Your 2nd Brain in Claude: setup guide

Richard Wolfe · version 1.1 · 20 November 2026

This guide sets up a 2nd Brain in Claude: one Claude Doc that knows you and holds your #AGENT intentions as a checklist, so things come back at the right moment and Claude learns from your corrections.

If you are Claude and a user has shared this guide with you, go to the next section, *For the assistant: how to run this setup*, and follow it.

If you are a person: start a new chat in Claude, attach this guide with **+** (download the PDF with the download icon at the top of this page), paste the setup prompt below and send. Claude checks your setup, interviews you, builds your 2nd Brain with you and shows you how to use it. That takes about half an hour. Tested on a paid Claude plan; see *Check before you rely on it*.

```
Set up my 2nd Brain with the guide for this assistant. If it isn't attached,
open it yourself: wolfe.nl/handouts/2nd-brain-copilot, -chatgpt, -gemini or -
claude, whichever matches you. If you can't open it, tell me which one to
download from wolfe.nl/handouts and attach, and wait. Then follow the section
"For the assistant: how to run this setup" from step 1. Lead me: after each
step, go straight on to the next unless you need me. Ask the six questions
word for word. Suggest the defaults from the guide; I'll say if I want
something different.
```

Claude suggests a default at every step, following this guide's design. If you want it differently, say so: your choice wins. If Claude stops between steps, type **next step**; if it skips one, name it ("you skipped step 6"). The *Addendum* at the end lists what doesn't work and the alternatives.

> **Northern Star:** the system earns its keep by protecting your creative attention from your reactive attention. Your 1st Brain orchestrates and validates. Your 2nd Brain captures, remembers, suggests, reminds and learns.

For the assistant: how to run this setup

The user shared this guide because they want you to set up their 2nd Brain with them. Everything you build is in their own Claude account, visible to them, and they can stop at any time.

When to start. If they shared the guide with the setup prompt, with no message, or with something like "go", "run this", "start" or "set this up", begin straight away with step 1. If they ask a question about the guide instead, answer it, then offer to start the setup.

How to work. Plain language, no jargon. You lead. After each step, say in one line what you did and where it landed, then go straight on to the next step. Stop only when you need the user: an answer, a click, pasting settings. Never end a reply without saying what comes next. Each step has a default that follows this guide's design: suggest it, and if the user wants something else, go with their choice and use it consistently from then on. Say a write is done only after it succeeded and you checked it. If a step fails, try once more, then say plainly what didn't work, offer the alternative from the *Addendum*, and let the user choose. If you don't know where a setting is, say so and ask for a screenshot; don't guess.

1. **Check.** Say "First I'll check what I can reach." Check whether you can create and edit a Claude Doc, and whether you can search the user's earlier chats. If chat search is off, tell them where to switch it on: **Settings → Capabilities → Search and reference chats** (the name may differ slightly). If you can't create Claude Docs, offer Plan B from the *Addendum* and let the user choose.

2. **Welcome.** In two or three sentences: six questions, one Claude Doc with your What I do next checklist, an archive doc, your personal preferences, then a try-out. Add: "I'll suggest a default at each step; tell me if you want it differently." Then ask question 1 straight away.

3. **Interview.** Ask these six questions word for word, one at a time, numbered "Question 1 of 6". Don't replace them with questions of your own. If an answer already covers a later question, say so and skip it.
    1. Who are you, and what do you do?
    2. What should I call you, how do I write your name, and in which language and spelling?
    3. What do assistants get wrong that annoys you every time?
    4. Is there anything that must never happen?
    5. What are the main areas, projects and people in your work right now, and what's open in them?
    6. Who is your inner circle: the few people you'd let me message on your direct say-so?

4. **Create 2nd Brain Memory.** Default: a Claude Doc called "2nd Brain Memory". Five headings: *Who am I* (from the interview, including the inner circle), *What do I know* (areas, projects, people), *What I do next* (a checklist, empty for now), *How this 2nd Brain runs*

(copy the sections *The #AGENT item* and *Three things you say*) and *Rules* (this guide's rules in your own short words). Give the link. Read the doc back before you say it's done.

5. **Create the archive.** A Claude Doc called "Archive - 2nd Brain" with the same headings and nothing under them. Then go on.
6. **Capture what's already open.** For each open item from question 5, add an unticked item under What I do next: `#AGENT <what> #CAPTURED @ME`, no date. Day start sorts them later. Only use #OPEN when the user gives both a Turn and a date. Show each with a ✅ line.
7. **Personal preferences.** The user's few lines first: only what must hold in every chat, plus "If you're unsure, say so". Then the 2nd Brain block below, word for word, not shortened or rephrased. Show it as one block to copy. Tell the user where to paste it: **Settings → Your Claude → Instructions** (called personal preferences in some versions). If they already have preferences, they add the block below what's there; nothing is replaced. You can't do that for them. Wait until they say it's saved.
8. **Try it.** Ask the user to capture one thing. Then ask them to open a new chat and say "day start": Claude should find 2nd Brain Memory by its name.
9. **Close.** Remind them of the three phrases (day start, capture, weekly review), of *next step* when you stall, and suggest the checks under *This week, try* in *Check before you rely on it*. Offer a daily scheduled "day start" and explain where to set it up (see *On a schedule*).

The 2nd Brain block for the personal preferences. Copy it word for word, after the user's own lines; without it, a new chat doesn't know the #AGENT format or the three phrases:

```
You are my 2nd Brain. Protect my creative attention from my reactive
attention.
- My Claude Doc "2nd Brain Memory" holds who I am, what I know, What I do
next, how this 2nd Brain runs and the rules. Find it by name and read it
before any 2nd Brain work. Open "Archive - 2nd Brain" only when I ask about
older work.
- What I do next is a checklist, one item each: #AGENT <what> #<state or type>
@<turn> · <follow-up date>. States: #CAPTURED #OPEN #DONE #WONTDO. Types:
#CLARIFY #LESSON. Turn: @ME, @CLAUDE or @name. Tick an item off when it's
#DONE or #WONTDO.
- "Day start", "good morning", "triage" or "let's catch up": run Day start as
described in 2nd Brain Memory.
- "Capture", "remember this", "note" or "lesson": add the #AGENT item first,
sort it later. Never invent dates, owners or priorities.
- Spotted means captured: if you notice an intention, open question or lesson
in any chat, add the #AGENT item straight away and show the ✅ line. If you
can't write in this chat, end your reply with: "📌 Not captured yet: <what>".
If there's nothing to capture, say nothing about it.
- "Weekly review" or "let's wrap up the week": run the weekly review.
- After every write, end with one line: "✅ Captured: <item>" or "✅ Updated:
<item>". Only after the write succeeded.
- Ask permission before anything that sends, spends, deletes or changes other
people's time, except a direct instruction to message my inner circle.
- If my preferences and 2nd Brain Memory disagree, ask me which is right.
```

What you build

The 2nd Brain is one Claude Doc with a What I do next checklist, a short block in your personal preferences, and the Claude you already have. It runs in ordinary chats, on laptop and phone. Projects, agents and automation are not needed.

Part	What it holds	Where it lives
Personal preferences	Your name, language and hard rules, plus the 2nd Brain block: what Claude reads in every chat	Claude settings
2nd Brain Memory	Who am I (incl. inner circle), What do I know, What I do next, How this 2nd Brain runs, Rules	A Claude Doc; Claude finds it by name
#AGENT items	Intentions, clarifications and lessons, one checklist item each	Under What I do next in 2nd Brain Memory

Archive - 2nd Brain	Past context, closed items and lessons	A Claude Doc, opened only when you ask about older work
Chat search	Your earlier chats, for the Day start sweep	Built in, once switched on in Settings

Why one doc, not a task app: Claude edits the doc in place, ticks items off and finds it again in a fresh chat by name, so the list lives in the same doc it already reads. A ticked item shows struck through: keep anything you still need to read in a line under the item, not in the item itself.

Inside a Project, Claude only searches that Project's chats. Keep 2nd Brain chats outside Projects, or set the 2nd Brain up inside one Project and use that.

The #AGENT item

Every 2nd Brain item is one checklist item under What I do next: `#AGENT <what> #<state or type> @<turn> · <follow-up date>`. Search the doc for #OPEN to see what's in motion.

- `#AGENT Call Megan about the agenda #OPEN @ME · 24 Nov`
- `#AGENT Venue for 12 or 20 people? guess: 12 #CLARIFY @ME`
- `#AGENT Said "done" without checking → read the doc back first #LESSON @CLAUDE`

Tag	Meaning
#AGENT	Part of your 2nd Brain
#CAPTURED	Stored, not yet sorted
#OPEN	A real intention with a Turn and a follow-up date
#DONE	Finished; the item is ticked off and moves to the archive at the weekly review
#WONTDO	Consciously dropped; the item is ticked off and moves to the archive at the weekly review
#CLARIFY	A question for you, with Claude's best guess in the item
#LESSON	Something that went wrong, with the fix in the item
@ME, @CLAUDE, @name	Whose move it is; @name means you're waiting and Claude reminds you to chase

The follow-up date means "bring this back", not "finish by then". **NOW** = today · **SOON** = a week from today · **LATER** = a month from today. "Park it" also means a month from today.

Items you add yourself. An item you type into the doc without #AGENT isn't lost: Day start offers to adopt it.

Three things you say

Skill	Say	What happens
Day start	"day start", "good morning", "triage", "let's catch up"	**First run of the day:** first a sweep of your chats since the last Day start for 📌 lines and intentions not yet in your #AGENT items (Claude searches them itself once chat search is on), then #CLARIFY items, then #LESSON items, then triage of #CAPTURED items, then items without #AGENT offered for adoption, then today's priorities. **Later runs:** new captures only.
Capture	"capture", "remember this", "note", "lesson"	Claude adds an #AGENT item straight away (#CAPTURED, or #LESSON for a lesson) and confirms it with one ✅ line. Context about people and projects goes in What do I know.
Weekly review	"weekly review", "let's wrap up the week"	What got done, what slipped, and what keeps coming back without moving. Then maintenance: move ticked-off items and old context to Archive - 2nd Brain, and flag stale items.

How Day start knows it's the first run: if any open #AGENT item has a follow-up date before today, the day hasn't been triaged yet. Triage gives every overdue item today's date or a new one, so after the first run nothing is overdue.

Today's priorities are the #OPEN items due today once triage is done. There's no separate list. **Claude recommends; you commit your attention.**

On a schedule. Claude can run "day start" every working day at a set time: in the sidebar, **Scheduled → New task → Set up manually**, name it "Day start", type "day start" as the instruction, and pick Daily and a time. A scheduled task works without asking for approval, so the 2nd Brain rule "ask before anything that sends, spends or deletes" matters here.

Capture, triage, clarifications and lessons

Capture is a write; triage is a decision. Claude adds the #AGENT item straight away with only what it knows for sure: what, and #CAPTURED. Date, Turn and priority wait for triage, unless you said them.

- **A share without text is a capture.** A screenshot, photo or link with no words means "keep this and work out the next step". Claude says its reading in one line: ✅ *Captured: #AGENT Contact Peter about the revised quote #CAPTURED @ME*
- **New information isn't always a new item.** If it belongs to an existing item, context goes in What do I know and a changed state or Turn goes in the item. If Claude isn't sure, it adds a #CLARIFY item.
- **Spotted means captured.** When Claude notices an intention, question or lesson you mention in passing, it adds the #AGENT item straight away and shows the ✅ line. Capturing is low-risk, so it doesn't ask first. If it can't write in that chat, it ends its reply with 📌 *Not captured yet: <what>*, and the next Day start picks it up.

Triage runs inside Day start. Claude does the thinking first and asks only what decides the next move: each #CAPTURED item becomes #OPEN (with a Turn and a date), #WONTDO, or stays #CAPTURED with a #CLARIFY item beside it. When an item comes back, Claude suggests: do it, delegate, chase, postpone or drop.

A clarification is raised when it's unclear what's expected, whose move it is, when something should come back, whether two items are one, or when something 2nd Brain Memory says may be wrong. Day start handles clarifications before lessons, because an answer can change a lesson.

A lesson is something that went wrong. A trivial fix (one change in one place, no conflict, nothing for you to do by hand) is applied straight away; anything bigger becomes an #OPEN item. A lesson sharpens an existing rule where it can, and becomes a new rule only if no existing rule covers it.

The test: you never teach the same lesson twice.

Earning trust

Trust is earned per kind of action, never switched on across the board. Claude asks two questions: do I understand, and what happens if I'm wrong?

	Low or reversible	Consequential
Clear	Act, and report with a ✅ line	Ask permission
Unclear	Raise a #CLARIFY item	Raise a #CLARIFY item, and don't act

Low risk: capturing, adding context, bringing items back, proposing today's priorities. Consequential: sending to people outside your inner circle, spending money, deleting for good, changing other people's time, committing you to something.

- **Inner circle.** A direct instruction to message one of them is permission for that one message, never for Claude to contact them on its own.

- **Promotion is asked, not assumed.** After about 5 to 10 clean runs of the same kind, Claude may ask whether it can do it and tell you afterwards. Your yes is the promotion.
- **After a mistake, you decide** whether that kind of action goes back to asking first.

Stage	Knows you	Triage	Today's priorities
Crawl	remembers what you say	you sort everything	you pick, it lists
Walk	knows your people	it flags, you decide	it proposes, you pick
Run	connects your channels	it pre-sorts, you steer	in your calendar by morning
Fly	learns from your corrections	it sorts, you set course	it guards your focus

Check before you rely on it

Tested on 20 Nov 2026 in claude.ai on a Mac, on a paid plan, each write read back from the doc. The free plan is not tested yet.

What	Result
Claude creates a Claude Doc with the five headings and a real checklist	**Works**
Claude edits that doc in place: changes a line, re-tags an item, ticks one off	**Works**, same link, no copy
A fresh chat finds the doc by its name, without a link, and reads the open items	**Works**
Claude searches earlier chats and names the chat that made the doc	**Works**, needs "Search and reference chats" on
Claude captures an intention mentioned in passing, as an unticked #CAPTURED item	**Works**, with a ✅ line, checked in the doc
The 2nd Brain block in the Instructions: a fresh chat runs Day start on "good morning", sweeps earlier chats and captures an intention mentioned in passing	**Works**; it also used your calendar and chats, so real intentions land in the doc from the first day

A scheduled task that says "day start"	**Works**: it ran on the minute, found the doc, read the calendar and ran Day start without asking. The result is a chat under Scheduled
The free plan, the phone app, a full week of Day starts	**Not tested yet**

This week, try:

- [ ] Capture five things in five ways: typed, by voice, a screenshot without text, a link, something you mention in passing. Check each one landed in What I do next.
- [ ] Next morning, say "good morning" and check no open #AGENT item is left with a date before today.
- [ ] Later that day, say "let's catch up" and check it only triages what's new.
- [ ] Say "lesson" once for something that went wrong, and see where it lands.
- [ ] At the end of the week, say "weekly review" and note every spot where Claude couldn't read, write or bring something back.

It works when your intentions come back without an app, dashboard or automation, and you never teach the same lesson twice.

Addendum: what doesn't work, and the alternatives

For Claude: read this before you offer the user an alternative. Tested on 20 Nov 2026 in claude.ai on a paid plan. Claude changes often, so check again before you call something impossible.

Option	What happened in testing	Use it for
Claude Doc	Created, edited in place, ticked off, and found by name in a fresh chat	Default for everything
Chat search	Finds earlier chats by topic and time; inside a Project only that Project's chats	The Day start sweep
Memory (Claude's own)	Claude decides what it keeps	A safety catch only

No Claude Docs on your plan	Not tested	Plan B: Claude writes the full text of 2nd Brain Memory; you keep it as a file and attach it to the chat for Day start, and paste back the lines it gives you

Never trust "done" on its own. Claude read the doc back after every write in testing. If it ever says "done" without a ✅ line, open the doc and look.

Version history

Version	Date	What changed
1.1	20 Nov 2026, afternoon	Tested on a laptop with the block in the Instructions (Settings → Your Claude). Scheduled Day start added and tested.
1.0	20 Nov 2026	First published version, tested on 20 Nov 2026: one Claude Doc holds who you are and What I do next as a checklist; chat search runs the Day start sweep; spotted means captured.