HNST · FREE PLAYBOOK

The HNST AI Setup Playbook

Our whole setup, free.
Copy anything.

- 7 · An outside judge
- 6 · A desk and a brain
- 5 · Connections
- 4 · A team of sessions
- 3 · Skills and commands
- 2 · An Ask Ledger
- **1 · A rules file (start here)**

EVERY PROMPT, ONE TAP TO COPY

tools.hnstgroup.com

FREE TO USE, COPY AND SHARE

BEFORE YOU START

Why we give it away

HNST is a small company. One founder, plus a team of Claude chats doing real work every day. It took a lot of trial and error to get this setup right. We don't think that knowledge should sit behind a paywall. So here it is: the method, the rules, the structure and the templates. Take what helps you.

What you get, and what you don't

IN HERE

- ✓ The method (7 layers)
- ✓ The rules we run on
- ✓ How the team is shaped
- ✓ Templates and starter prompts

NOT IN HERE

- × Our private files
- × Names, clients, accounts
- × Money, numbers, links
- × Anything about our people

How to use this PDF



Start with layer 1
Write your rules file.
Everything else builds on it.

Add one layer a week
Paste the starter prompt.
Use it for a week. Then add more.

Check it worked
Each layer ends with
"You'll know it works when..."

Works with a Claude Project (paste rules into its instructions) or with Claude Code (a CLAUDE.md file). No coding needed for layers 1 to 3. Short on time? Paste the **Claude Setup Optimizer** prompt from tools.hnstgroup.com. It interviews you and builds a first draft of all seven layers.

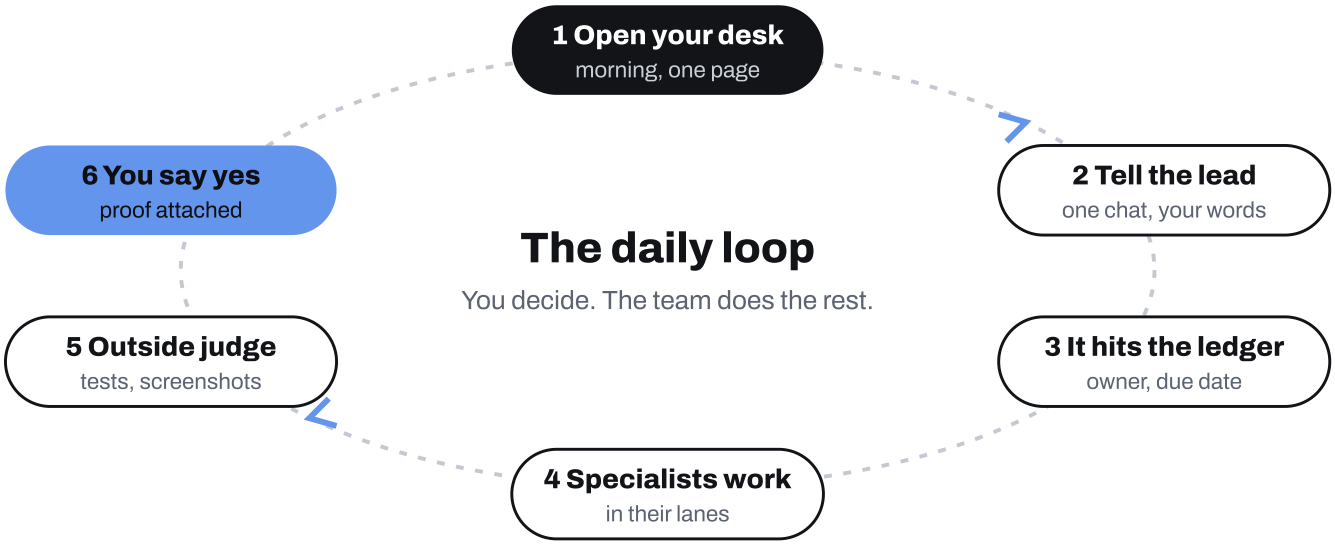
THE WHOLE SETUP ON ONE PAGE

The 7 layers

Each layer stands on the one below it. Start at the bottom.

7 An outside judge	proves it's right
6 A desk and a brain	where you start the day
5 Connections	checks the real thing
4 A team of sessions	one lead, specialists behind
3 Skills and commands	same great result every time
2 An Ask Ledger	nothing falls through the cracks
1 A rules file	**START HERE · who you are, your lines**

How a day runs



1 A rules file

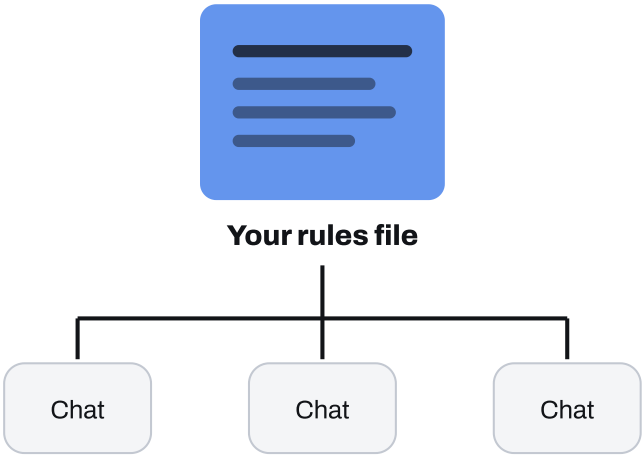
One file that tells Claude who you are, your goals, how to answer you, and what it must never do without asking. Every chat starts from it.

WHY WE DO IT

Without it, every chat starts from zero and you repeat yourself. With it, a brand-new chat already knows your goals and your lines.

HOW IT WORKS FOR US

- One file sits at the top of the whole setup. Every chat reads it first.
- Each rule has the date it was set, so we can see when and change it.
- Rules are testable: "report in 3 parts", not "be clear".
- New rules go in the day they're said. Old ones are marked "replaced", not deleted.



every chat starts already knowing you

STARTER PROMPT PASTE INTO CLAUDE

Help me write my rules file (I'll paste it into a Claude Project's instructions, or a CLAUDE.md if I use Claude Code). Interview me in short rounds, then write it with these sections: Who I am. My goals (12 months and this week). How to answer me (length, format, pictures vs text). Always do. Never do without asking me first (money, going live, sending messages as me, signing, accounts). Privacy lines. How to report back (exactly 3 parts: where we're at, where I need you, what you should know). Keep it under 600 words and make every rule specific enough to test. Date every rule so we know when it was set.

YOU'LL KNOW IT WORKS WHEN...

- ☑ A brand-new chat answers the way you like on the first try.
- ☑ You stop typing the same background into every chat.
- ☑ When something goes wrong, you can point to the rule that should have caught it, and fix that rule.

A fill-in-the-blanks version is on page 13.

2 An Ask Ledger

Every ask gets logged the moment you make it: owner, due date, status, proof. Done only counts with proof. Nothing falls through the cracks.



WHY WE DO IT

Asks get lost in long chats and in handoffs between chats. Written down with an owner and proof, they can't quietly disappear.

HOW IT WORKS FOR US

- Any chat that hears an ask logs it before starting work.
- Questions for you sit on the row: the question, options and a pick.
- A scheduled check a few times a day compares rows to reality and chases late ones.
- Rows are never closed just to shrink the count.

STARTER PROMPT PASTE INTO CLAUDE

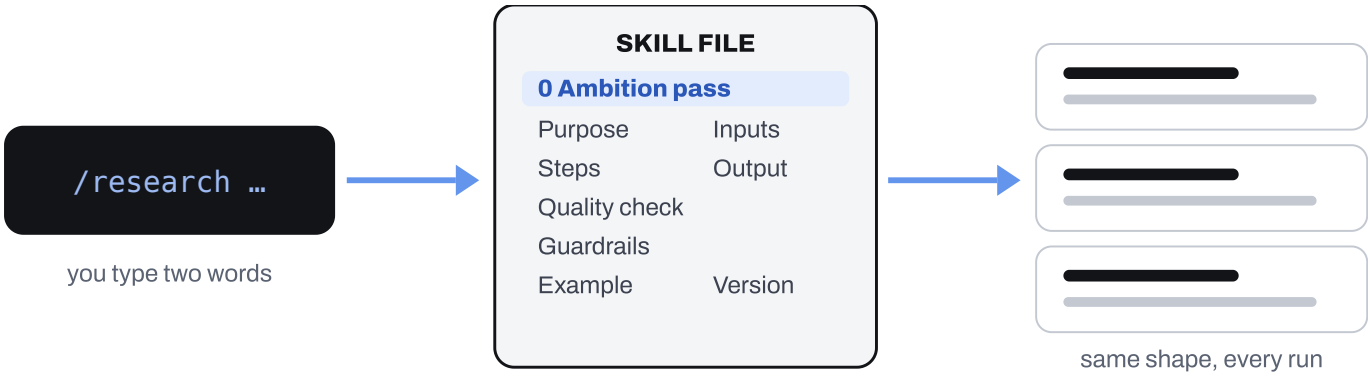
Build me an Ask Ledger as an interactive page (artifact) that saves its rows: columns title, area, owner, asked on, due, status (new, in progress, waiting on me, blocked, done), proof link, note. Waiting-on-me rows float to the top. A row can't be marked done without a proof link. Add a one-line 'how it was checked' field. Then add a rule to my setup: every time I ask for something, log it here first.

YOU'LL KNOW IT WORKS WHEN...

- [x] You can answer "what did I ask for last week?" in seconds.
- [x] Every "done" row has a link you can open.
- [x] Nothing you asked for turns up forgotten a week later.

3 Skills and commands

Your repeat jobs written once as skills, so a two-word command gets the same great result every time.



WHY WE DO IT

Good instructions are hard to write twice. A skill saves them once, so Monday's result and Friday's result look the same.

HOW IT WORKS FOR US

- Every skill uses the same 8 parts (filled example on page 12).
- Step 0 is always "deeper, not bigger" (see the rules on page 11).
- Each has a quality check done by something outside the model.
- Each has a version. A change gets a re-test before it's used.

STARTER PROMPT PASTE INTO CLAUDE

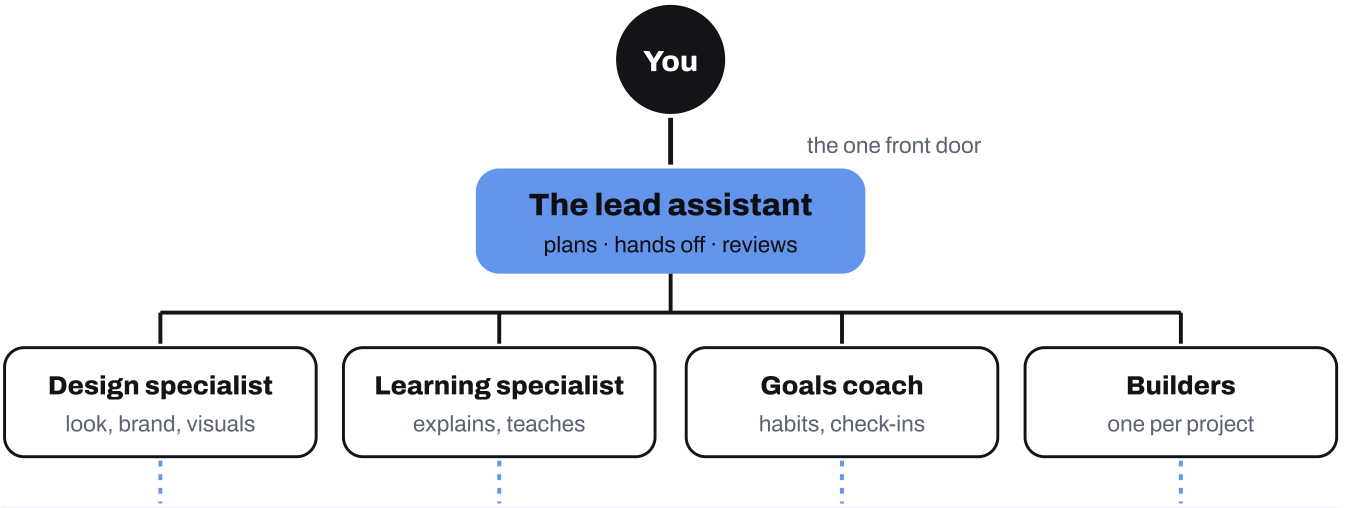
Turn my repeat work into skills. Ask me for my 5 most repeated tasks. For each, write a skill file with: name, a one-line trigger (when to use it), inputs, steps (step 0: ask what a top pro would do differently), output format, a quality check run by something outside the model, guardrails (what needs my OK), one example, and a version number. Then give me a one-page command list: the short phrase I type, what comes back.

YOU'LL KNOW IT WORKS WHEN...

- [x] Two runs a week apart come back in the same shape.
- [x] You type a short command instead of a long paragraph.
- [x] Every result ends with how it was checked.

4 A team of sessions

One lead you talk to. Specialists behind it for design, learning, sales, building. They report on the ledger; you only talk to the lead.



WHY WE DO IT

One chat can't hold everything. But many chats means you forget what you asked where. So: one front door, specialists behind it.

HOW IT WORKS FOR US

- Specialists never ask you directly. Questions go on the ledger; the lead brings them to you.
- Big jobs: the lead runs helpers side by side (builder, reviewer, fact-checker, tester), then reviews.
- Asked the wrong chat? It logs your words and passes it on. You never repeat yourself.
- The lead keeps one map of where everything is.

STARTER PROMPT PASTE INTO CLAUDE

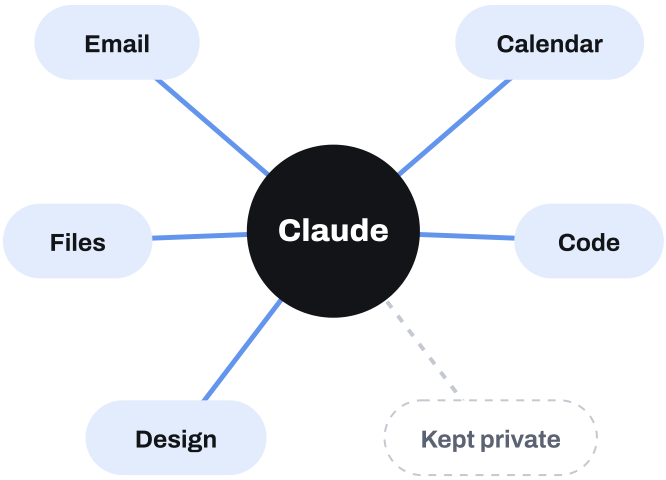
Design my AI team. I want one lead assistant I talk to, and specialist chats or projects behind it. Ask me what areas of my life and work need help, then propose: the lead (what it owns), 3 to 5 specialists (each with a lane, what it never does, and how it reports back), the rule that I only talk to the lead, and how work gets handed off without losing anything (a ledger row with my exact words). Give me the starting instructions for each one, ready to paste.

YOU'LL KNOW IT WORKS WHEN...

- [x] You open one chat a day, not five.
- [x] You never have to remember which chat did what.

5 Connections

Link the apps you already live in (email, calendar, files, code, design) so Claude checks the real thing instead of asking you.



WHY WE DO IT

If Claude can see your real calendar, it doesn't need to ask what's on it. Our rule: check before you ask. If an assistant can check it, that's its job, not yours.

HOW IT WORKS FOR US

- Connected apps get checked before any question comes to you.
- Every item on your list says how and when it was checked ("checked just now").
- Email gets drafts, never sends. You send.
- Some things stay disconnected on purpose. You pick them.

STARTER PROMPT PASTE INTO CLAUDE

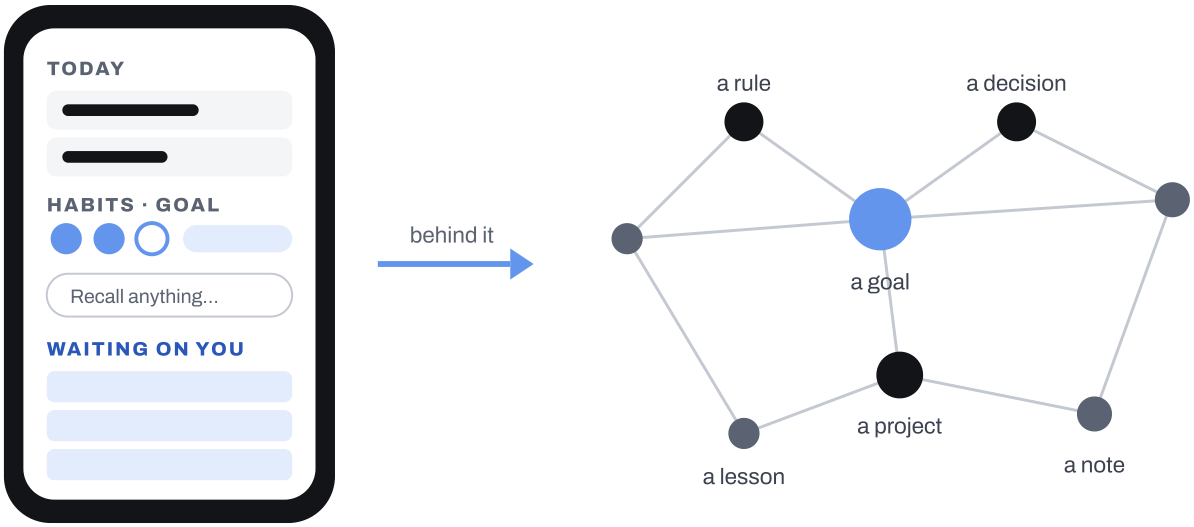
Help me connect Claude to the apps I already use. Ask which apps I use for email, calendar, files, notes, code, design and money. For each one, tell me: whether a Claude connector exists for it (check current docs), what it would let you do for me, what I should keep disconnected for privacy, and the exact steps to connect it. Then write one rule for my setup: 'Check connected apps before asking me anything.'

YOU'LL KNOW IT WORKS WHEN...

- [x] Claude asks you fewer questions it could have answered itself.
- [x] Items on your list say how they were checked, and when.
- [x] You know exactly what is connected, and what isn't.

6 A desk and a brain

A home page you open every morning: today, your goals, what needs you. Behind it, a map of everything you know, linked like a brain.



WHY WE DO IT

One calm place to start the day beats opening six apps. And the brain means an old decision is one search away, not lost in a chat.

HOW IT WORKS FOR US

- Top: today, a few daily habits, one big goal.
- Middle: one "recall anything" box.
- Bottom: only what's waiting on you, pulled from the ledger.
- Phone first, light and dark.

STARTER PROMPT PASTE INTO CLAUDE

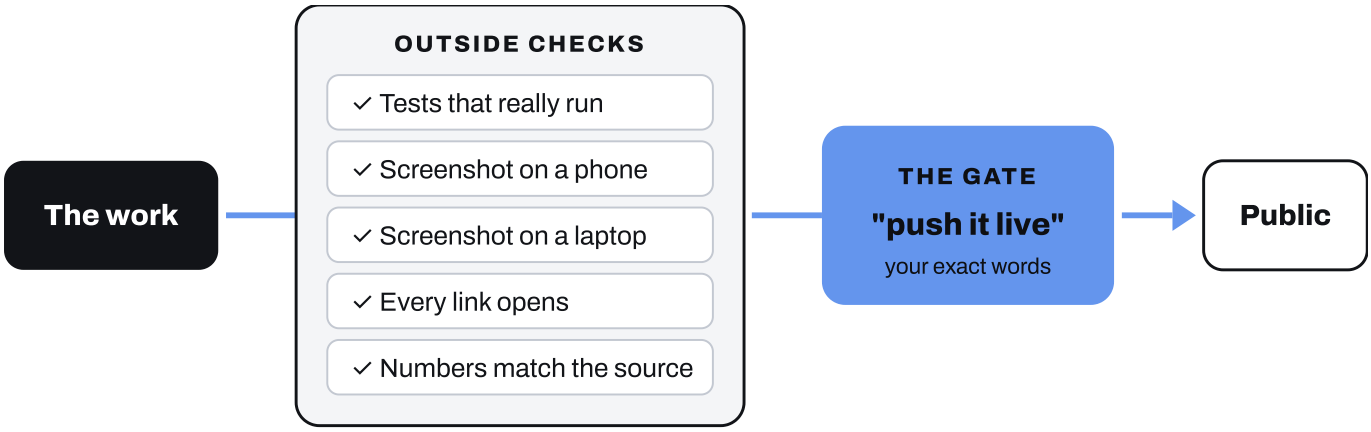
Build me a personal desk as an interactive page (artifact). Top: today's date, my events (from my calendar if connected), my 3 daily habits and one big goal with a countdown. Middle: a single 'recall anything' search box. Below: the 3 things waiting on me. Then build a second page: a knowledge map of my notes as a force graph (dots = notes, rules, decisions, lessons, projects; lines = links between them), clustered by area, that I can drag and search. Calm design, phone first, light and dark.

YOU'LL KNOW IT WORKS WHEN...

- [x] One page in the morning tells you your day.
- [x] You find an old decision with a single search.

7 An outside judge

Nothing ships because Claude says it's good. Tests, screenshots on a phone and laptop, live checks, real replies. Then you say the words.



WHY WE DO IT

A model grading its own work will usually say it's great. Proof has to come from something else: a test, a screen, a real reply.

HOW IT WORKS FOR US

- Every "done" says how it was checked, with the result.
- Anything a check can't reach is marked "unconfirmed".
- "Push it live" covers one item only. The next item needs its own yes.

STARTER PROMPT PASTE INTO CLAUDE

Add an outside-judge rule to my setup. Write it so that before anything is called done, Claude names and runs checks from outside the model (tests, a screenshot at phone and laptop width, a live link check, the source number, a real reply), shows the results, and says how it was checked. Anything public waits for my exact words 'push it live'. Give me the rule text plus a 6-item checklist I can reuse.

YOU'LL KNOW IT WORKS WHEN...

- [x] Every "done" comes with a line saying how it was checked.
- [x] Nothing goes public that you didn't say yes to.
- [x] You catch problems before your customers do.

THE RULES WE RUN ON

12 rules

Each one is a line in our rules file. Copy the ones that fit you.

1 One front door.
You talk to one lead assistant. So you never forget what you asked where.

2 Log every ask the moment it's said.
Handoffs lose asks. Your exact words go on the ledger first.

3 Done needs proof.
A link, a file or a screenshot. "I did it" is not proof.

4 Check before you ask.
If an assistant can check it, that's its job. Only approvals come to you.

5 Deeper, not bigger.
Before real work, ask what a top pro would do, start from raw sources, use real tools.

6 An outside judge.
Never let the model grade its own work. Tests, screenshots, live checks.

7 Nothing live without "push it live".
Sites, posts, stores: your exact words, on that one item.

8 Spending caps you set.
No new cost without your yes on the exact item and price.

9 Your name, your review.
Nothing is sent as you until you've seen the exact words. Drafts only.

10 Truth only.
No invented numbers, reviews or results. Unknown is written "unknown".

11 Pictures over paragraphs.
Three or more moving parts? Make a picture page, not a wall of text.

12 Phone first, ages 15 to 75.
Plain words, big taps, one thing per screen. If a teen and a grandparent can't use it, it isn't done.

RULE 5 IN FULL · "DEEPER, NOT BIGGER" PASTE BEFORE ANY REAL TASK

I want a deeper version of this task, not a bigger one. Before rewriting it: (1) say what a top professional in this field would do here that I'm not asking for; (2) name the rawest material they would start from instead of what has already been digested for me; (3) name the specific methods and tools they would use; (4) say what, other than you or me, would show the result is right. Then rewrite my prompt to ask for exactly that.

TEMPLATE · FILLED EXAMPLE

The skill template

Every skill we write has these 8 parts, in this order. Here is one filled in.

`/research [question]`	`v1.1.0`
PURPOSE	A decision made on solid ground. Not "find some links": an answer you could defend.
INPUTS	The question. Defaults: normal depth, answer today. Ask once, in one message, only for what can't be found or assumed.
STEP 0	**Ambition pass.** What would a top researcher do here? What's the rawest source? Which methods and tools? What outside the model proves it right? Do that version.
STEPS	1. Check the question's premise first. If it's wrong, say so at the top. 2. Go to primary sources: official docs, the original study, the real pricing page. 3. Money, legal or date claims need one primary source or two independent ones. 4. If a source is blocked, mark it and lower confidence. Never fill the gap with a guess.
OUTPUT	Same every run: **Answer** (3 sentences) · **Evidence** (each claim + link) · **Confidence** (high / medium / low, and why) · **What would change it** · **Not verified** · **How this was done** (4 lines).
QUALITY CHECK	Outside the model: every link opened, each key line quoted with its date, any math re-run in code. Written as pass / fail lines.
GUARDRAILS	No invented numbers ("unknown" instead). No paid sources without your yes. Nothing sent or posted. Private details stay out.
EXAMPLE	`/research what does it really cost to run a small website for a year?` Back comes: the answer up top, each cost line with its pricing-page link, confidence "medium" with the reason, and what would change it.
VERSION	1.1.0 · added the blocked-source check. 1.0.0 · first version. Small wording fix = patch. New step or check = minor. Output changes = major.

TEMPLATE · FILL IN THE BLANKS

Your starter rules file

Fill this in, then paste it into a Claude Project's instructions or a CLAUDE.md. Keep it under a page. Date each rule.

WHO I AM

Call me _____. I am a _____ and I spend my days on _____.

MY GOALS

In 12 months: _____. This week: _____.

HOW TO ANSWER ME

Keep answers to _____ sentences. Lead with the answer. If there are 3 or more moving parts, (show me a picture / give me a list). Use plain words.

ALWAYS

Log every ask on my ledger before starting. Check connected apps before asking me anything. Say how you checked. Write "unknown" instead of guessing. Also: _____.

NEVER WITHOUT ASKING ME FIRST

Spending money (my cap is _____ a month; every new cost needs my yes on the item and price). Going live (wait for "push it live"). Sending anything as me. Signing. Opening accounts. Deleting my things. Also: _____.

PRIVACY LINES

Never put these in anything shared: _____. Keep these apps disconnected: _____.

WHO DOES WHAT

I talk only to _____ (my lead). Specialists: _____. They report on the ledger.

HOW TO REPORT BACK

Exactly 3 parts: **Where we're at** (a few lines) · **Where I need you** (only what I must decide, ranked) · **What you should know** (money, risk, deadlines).

RULE LOG

Rule _____ set on _____ · replaced on _____ by rule _____

EVERY TOOL, ONE LINE EACH

The tool index

START HERE

Claude Setup Optimizer
Interviews you, then builds your own rules file, skills and daily routine so Claude works the way you work.

Deeper, not bigger
Rewrites any task the way a top pro in that field would attack it, before you start.

DAILY

Ask Ledger
Logs every ask the moment you make it. Nothing counts as done without proof.

Needs
Shows only the 3 decisions you personally have to make right now.

Status
Where things stand, checked live, on one screen.

Morning brief
A calm one-page start to your day: calendar, what's due, one focus.

Clean handoff
Hands a long chat to a fresh one without losing anything.

RESEARCH

Research with receipts
Any question answered with sources, a confidence line, and what would change the answer.

Should we use this?
Any app or service judged on cost, safety, fit, and the cheapest way to try it.

Study a video
Breaks down a video or screen recording: the tools in it, checked, and habits worth copying.

SALES

Draft in my voice
A reply, thank-you or follow-up that sounds like you, ready for you to send.

Proposal deck
A 10-slide swipe deck plus a short cover email for a client or partner.

Prospect card
Who they are, how well they fit, a warm way in, and a first message drafted.

Quote compare
Vendors ranked on the real job, inquiry emails drafted, quotes compared side by side.

BUILD

3 real options
Three truly different ways to solve something, shown side by side, with one pick.

Setup runner
Every step of a setup done or written out for you, up to the point where your card, ID or signature is needed.

Outside judge
Proves work is right with something other than Claude grading itself.

Go-live gate
Nothing goes public until it passes a real checklist and you say the words.

CONTENT

Post drafts
Ready-to-post drafts for the platform you pick. You post, on your schedule.

LEARN

Explain it like I'm 5, as a picture
Any idea explained with a big visual and very few words.

tools.hnstgroup.com
Every tool here, plus every layer's starter prompt, has a copy button. One tap, then paste into your own Claude.

Copy ▸