


The AI Agent Fleet Playbook

My business runs on a fleet of Claude Code agents: a reel agent that ships 2 videos a day, an email agent that drafts every reply, a trend agent that spots what's next, and one dashboard directing them all. This is the system for going from one agent to a fleet — the split, the trust ladder, the log, the dashboard.



01
01Rung 1
Finish one agent before you add one

A fleet of half-built agents is just a pile of scripts. One agent is **finished** when three things are true: it runs on a **schedule** without you, it **logs** what it did, and you haven't touched it for a week. Mine was content: 2 reels a day, on a timer. Until your first agent passes that test, adding a second only doubles your babysitting.

02
02The Split
One job, one brief, one output
Never brief an agent to "do marketing". Split by job: **research**, **script**, **publish**, **email**, **follow-up**. Each agent gets its own 20-line brief file in an **agents/** folder and one place its output lands. When something breaks — and it will — you know exactly which brief to open. Fleets scale by adding files, not by making one agent smarter.

03
03Trust Rule
The trust ladder
Every agent sits on one of three rungs: **DRAFT** (writes, saves, touches nothing live), **APPROVE** (queues it, a human clicks send), **AUTO** (ships and logs). Anything outward-facing — email, DMs, posts — starts on DRAFT, no exceptions. An agent earns the next rung after ~20 clean runs. My email agent still drafts every reply and **never sends**. A mistake on DRAFT costs nothing; the same mistake on AUTO repeats itself at scale.

04
04Visibility
One log, or no fleet
Every run appends one line to a single log: **agent · time · what shipped · link**. From day one, no exceptions. What you can't see, you can't trust — and you can't promote an agent up the ladder without its track record. My daily check is one command: **gh run list** — three green runs before 1pm today, none of them mine.

05
05Scale
Put Mission Control on top
At 3+ agents, open the log in one page: every agent, its rung, its last run, its last output. Five minutes each morning: scan, approve the APPROVE queue, fix anything red. That's the whole unlock — you stop doing the work and start **directing** it. The dashboard isn't a vanity screen; it's what lets the fleet grow without growing your hours.

The fleet rules · paste at the top of your CLAUDE.md

Copy


- # FLEET RULES — every agent reads this before every run.
1. You have ONE job. It is in your brief file. Do nothing else.
 2. Log every run to log.md: name · time · what shipped · link.
 3. Your trust rung is in your brief: DRAFT / APPROVE / AUTO.
DRAFT = write it, save it, touch nothing live.
APPROVE = queue it, a human clicks send.
AUTO = ship it, then log it.
 4. When unsure, drop to DRAFT and say why in the log.

Want this running for you?

Build it step by step inside AI Creative OS, with me in the thread.

[Implement this today →](#)

 Zdeno Mucina



Not ready for that?

Drop your email and 3 more guides open in the Library, free.

Unlock 3 more free →

[Download as PDF ↓](#)

Unlock all **56** guides

[Implement this today →](#)