

the ai employee playbook

The system behind Ada, the AI teammate that ships code for Classentra overnight. Six rules, the stack, and what broke. Every rule exists because something went wrong without it.

THE SIX RULES

- 1 Write specs, not wishes**
A wish: "add validation." A spec ends with **"done means:"** criteria a command or test can prove. Can't write that line? The task isn't ready to hand over.
- 2 Cut the work small**
One bounded slice per run. Small tasks fail small; big ones fail in ways you find three days later. Task size is your main risk dial.
- 3 Build gates the AI cannot skip**
Don't trust the AI, trust the gates. Nothing counts as done until tests, type check, lint, and build all pass. It never marks its own homework.
- 4 Review like it's a coworker**
Every morning, line by line, like a pull request. AI output looks confident whether or not it's right. Review is where you catch the difference.
- 5 Keep a learnings log**
Every failure becomes a dated note: what happened, what rule would have prevented it. Feed it back into the standing instructions. This is the step that compounds.
- 6 Keep the final door human**
Staging all night long. Production waits for one specific word from a human, and nothing automates around it.

THE SPEC TEMPLATE

Goal (one paragraph) · **Done means** (verifiable criteria) · **Out of scope** · **How to verify** (exact commands)

THE STACK (NOTHING EXOTIC)

- **Brain:** Claude via Claude Code. Any capable agentic CLI works; the loop matters more than the brand.
- **Queue:** a markdown file of tasks, each pointing at a spec and a plan.
- **Loop:** a script that watches the queue and hands tasks to the AI on their own branch.
- **Gates:** tests, strict types, lint, production build. All four, every time.
- **Ping:** a Telegram bot. The 3am buzz.
- **Cost:** a Claude subscription plus roughly zero infrastructure.

WHAT BROKE (SO YOURS DOESN'T)

The overnight shutdown

Machine died mid-run, orphaned half-finished work. Now the loop checks for stale runs on startup and recovers them.

The wrong branch

A cleaned-up branch meant a whole feature quietly landed where review would never see it. Now the loop checks where it's standing before its first commit, every time.

The zombie processes

Dev servers that never died stacked up until the machine crawled. Nothing gets assumed dead; kill by process ID.

NO CODE? SAME METHOD

- One standing instructions doc: who you are, what good looks like, what to never do.
- A spec with a "done means" line for every real task.
- Tasks small enough that a bad output costs minutes, not days.
- A checkpoint between AI output and the real world. Always.
- Every miss becomes a note, fed back in.
- Review everything that goes out with your name on it.

The instructions are the impressive part, not the AI. She's only as good as the spec, and getting the specs right was the actual engineering. Full method, starter prompt included: the blog post below.