A SLOAN FIELD GUIDE · Nº 04

The Missing Half

YOUR SECOND BRAIN REMEMBERS · THIS IS THE PART THAT DELIVERS

You've seen the second-brain videos. Here's the half they skip: remembering was never the problem. This guide builds the rest, from a board that gets work out the door to a small AI staff that brings you only the decisions. The first version takes about thirty minutes. No experience assumed.

SECOND EDITION · OCTOBER 2026
SLOAN COACHING & CONSULTING
SLOANCOACHINGCONSULTING.COM

WHY THIS EXISTS

Your second brain needs hands

You've heard the pitch: build a second brain. Capture everything, organize it beautifully, never lose an idea again. The videos are good, and some of the systems are too. Here's the thing nobody in that genre says out loud: **a brain, on its own, doesn't do anything.** It remembers. Remembering was never your problem.

Meanwhile your actual AI use is one chat window. It's capable, tireless and wired to nothing. Its output has nowhere to live, no check on whether it's true, and no path to becoming something another person receives. One brilliant component, no body.

You may also have seen, or been, the second-brain failure case: a weekend of enthusiastic organizing, six months of drift, four hundred unconnected notes and three half-finished projects. Every system like this dies the same death, because every one assumes a human will do the maintenance, and humans don't. The fix is in Part Four, and it's the reason AI changes this genre: for the first time, the maintenance can be someone else's job.

So this guide builds the full body: a brain, hands, and a staff that brings you only what needs you.

PART	ITS JOB	WHAT I USE
The record, *the brain*	Hold what's worth keeping, in your words, outliving the tool that made it	Obsidian
The gate, *the brain's filter*	Keep the record from filling with confident nonsense	One line per note
The board, *the hands*	Turn what you know into things that reach people	Notion
The staff, *the reflexes*	Do the checking and the drafting while you're away	Scheduled tasks in Claude
The decision page, *the yes*	Hold everything that needs you, one press each	A page with a button per decision

The tools are swappable. The jobs are not. If you already like your note app and your task app, keep them and read this for the jobs. Bring your own AI, too. I recommend Claude, and ChatGPT works. Both can now run tasks on a schedule.

> **WHAT THIS GUIDE WILL NOT DO**
>
> It won't make you productive. It removes three failure points: forgetting, believing and never shipping. The last one stays hard after the system is built. Mine has run for months, and I still have finished things nobody has read.

The record

The job: hold what's worth keeping, in your words, in a form that survives the tool that produced it. This part *is* the second brain, the one the videos taught you, kept deliberately small because it's one part of a body and not the whole project. A chat log fails the job by design. It's a transcript, organized by nothing, and the good idea from three weeks ago is sitting between a meal plan and a question about a rash.

Set it up in fifteen minutes

1. Install Obsidian (free, obsidian.md) and create one vault. A vault is just a folder of plain text files on your own disk, which is the reason to pick it: no company's product decisions can take your notes away.
2. Make four folders, no more: `Inbox`, `Notes`, `Projects`, `Archive`. You'll want an elaborate taxonomy. Resist until you have two hundred notes and know what you actually have.
3. Turn on the Daily Notes plugin. Today's note is your inbox. Capture there, file later. One reliable capture point beats five clever ones.
4. When you're ready, point your AI at the folder so it can read and write your notes. If that's beyond today, skip it. The vault works alone, and the AI can join later.

The rule that makes it work

Nothing enters the record in the AI's words. Rewrite every note in yours, even badly. It's the only reliable test of whether you understood the thing. If you can't restate it, you didn't learn it, and storing it will feel like learning while producing none of it.

The cost: capture gets slower and you keep less. That's the trade. A smaller record you can use beats a larger one you can't.

The gate

The job: keep the record from filling with confident nonsense. The model never says *I'm not sure*, and early in your career that's more dangerous for you than for anyone else, because catching wrong answers is what experience normally does, and you don't have it yet. So you substitute discipline for instinct.

Set it up with one line per note

At the bottom of anything you keep, say how far you trust it and why. Four grades cover almost everything:

GRADE	MEANS	YOU MAY
Solid	Well established, checked, uncontested	Cite it
Contested	Real debate among people who know	Cite it with the debate named
Busted	Tested and didn't hold up	Cite it only as a correction
Mine	Your own inference, unverified	Say so, every time

Ask the AI to grade its own claims. It does this reasonably well. Then check the ones that matter. The habit outlives the tool, which is the real return: the person who can say *here's what's solid, here's what's contested, and here's what I'd do anyway* is valuable in every room, with or without a model.

> **THE TEST THAT CATCHES THE MOST**
>
> Ask for the source. Not a summary. The name, the year, the claim it actually made. Confident nonsense collapses at that question more reliably than at any other.

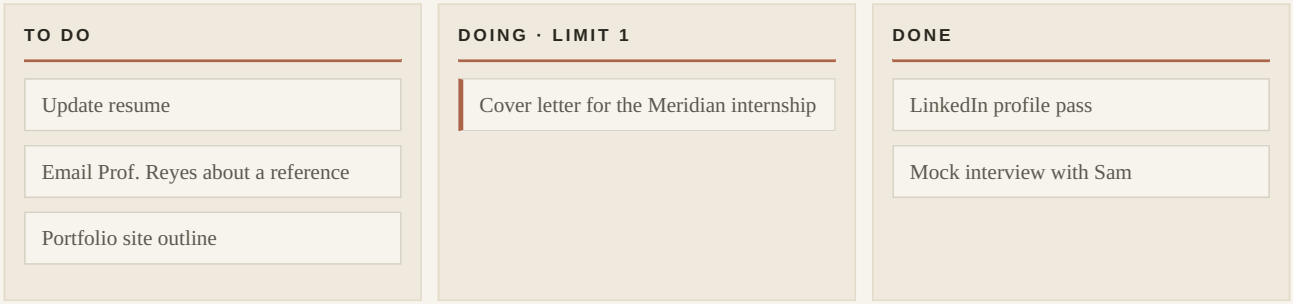
The board, from zero

The job: turn what you know into things that reach people. These are the hands, the part every second-brain guide skips and the reason this one exists. Capture and organize everywhere, deliver nowhere: that's the genre. And it matters most, because insight doesn't become action on its own. Give someone an engine that manufactures insight on demand and the gap between what they know and what they do gets wider.

The tool for the job is a kanban board, and the term deserves a real explanation. A kanban board is the simplest honest picture of your work: **every task is a card, every card sits in a column, and work flows left to right.** That's the whole idea. The personal version is two rules, worked out by Jim Benson long before AI made them urgent:

Make the work visible. Everything you've agreed to do goes on the board. Not most things. Everything. Commitments that live only in your head are the ones that ambush you.

Limit your work in progress. "WIP" is the count of things you've started and not finished. Cap it, hard. This is the rule people resist and the rule that does all the work: every open task charges rent in attention whether or not you're touching it, and switching between five of them burns real time before any work happens. **Start less, finish more.** Three things finished beat seven at eighty percent, because nothing at eighty percent counts. Nobody can read eighty percent of a cover letter.



A student's board, mid-week. One card in Doing. That isn't because only one thing matters. It's because only one thing is actually being worked. If you build nothing else from this guide, these three columns and that limit are a working system.

Two upgrades that earn their keep

Split "Done" in half

The most important distinction on the board: **Completed** means you finished it. **Accepted** means it reached the person it was for. The cover letter polished in your drafts folder is Completed. It's Accepted when it's *sent*. The gap between those two states is where work quietly dies, and a board that can't show the difference will hide your most expensive habit from you. Mine hid it for months.

COLUMN	THE QUESTION IT ANSWERS	THE COVER LETTER IS HERE WHEN…
Backlog	What have I agreed to, someday?	"I should apply to Meridian"
Defined	What's clear enough to start?	Post read, deadline known, plan made
In progress	What am I actually working? (limit: 1)	You're writing it
Completed	What's finished on my side?	Final draft, sitting in your folder
Accepted	What reached its person?	**Submitted**

Give work two honest exits

A board with no way out fills with zombies. **Parked** means on hold with a trigger, a condition that reawakens it ("when interviews start"), never a date you'll just move. **Killed** means you decided not to do it, with one line of reason. Killing a stale commitment on purpose isn't failure. Carrying it silently for six months is. If your Killed column is empty after a quarter, you aren't deciding. You're accumulating.

Set it up in twenty minutes

1. One board in Notion (free), or Trello, or paper. One, not a workspace of them.
2. Status columns: `Backlog → Defined → In progress → Completed → Accepted`, plus `Parked` and `Killed`.
3. Four more fields: Priority, Effort in hours, Due, Notes. Stop there until something breaks.
4. Load everything you're carrying, including what you've been avoiding. All of it, visible. Then hold the two limits: one card In progress, and a weekly hours cap where adding something means trading something out.

The staff and the decision page

A brain and hands with no reflexes are three good intentions, and this is the answer to the graveyard problem from page two. Organizing systems die because they depend on a human remembering to check them. The fix isn't more discipline. It's making the checking someone else's job, and AI is the first assistant that will actually do it daily without being managed.

The first version of mine was a morning brief: a scheduled task that read my board overnight and emailed me what was overdue, who was waiting on me, and what sat finished but never delivered. It worked until it didn't. The brief asked me to read a wall of text before I could decide anything, and on a hard day I didn't. Six briefs in a row went unanswered. The afternoon I moved every decision onto one page with a button for each, I answered fifteen of them in eleven minutes. The email retired that night.

So the reflexes are now two pieces. **The staff** is a few scheduled AI tasks, each with one job, that check and draft while you're away. **The decision page** is the one place you answer them, one press per decision.

Sort your work into three piles

Before you build anything, sort what you carry: what the AI can do by itself, what needs a decision from you, and what is yours alone to finish. The staff takes the first pile. The decision page holds the second. The board holds the third.

Build yours in about an hour

1. **Write the rules on one page, in your words.** What the staff may do alone, and what always comes back to you. Anything that leaves your desk (an email, a post, a payment) waits for your yes.
2. **Give each task one job and a schedule.** Start with one: read the board each morning, report what's overdue, **name who's waiting**, and draft anything that's due. Add a second only after the first has run for two weeks.
3. **Put every decision on one page.** Each item is one plain sentence with three answers: yes, hold or no. In Claude, that page can be an artifact with a button per decision. In ChatGPT, a scheduled task pauses for your approval before it sends. In any tool, a status column on your board works.
4. **Cap the queue at ten.** When ten decisions are waiting, the staff stops asking and only finishes what you've approved. Pull, not push.
5. **Let the alert point.** A short notification with a link to the page, never a wall of text.
6. **Undo until it acts.** A yes is carried out at the staff's next run, so you can take it back until then.

If your AI can't run scheduled tasks yet, a recurring reminder where you paste the board into a chat does the checking. The decision page still works as a filtered view of your board where everything waiting on you sits in one list.

Adding the next workflow

Every new workflow starts the same way, and the order matters.

1. **Name what keeps landing on you.** Following up after calls. Chasing an invoice. Getting ready for a meeting.
2. **Let the AI look before it proposes.** Point it at your real email, calendar and board, and ask how often this actually happens. Most people describe their work one way and do it another. Build for the way you do it.
3. **Write it as a card.** What starts it, what it reads, what it does, the one moment it needs you, and how you'll judge it.
4. **Run it in practice mode for two weeks.** Drafts only. Then keep it, change it or retire it.

A CARD, FILLED IN	FOLLOW UP AFTER EVERY CALL
When	Each evening at six
Reads	Today's calendar and your call notes
Does	Drafts a short follow-up for each call
Needs you	One press per email, before anything sends
Judged by	Replies within a week. Retired if you rewrite most of the drafts.

> **THE STAFF HAS A HEADCOUNT**
>
> Every workflow you add is one more thing that can break and one more source of decisions. Set a number, and when you hit it, adding one means retiring one. The structure's appetite has to stay smaller than the things it holds up.

> **THE MOVE MOST PEOPLE MISS**
>
> Everyone points AI at the half that was never the bottleneck: more ideas, more plans, more drafts. The bottleneck is your yes. **Make every decision one press, and point the staff at the doing side.**

START HERE

The thirty-minute version

A partial system beats a planned one. In order of return:

1. **The three-column board, limit one.** Ten minutes. The highest-leverage page in this guide.
2. **Split Done into Completed and Accepted.** Five minutes. It will show you something uncomfortable within a week.
3. **One staff task and one decision list.** A scheduled check if your AI can run one, a recurring reminder if it can't, and one place where everything waiting on you sits. Start with one question: *what reached another person this week?*

Build the record when the chat scroll loses its first good idea. Add the grades when the record grows past what you can remember. Add the second workflow when the first has earned its place.

What it costs

Setup, about an hour for the full version. Upkeep is real: minutes a day to file things in your own words, a few presses each morning, and a weekly review you'll skip and have to restart. Claude's scheduled tasks need a paid plan. ChatGPT's run on any plan, with more tasks on paid ones. The system doesn't create motivation and doesn't replace judgment. It removes excuses, which is less comfortable than it sounds, because once it's running, the reason something hasn't shipped is no longer that you forgot.

The words, in one place

Kanban board	Cards in columns. Work flows left to right. The whole idea.
WIP limit	A hard cap on how many things you've started and not finished. The rule that does the work.
Completed	Finished on your side. Not the end of the board.
Accepted	Reached the person it was for. The only column that counts.
Parked, Killed	On hold with a trigger, or dropped on purpose with a reason. Decisions, not failures.
The staff	A few scheduled AI tasks, each with one job. They check and draft. Nothing leaves without you.
Decision page	One place for everything waiting on you, one press each: yes, hold or no.
Practice mode	A new workflow's first two weeks. Drafts only, then keep, change or retire.

WANT IT SET UP WITH YOU?

Personal OS is this guide, built in your own accounts with me over about a month: the record, the board, your rules, a decision page and your first workflows, with a retro at day thirty. Bring your own AI. Claude is what I recommend, and ChatGPT works too. **start.sloancoachingconsulting.com**

ABOUT

Parker Sloan runs agile delivery and process work at a Fortune 500 company. Through Sloan Coaching & Consulting he works on the distance between knowing and doing, for individuals and for the teams they work in. He writes *The Architecture of Transformation*.

sloancoachingconsulting.com · **sloanpk.substack.com**