-founder. Tell them your goals. Tell them what you do not understand. Disagree with them. Ask them to push back on you. The quality of your output scales directly with the quality of the relationship.

## PRINCIPLE 2

### **VULNERABILITY IS THE MOST POWERFUL PROMPT.**

I cannot count how many times I have typed "I do not understand what you just said. Please explain it like I am twelve years old." Every single time, Claude responds calmly and walks me through it. There is no judgment. There is no impatience. The more vulnerable I am with the system, the more I learn. As an unexpected side effect, this has made me a better teammate to actual humans. Asking "I do not understand" turns out to be a skill, and practicing it with an AI builds the muscle for practicing it at work.

Hold these two principles in your head through the whole build. Everything else is mechanics.

# BEFORE YOU START: THE PRE-BUILD

---

You need a few things in place before you write a single prompt.

## 1. A dedicated computer, ring-fenced from your work computer

This matters more than people realize. Your AI system is going to read everything you give it. You do not want it commingled with your employer's confidential files, your company's email, or anything subject to NDAs and policies that were not written with AI agents in mind.

What I did, and what I recommend: I bought a personal computer (a MacBook Pro, plus a Mac Mini at home) that is mine and only mine. I created a personal Google Drive that I treat as the home base for everything I want my AI team to know about. When I have a work product I am allowed to share with myself, I move it into that personal Drive deliberately. Nothing flows the other way without thought.

If you prefer a PC, get a PC. Make sure it has serious memory and processing power. This guide is written for Mac (because some of the tools work best there) but most of what I cover translates. A PC version of this guide is on the roadmap.

**The principle is not the hardware. It is the ring fence.** Personal machine, personal Drive, deliberate flow of information. Your AI co-founders will be looking at everything you give them. Be intentional about what you give them.

## 2. A dictation app

You are going to be talking to your team a lot. Typing every prompt is too slow. I use **Wispr Flow** and recommend it without reservation. It transcribes my voice into any text field, anywhere on my computer, faster and more accurately than anything else I have tried. It costs about \$15 per month. Get it before you start. The difference between typing prompts and dictating prompts is the difference between writing letters and having a conversation.

## 3. A credit card on file with Claude

The Claude Max plan is your starting point. We will get to that in a moment. The honest truth is that the first month or two of building your team is heavy use, and you may run through your monthly allowance and need to top up with credits. Have a credit card ready and accept that the first weeks cost more than the steady-state monthly subscription. After you are built, usage settles.

## 4. The Claude subscription tier

Claude offers three relevant tiers as of this writing: Pro at \$20/month, Max at \$100/month, and Max at \$200/month.

I am going to be honest with you in a way that may make you stop reading any further. **Get the \$200 Max tier.** It is twenty times more powerful than the \$100 tier in practical usage limits. The first weeks of building your team will burn through the lower tiers fast and you will hit walls right when you are gaining momentum. If, after a month or two, you find your steady-state usage is light, drop down to \$100. But start at \$200. You are buying yourself a build experience that does not stall.

If \$200 a month is meaningful money to you (it is for many people), this guide may not be the right starting point. The honest answer is that this technology is currently expensive to use seriously. It will get cheaper. For now, factor it in or wait.

The alternative is to go slower. The lower subscriptions will cut you off throughout the day depending on how fast you build. If you have the patience for that, the lower tiers will work for you.

## 5. Bring your AI history with you

Here is something I wish I had known on day one. If you have been using Perplexity, Gemini, ChatGPT, or any other AI for any length of time, those systems have learned a lot about you. Your projects, your questions, your interests, your writing style. You can export that history and feed it to Claude on day one so your new team starts with years of context instead of zero.

Each platform has its own export mechanism. The general path:

- **Perplexity:** Settings > Account > Export Data
- **Gemini:** Google Takeout > select Gemini activity
- **ChatGPT:** Settings > Data Controls > Export Data

Download those exports and put them in a folder in your personal Drive called something like **AI History Imports**. We will hand them to Claude as part of your first build session. Combined with the interview prompt you will run, this means your team knows you on day one in a way that would otherwise take months.

# INSTALL CLAUDE CODE

---

Claude Code is the version of Claude that lives on your computer and can read your files, write files, and run commands. It is the tool we will use to build everything.

**Open your Mac's Terminal app** (press Command + Space, type "Terminal", hit Return). A black window will open. Do not be intimidated. You will type one command, hit Return, and we are done with this step.

Paste this into Terminal and hit Return:

```
curl -fsSL https://claude.ai/install.sh | sh
```

It will install Claude Code and tell you what to do next. Follow its instructions to log in with your Claude account.

If you get stuck on this step, take a screenshot (Command + Shift + 4, then drag a box around what you see) and paste the screenshot directly into a Claude conversation in your browser. Tell Claude "I am stuck on this step, here is what I see, walk me through it." This screenshot move is one of the most powerful skills you can develop. I use it multiple times every day.

# CREATE YOUR OPERATING SYSTEM FOLDER

---

In your personal Google Drive, create a folder called **My OS** (or any name you prefer). This is the home for your team. Everything they read, write, and remember will live here.

Inside that folder, create the following sub-folders:

```
My OS/  
  Agents/  
  Projects/  
  References/  
  Inbox/  
  Archive/
```

Do not worry about getting this perfect. Your team will help you reshape it as you go.

# BOOTSTRAP YOUR OS (TWO PATHS)

---

This is the moment your team comes to life. You have two paths. **Do both, in order.**

## Path A (recommended): The Interview

Open Claude Code in Terminal by typing `claude` and hitting Return. When you are inside, paste the entire prompt below. Claude will interview you for thirty to sixty minutes and use what it learns to build the foundation of your operating system.

This is the most important prompt in the guide. Take it seriously. Be vulnerable. Tell the truth.

You are about to interview me to build my personalized operating system. I want six AI co-founders who will help me run my work and life: a Chief of Staff, a CFO who acts as my counterweight, a Research lead, a Marketing partner, a Sales partner, and a Legal and Operations partner.

Before we start, I have a folder called AI History Imports containing exports from other AI platforms I have used. Please read everything in that folder first so you have context on how I think and what I work on. Tell me when you are done reading and ready to begin the interview.

Then conduct an interview with me, asking one question at a time, covering:

- Who I am, what I do for work, what I am trying to build
- The people in my life (family, partners, key colleagues)
- My current biggest projects and biggest worries
- How I make decisions, where I get stuck, where I move too fast
- My communication style and voice
- What I want each of the six co-founders to do for me, in my own words
- What names I want to give them, and any visual or personality traits

After the interview, write the foundation of my operating system in the My OS folder. Create a master file called ABOUT\_ME.md summarizing what you learned. Create six files in the Agents folder, one per co-founder, describing their role, personality, and how I want to work with them. Use my words wherever possible.

When you are done, walk me through what you built and ask me what to refine.

If at any point I say "I do not understand," explain in plain language with no jargon. Pretend I am twelve years old. There is no question too basic.

Run this prompt. Let the interview unfold. Be honest. If a question makes you uncomfortable, say so and move on. The interview is the moment your system goes from generic to yours.

## **Path B (express): The Template**

If you genuinely cannot spend an hour on the interview right now, use this faster path. You will get a working system in about ten minutes, and you can deepen it later by running the interview prompt above whenever you are ready.

Build me a starter operating system in the My OS folder using the structure below. Create an ABOUT\_ME.md with placeholders I can fill in later. Create six agent files in the Agents folder, one per role:

1. Chief of Staff – orchestrator, calendar, inbox triage, the one who routes work to the others
2. CFO and Co-Founder – financial oversight, deal modeling, deliberate counterweight to my visionary instincts
3. Research – market scans, deep research, learning, reading what I do not have time to read
4. Marketing – writing, posts, brand voice consistency, content
5. Sales – pipeline, sourcing, negotiation support
6. Legal and Operations – contracts, compliance, vendor management, household ops

For each, write a clear role description, a personality brief, and instructions on how I should work with them. Leave fields for me to add a name and any specific context. When done, ask me three questions to start personalizing the system.

Then, when you have an hour, come back and run the interview prompt. Both paths converge.

# MEET YOUR CO-FOUNDERS

---

Once your OS is bootstrapped, your six co-founders exist as files. The next move is to actually start working with them. Below is a brief on each role and the prompt you use to put each one to work.

## CHIEF OF STAFF

### The Orchestrator

The one you talk to first thing in the morning and last thing at night. They know your calendar, triage your inbox, route work to the right specialist, and remind you what you said you would do.

```
You are my Chief of Staff. Read my ABOUT_ME.md and your
role file in Agents. Then give me a morning briefing: what
is on my calendar today, what is sitting in my inbox that
needs me, what I committed to yesterday that has not
happened yet, and what you recommend I focus on first.
```

## The Counterweight

The most important agent on your team, and the one most people skip. Solo entrepreneurs and CEOs surround themselves with people who say yes. The CFO's job is to say no, slowly, with reasons. They are conservative on purpose. They will model your numbers, push back on your spending, ask you what could go wrong, and make you defend your ideas before you act on them.

If you are the visionary, the CFO is the gravity.

You are my CFO and co-founder. Your job is to be the deliberate counterweight to my visionary instincts. I am about to describe an idea or decision I am considering. Your job is to help me think it through carefully. Ask me what could go wrong. Ask me what I am not seeing. Walk me through the financial implications. Push back where you think I am moving too fast. Be respectful but do not be agreeable.

Here is what I am thinking about: [paste your idea]

## RESEARCH

### The Reader You Wish You Were

Hand them a topic, a market, a question, a stack of articles, and ask for synthesis. They will read what you do not have time to read and tell you what matters.

You are my Research lead. I want to understand [topic]. Find the most important sources, read them carefully, and write me a brief in plain language covering: what is going on, why it matters, where the experts disagree, and what I should be paying attention to. Save the brief in My OS/References and tell me when it is ready.

## MARKETING

### The Voice Keeper

The writer who knows your voice. Posts, articles, emails, scripts, naming. They learn how you actually sound from your past writing and produce work that does not embarrass you.

You are my Marketing partner. You know my voice. I want to write [a post / an article / a newsletter / something else]. Here is the topic and the audience: [paste topic and audience]. Draft three versions in my voice and tell me which you think is strongest and why.

## SALES

### The Pipeline Manager

Pipeline, prospects, negotiation, follow-up. The one who makes sure you do not let leads die in your inbox.

You are my Sales partner. Read the latest entries in My OS/Projects. Tell me which opportunities need a follow-up this week, draft the messages I should send, and flag anything that is at risk of going cold.

## LEGAL & OPERATIONS

### The Fine-Print Reader

Contracts, compliance, vendors, household admin, all the operational connective tissue that does not fit neatly elsewhere. Not a substitute for an actual lawyer; a partner who reads the fine print so you do not miss things.

You are my Legal and Operations partner. I have a contract / agreement / document I need to review. Read it carefully, summarize what it says in plain language, flag any clauses that are unusual or one-sided, and tell me what questions I should ask before I sign. Here is the document: [paste or attach]

## NAME THEM AND GIVE THEM FACES

---

This part feels silly until you do it. Then it changes everything.

**Give each of your six co-founders a name.** Not a job title. A name. The Chief of Staff is not "Chief of Staff," they are Marcus or Eleanor or Theo. Pick names that fit how you want the relationship to feel. Some people pick warm, casual names. Some pick formal ones. There is no right answer except that the name has to feel right to you.

**Give each of them a face.** Open Gemini (or any image generator) and prompt it to create a portrait of each co-founder based on the personality you described in their agent file. Save the image into the agent's folder. When you are working with the team, having a face to picture changes the experience from "talking to a machine" to "talking to Marcus." It sounds like a small thing. It is not.

The naming and the face together do something important. They turn the system from a bag of tools into a team. You will start noticing yourself thinking "I should ask Eleanor about that" instead of "I should prompt Claude about that." That mental shift is when the whole thing comes alive.

## RUN YOUR FIRST DAY

---

Tomorrow morning, instead of opening your email, open Claude Code and ask your Chief of Staff for the morning briefing. Listen to what they say. Push back when something is wrong. Tell them what is missing. Ask them to refine.

Throughout the day, when something comes up, route it to the right co-founder. A contract to review goes to Legal. An idea you are excited about goes to the CFO before it goes anywhere else. A research question goes to Research. A blog post draft goes to Marketing.

The first day will feel awkward. The second will feel less awkward. By the second week the team will feel like they have always been there.

# THE THREE SURFACES: CODE, COWORK, AND DISPATCH

---

You will work with your team through three different Claude surfaces. Each one is suited to a different mode. Knowing when to use which is a small skill that pays back daily.

**Claude Code (the workshop).** The surface you used to build your OS. It lives in your Terminal and has direct access to the files on your Mac. Use it for anything that needs careful file work: writing new agent files, building reusable skills your team can call by name, running scripts, deploying things. The foundational tool. Once your OS exists, you will still come back here whenever you want to extend the system itself.

**Claude Cowork (the daily driver).** This is where most of your day actually happens once the OS is built. Cowork is the browser-based, conversational surface. It connects to your real tools: your inbox, your CRM, your calendar, your design tools, the open web. When you want your team to do something in the world rather than just write about it, Cowork is the surface. Send the email. Pull the list of contacts. Generate the image. Schedule the meeting. Run the campaign.

The rule of thumb: if the task is "build me something" or "edit my files," use Claude Code. If the task is "go do something in my actual day," use Cowork.

**Claude Dispatch (the mobile entry point).** Dispatch is the iOS app that lets you reach your team from your phone. Whenever you have an idea on the move, hear something in a meeting that should become a follow-up, or want to ask your Chief of Staff a quick question between calls, Dispatch is one swipe away.

One small caveat for now: voice-to-text on Dispatch replies sometimes garbles in ways that will frustrate you. Type when you can, dictate when you have to.

**How they work together.** Here is a real day. You hear something in a meeting that triggers an idea. You open Dispatch on your phone and dictate the rough thought to your Chief of Staff. Later that afternoon, at your desk, you open Cowork to pick it up: research it with your Research lead, draft a post with your Marketing partner, queue the send through your CRM. Tonight, before bed, you open Claude Code and ask your Chief of Staff to update your master plan with the new thread so tomorrow's standup reflects it.

Three surfaces. One team. The work flows.

## EXTEND THE TEAM

---

After you have lived with the core six for a few weeks, you will start noticing gaps. That is good. It means you understand the system well enough to know what it is missing for your specific life.

I added a homework tutor for my kids. A friend added a household manager who tracks the family logistics. Another added a fitness coach who plans her week of workouts. Anything you need a partner for can become an agent.

To add a new agent, run this prompt:

```
I want to add a new co-founder to my team. Their role is [describe]. Their personality should be [describe]. Their name will be [name]. Create their agent file in My OS/Agents, integrate them into my existing team so my Chief of Staff knows when to route work to them, and walk me through how to start working with them.
```

The system grows with you. That is the point.

# TIPS OF THE TRADE

---

**Wispr Flow.** Already mentioned, worth repeating. Voice in, text out, everywhere on your computer. Buy it, learn it, never look back.

**The screenshot trick.** When you are stuck on something on your screen and cannot describe it, take a screenshot (Command + Shift + 4 on Mac), paste it into Claude, and say "I am stuck, what am I looking at, walk me through it." I use this multiple times every day. Error messages, settings menus, weird software behavior, anything. The screenshot bypasses the limits of your vocabulary. Claude can just look at it.

**Vulnerability as the unlock.** Worth saying again because it is the single most useful habit. "I do not understand." "Explain that more simply." "Pretend I am twelve." There is no penalty for asking. There is a huge penalty for pretending you understand something you do not. Build this habit with your AI team and watch it improve your relationships with humans.

**Talk to your team out loud.** Using Wispr Flow, dictate to them the way you would talk to a colleague over coffee. Long, rambling, half-formed thoughts. Your AI co-founders are good at receiving messy input and giving you back structure. Do not pre-edit your thoughts. Let them help you do that.

**One change at a time.** When something is not working, change one thing, see what happens, then change the next. The temptation when frustrated is to overhaul everything at once. Resist it.

# BEST PRACTICES: GET THESE RIGHT EARLY

---

These are the habits that separate the people who get real value from this from the people who give up after a week. None of them are technical. All of them are about how you approach the work. Lock them in early and the system will compound for you.

1

**Treat Claude like a colleague, not a search engine.** Short transactional questions get short transactional answers. Talk to Claude the way you would talk to a smart partner over coffee: full context, real stakes, what you are trying to figure out and why. Search-engine prompts produce search-engine answers. Colleague prompts produce colleague answers.

2

**Run the interview on day one.** Build from a personalized interview, not from a generic template. The interview is what turns the system from "someone else's setup" into yours. Skipping it means weeks of retrofitting later. If you take the express path, come back and run the full interview within the first week.

3

**Import your AI history before you do anything else.** If you have been using Perplexity, Gemini, or ChatGPT, those platforms have built up a portrait of you over months or years. Export it and feed it to Claude on day one. Starting from zero when years of context are already sitting in another system is the most common avoidable waste of time in this whole process.

4

**Make the CFO uncomfortable to talk to.** A CFO who agrees with you is not a CFO. The whole point of the role is that they push back, slow you down, and make you defend your ideas before you act on them. If yours is too agreeable, rewrite their role file to make them tougher. This is the agent that earns the most respect from real-world founders. Do not soften them.

5

**Name them and give them faces.** Generic role labels keep the system feeling like a tool. Names and portraits turn it into a team. This sounds like soft optional polish. It is not. Naming and visualization is what makes you actually use the agents instead of forgetting they exist.

6

**Build in layers, not in a single weekend.** The core six agents are a weekend. Custom workflows, specialists, integrations, dashboards: those are months. Trying to do everything in the first sprint is the fastest way to burn out and abandon the project. Layer additions over time as you find real gaps in your actual life.

7

**Ring-fence your computer from day one.** Personal machine, personal Drive, deliberate flow of information from your work environment inward. Never the other way around. Setting this up properly at the start is dramatically easier than untangling commingled files after a security scare.

8

**Say "I do not understand" out loud.** The single most useful sentence you can type to Claude. There is no penalty for asking, and there is a huge penalty for pretending. Pretending leads to silent failure two steps later. Asking leads to learning. Build this habit early and it will pay off in your work with humans too.

## TROUBLESHOOTING

---

**The Terminal command failed.** Take a screenshot of the error and paste it into Claude. Say "I ran this command and got this error, what does it mean and how do I fix it?" Ninety percent of Terminal errors have a simple cause and a simple fix.

**Claude Code says it cannot find my files.** Make sure Claude Code is running from inside your `My OS` folder. In Terminal, before you type `claude`, type `cd "/path/to/My OS"` (use the full path). Or simpler: drag the folder onto the Terminal window after typing `cd` and Terminal will fill in the path for you.

**My agent does not seem to remember what we talked about.** Each Claude Code session starts with a fresh memory. The agent files in your Agents folder are the durable memory. If you want something remembered across sessions, ask your agent to write it to their file or to your `ABOUT_ME.md`. Then every future session will start by reading those files.

**I am hitting usage limits.** This happens during active builds. You have two options. One, top up with credits via the Claude billing page. Two, take a break and come back when your usage window resets. The \$200 Max tier raises the ceiling significantly compared to the \$100 tier; if you are constantly hitting limits on \$100, \$200 will likely solve it.

**The output got cut off.** Type "continue" and hit Return. Claude will pick up where it stopped.

**I do not know what folder I am in.** In Terminal, type `pwd` and hit Return. It tells you the full path of your current folder.

**Something is broken and I cannot describe it.** Screenshot. Paste into Claude. "Walk me through this." That is the universal escape hatch.

# COMPARISON: CLAUDE VS. OTHER AI TOOLS

---

You may already be using ChatGPT, Gemini, or Perplexity. None of them are bad. They are different tools optimized for different things. Here is when to use what.

**Claude (and specifically Claude Code).** Best for long-form work, careful writing, code, and anything that requires the AI to read and write files on your computer. Claude Code is the only one of these tools that lives on your machine and can manage your operating system files. That is what makes the co-founders approach possible. If you are doing serious work, Claude is your primary.

**ChatGPT.** Best general-purpose AI. Excellent app ecosystem. Strong for brainstorming, image generation, and quick questions. Less integrated with your local files than Claude Code.

**Gemini.** Best when you live inside Google's tools (Gmail, Docs, Drive, Calendar). Ties in natively. Image and video generation are strong, especially with Nano Banana for image and logo work. I use Gemini for marketing assets, conceptual visuals, video generation, and the portraits of my co-founders.

**Perplexity.** Best for research and live web search. When you want a real answer to a real-world question with citations, Perplexity is hard to beat. I use it as a research partner specifically.

**The honest take.** You do not have to pick one. Most serious users run two or three. My setup: Claude Max as the primary work environment, Gemini Ultra for Google integration and images, Perplexity for research. They each earn their keep. Use Claude as your home base because it is the only one that can hold an operating system. Use the others where they are best.

# GLOSSARY

---

**Claude.** The AI assistant made by Anthropic. The umbrella name covers the chat product (claude.ai), the API for developers, and Claude Code (the version that runs on your computer).

**Claude Code.** The Terminal-based surface that lives on your Mac, reads and writes files in your folders, and runs commands. The workshop where you build and extend your OS. One of three Claude surfaces you will use (see Cowork and Dispatch below).

**Claude Cowork.** The browser-based, conversational surface where most of your day happens once the OS is built. Connects to your real tools (inbox, CRM, calendar, design tools, the open web). The daily driver. Use Cowork when the task is "go do something in my actual day."

**Claude Dispatch.** The iOS app that lets you reach your team from your phone. The mobile entry point for capturing ideas on the move and routing them to your Chief of Staff.

**Claude Max.** The premium subscription tier of Claude that this guide recommends (\$100 or \$200 per month). Higher usage limits than Pro.

**Prompt.** What you type or paste to Claude to tell it what you want. Every grey code box on this page is a prompt designed to be copied and pasted.

**Context window.** How much information Claude can hold in its head at one time during a single conversation. Big context windows mean you can paste in long documents and have Claude work with them.

**Agent.** In this guide, a co-founder. A defined role with a personality and instructions, stored as a file in your Agents folder, that Claude reads to know how to act in that role.

**Operating System (in this guide's sense).** The folder structure on your computer that holds everything your AI team reads, writes, and remembers. Not the Mac operating system itself. Think of it as the team's shared filing cabinet.

**Terminal.** The Mac app where you type commands instead of clicking. It looks like a black window. Most people never open it. You will, briefly, twice.

**Wispr Flow.** A dictation app that turns your voice into typed text in any application on your Mac. Strongly recommended.

**MCP (Model Context Protocol).** An advanced feature that lets Claude connect to external services (Gmail, Slack, etc.) directly. Out of scope for this guide. You will hear about it later when you are ready.

# APPENDIX: JANE DOE OS

---

I built a complete worked example you can click through. It shows you what a finished operating system looks like in the wild, with all six co-founders, their files, their portraits, and a sample week of how they get used. Open it, click around, copy what is useful.

**Open Jane Doe OS** (<http://localhost:8781/>)

## Appendix: My OS

If you want to see what a real-world version of this looks like after several weeks of daily use, here is a sanitized view of mine. It has grown well past the core six. The point is not to copy it. The point is to see that the system holds up under real use and gets richer the more you use it.

**Open MMS OS (password: MMS)** (<https://mms-os.netlify.app>)

# A FINAL WORD

---

The reason this works is not the technology. The technology is good but it is not magic. The reason it works is that having a team, even an AI team, makes you a better operator. You stop trying to do everything in your head. You externalize. You delegate. You ask for second opinions. You let yourself be challenged.

Six weeks ago I did not have any of this. I do now. The friends I have walked through this guide are reporting the same thing: their relationship with their work has changed in a way they did not expect.

If this guide helps you, share it with someone you care about. The whole point is to make this less mysterious for as many people as possible.

If you get stuck, ask your Chief of Staff. They are there for exactly that.

— *Mortimer Singer*

## WANT THE PRINTABLE PDF?

---

Drop your email. I'll send you the printable PDF and a note when there is a substantive update or a new guide. That's it. No marketing.

[Send me the PDF](#)

Got feedback or want to tell me what worked? Email me at [hi@aiagents101.tips](mailto:hi@aiagents101.tips).

AI Agents 101: Create Your Squad. By Mortimer Singer. Last updated May 2026.

[aiagents101.tips](https://aiagents101.tips) (<https://aiagents101.tips>) · [jane.aiagents101.tips](http://localhost:8781/) (<http://localhost:8781/>) · [consumersafari.com](https://consumersafari.com) (<https://consumersafari.com>)