Your 2nd Brain in Copilot: setup guide

Richard Wolfe · version 5.3 · 20 November 2026

This guide sets up a 2nd Brain in Microsoft 365 Copilot: one Word document that knows you, and To Do tasks tagged #AGENT that bring things back at the right moment and learn from your corrections.

If you are Copilot 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 Microsoft 365 Copilot, attach this guide with **+** (download the PDF with the download icon at the top of this page), paste the setup prompt below and send. Copilot checks your setup, interviews you, builds your 2nd Brain with you and shows you how to use it. That takes about half an hour. You need a Microsoft 365 Copilot licence; 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.
```

Copilot suggests a default at every step, following this guide's design. If you want it differently, say so: your choice wins. If Copilot stops between steps, type **next step**; if it skips one, name it ("you skipped step 6"). Each chapter ends with a short *If it gets stuck* list, and 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 Microsoft 365, 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, opening a file, 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. If a step fails, try once more, then say plainly what didn't work, offer the alternative from that chapter's *If it gets stuck* list or the *Addendum*, and let the user choose. You can't create or edit Copilot Pages or Loop: don't use them.

1. **Check.** Say "First I'll check what I can reach." List the user's To Do tasks. If you can't read To Do, say: this setup needs a Microsoft 365 Copilot licence with the Planner Agent switched on. Offer Plan B from the *Addendum* (everything in one Word document) and let the user choose.
2. **Welcome.** In two or three sentences: six questions, two Word documents, a few To Do tasks tagged #AGENT, the Custom instructions, 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 Word document called "2nd Brain Memory" in the user's OneDrive (it lands in Documents/Copilot/Created). Four headings: *Who am I* (from the interview, including the inner circle), *What do I know* (areas, projects, people), *How this 2nd Brain runs* (copy the sections *The #AGENT task* and *Three things you say*) and *Rules* (this guide's rules in your own short words). Give the link and ask the user to open it once.

If they want another name or folder, use that name everywhere, including the Custom instructions.

5. **Create the archive.** A Word document called "Archive - 2nd Brain" with the same four headings and nothing under them. Then go on.
6. **Capture what's already open.** For each open item from question 5, create a private To Do task: `#AGENT <what> #CAPTURED @ME`, no due 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. **Custom instructions.** 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 → Personalization → Custom instructions** (in the Mac app: **Settings → Personalisation → Customise Copilot → Manage**). 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": it should read 2nd Brain Memory by itself. Show them how to tap #AGENT in To Do to see only their 2nd Brain.
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*.

The 2nd Brain block for the Custom instructions. 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 Word document "2nd Brain Memory" holds who I am, what I know, how this
2nd Brain runs and the rules. Read it before any 2nd Brain work. Open "Archive
- 2nd Brain" only when I ask about older work.
- My intentions, clarifications and lessons are private To Do tasks titled:
#AGENT <what> #<state or type> @<turn>. States: #CAPTURED #OPEN #DONE #WONTDO.
Types: #CLARIFY #LESSON. Turn: @ME, @COPILOT or @name. The due date is the
follow-up date.
- "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": create the #AGENT task
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, create the #AGENT task straight away and show the ✅ line. If
you can't write in this chat, end your reply with: "📌 Not captured yet:
<what>".
- "Weekly review" or "let's wrap up the week": run the weekly review.
- After every write, end with one line: "✅ Captured: <title>" or "✅
Updated: <title>". 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 Custom instructions and 2nd Brain Memory disagree, ask me which is
right.
```

If the setup gets stuck (seen in a test run on 18 Nov 2026):

- Copilot stops after a step: type **next step**. It skipped one: name it, for example "you skipped step 6".
- Copilot asks its own questions instead of these six: answer anyway and add what the guide's question asked. What lands in the document matters more than the wording.
- The Custom instructions come out as one condensed paragraph: ask for "the 2nd Brain block from the guide, word for word".
- The new document doesn't appear when you type "/": the search index takes a while. Name it instead: "read my Word document 2nd Brain Memory". After typing "/", press Escape before you type a sentence, or your sentence goes into the file search.
- An older file with a similar name shows up under "/": rename or remove the old one, so Copilot reads the right file.
- Pressing Copy asks "Let Copilot read data you copy and paste?": Yes is safe. Or select the text and press Cmd+C.

- The document and your Copilot profile disagree on your name: Copilot asks which is right. Answer, then correct the document.

What you build

The 2nd Brain is one Word document, your To Do tasks tagged #AGENT, a short set of Custom instructions, and the Microsoft 365 you already have. It runs in ordinary Copilot Chat. Notebooks, agents, Pages and automation are not needed.

Part	What it holds	Where it lives
Custom instructions	Your name, language and hard rules, plus the 2nd Brain block: what Copilot reads in every chat	Copilot settings
2nd Brain Memory	Who am I (incl. inner circle), What do I know, How this 2nd Brain runs, Rules	Word document in your OneDrive
#AGENT tasks	Intentions, clarifications and lessons	Private tasks in To Do, in the default Tasks list
Archive - 2nd Brain	Past project context and closed lessons; the older an entry, the shorter it gets	Word document in your OneDrive, opened only when you ask about older work
Microsoft 365	Mail, Teams, calendar, meetings, files: the evidence	Where it already is; never copied
Copilot Memory	Whatever Copilot decides to keep	A safety catch; the system doesn't rely on it

The document explains, the tasks move. 2nd Brain Memory says who you are and how things are done, and hardly changes once things settle. The #AGENT tasks hold what's in motion and change every day. A ticked-off task is its own archive; only older context moves to Archive - 2nd Brain.

Why not Copilot Pages or Copilot Memory: Copilot can read Pages but can't write them, and Memory is too limited and too hard to check. Copilot writes To Do tasks and Word documents without asking for a click.

If this doesn't fit you:

- No To Do access for Copilot: Plan B. Put a section *What I do next* in 2nd Brain Memory and keep #AGENT lines there. Copilot edits the document in place. It's slower, and you can't tick items off on your phone.
- You want 2nd Brain tasks in their own list: not possible yet; Copilot always uses the default Tasks list. Tap #AGENT to filter.

- You prefer a table: a SharePoint list works, but every write needs a click on Confirm. See the *Addendum*.

The #AGENT task

Every 2nd Brain item is a private To Do task whose title carries everything: `#AGENT <what> # <state or type> @<turn>`. To Do turns each #word into a tag you can tap, so #AGENT shows your whole 2nd Brain in one view, on laptop and phone.

- `#AGENT Call Megan about the agenda #OPEN @ME`
- `#AGENT Venue for 12 or 20 people? guess: 12 #CLARIFY @ME`
- `#AGENT Said "done" without checking → check the file first #LESSON @COPILOT`

Tag	Meaning
#AGENT	Part of your 2nd Brain. Tap it in To Do to see only these
#CAPTURED	Stored, not yet sorted
#OPEN	A real intention with a Turn and a follow-up date
#DONE	Finished; the task is also ticked off
#WONTDO	Consciously dropped; the task is also ticked off
#CLARIFY	A question for you, with Copilot's best guess in the title
#LESSON	Something that went wrong, with the fix in the title
@ME, @COPILOT, @name	Whose move it is; @name means you're waiting and Copilot reminds you to chase

The due date is the follow-up date: "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.

Notes are written once. Copilot puts the source in a task's notes when it creates the task, but can't change notes afterwards. Anything that changes goes in the title.

Tasks you add yourself. A task you type straight into To Do, without #AGENT, isn't lost: Day start offers to adopt it by adding #AGENT, a state and a Turn.

If it gets stuck:

- Copilot wants to put a change in the notes: it can't. Ask it to change the title instead.

- A task is #OPEN without a due date: it never comes up as a priority. Say "give every #OPEN task a date".
- Copilot says it changed a task: tap #AGENT in To Do and look. Trust the ✅ line only once the task shows the change.
- You want a task gone: Copilot can't delete. It ticks it off as #WONTDO; delete it by hand if you like.

Three things you say

Skill	Say	What happens
Day start	"day start", "good morning", "triage", "let's catch up"	One workflow, two faces. **First run of the day:** first a sweep of your chats since the last Day start for 📌 lines and intentions not yet in To Do (Copilot can't find them in earlier chats yet, so paste any 📌 lines in), then #CLARIFY tasks, then #LESSON tasks, then triage of #CAPTURED tasks, then tasks without #AGENT offered for adoption, then today's priorities. **Later runs:** new captures and new untagged tasks only.
Capture	"capture", "remember this", "note", "lesson"	Copilot creates an #AGENT task 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 in 2nd Brain Memory.
Weekly review	"weekly review", "let's wrap up the week"	What got done (#AGENT tasks ticked off this week), what slipped, and what keeps coming back without moving. Then maintenance: tidy 2nd Brain Memory, move old context to Archive - 2nd Brain, flag stale tasks.

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

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

If it gets stuck:

- "day start" in a new chat does nothing special: the Custom instructions aren't saved. Check Settings, or say "read my Word document 2nd Brain Memory and run Day start". That works without Custom instructions too.
- Day start offers all your own tasks for adoption, one by one: say "skip the rest" or "adopt only tasks from this week".

- Day start ends without today's priorities: say "show today's priorities".
- Answer chips under a reply (NOW, SOON, LATER) only fill the message box: press Enter to send.

Capture, triage, clarifications and lessons

Capture is a write; triage is a decision. Copilot creates the #AGENT task straight away with only what it knows for sure: what, source (in the notes) and #CAPTURED. Due 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". Copilot says its reading in one line: ✅ *Captured: #AGENT Contact Peter about the revised quote #CAPTURED @ME*
- **New information isn't always a new task.** Peter's Teams message about pricing for 50 users is context for an existing task. Copilot checks the #AGENT tasks first; context goes in What do I know, and a changed state or Turn goes in the title. If it isn't sure, it adds a #CLARIFY task.
- **Spotted means captured.** When Copilot notices an intention, question or lesson you mention in passing, it creates the #AGENT task 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. Copilot does the thinking first and asks only what decides the next move: each #CAPTURED task becomes #OPEN (with a Turn and a due date), #WONTDO, or stays #CAPTURED with a #CLARIFY task beside it. When an item comes back, Copilot 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 in 2nd Brain Memory may be wrong. Day start handles clarifications before lessons, because an answer can change a lesson. Your answer updates the task or the document, and the #CLARIFY task is ticked off. When Custom instructions and 2nd Brain Memory disagree, Copilot can't tell which is right, so it raises a clarification instead of choosing: *#AGENT Custom instructions say Peter, Who am I says Petra: which is it? #CLARIFY @ME*

A lesson is something that went wrong. Say "lesson", and Copilot creates a #LESSON task with what went wrong and the fix.

- **Trivial fix: applied straight away.** One change in one place in 2nd Brain Memory, no conflict with other rules, nothing you have to do by hand. The #LESSON task is ticked off.
- **Anything bigger becomes an intention** (#OPEN), so you look at it properly.
- **Sharpen before adding.** A lesson updates an existing rule where it can. It becomes a new rule only if it's your decision, would still hold with another tool, and no existing rule covers it.

The test: you never teach the same lesson twice.

If it gets stuck:

- No ✅ line after a capture: maybe nothing was written. Tap #AGENT in To Do before you move on.
- Copilot asks for a date or owner when you capture: say "just capture it". Triage comes later.

Earning trust

You build a setup you can trust, step by step. Copilot may be as proactive as it likes; this guide neither holds it back nor pushes it.

Trust is earned per kind of action, never switched on across the board. Copilot 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 task	Raise a #CLARIFY task, 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. For the few people listed under Who am I, a direct instruction is permission for that one message: "tell Anouk I'll be ten minutes late" needs no second yes. It never lets Copilot decide to contact them on its own.

- **Promotion is asked, not assumed.** After about 5 to 10 clean runs of the same kind, Copilot may ask: *"I've done this ten times without correction. Shall I just 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

You need a Microsoft 365 Copilot licence with the Planner Agent switched on; Copilot Chat without a licence (Basic) can't run this setup. Everything below was tested on 18 Nov 2026 in the Copilot app on a Mac, each write checked in To Do or OneDrive itself.

What	Result
Copilot creates, renames, re-tags and ticks off To Do tasks	**Works**, no confirm click
#AGENT at the start of a title is a tag you can tap to filter in To Do	**Works**
Copilot finds #AGENT tasks by tag, and adopts a task typed into To Do by hand	**Works**
Copilot changes the due date, priority or status of a task	**Works**; the due date showed on the right day
Copilot changes the notes, start date or checklist of an existing task	**Doesn't work**: notes only when the task is created
Copilot puts tasks in a To Do list other than Tasks	**Doesn't work**: everything lands in Tasks
Copilot creates a Word document in OneDrive	**Works**
Copilot changes a line in place, adds lines and fills a section in that document	**Works** when the file is picked with "/"
A fresh chat finds a brand-new Word document	**Works when named** in plain words ("my Word document 2nd Brain Memory"); "/" didn't show it yet
Day start in a fresh chat reads 2nd Brain Memory and follows its order, including adoption	**Works**

The full setup, run by Copilot from this guide as a new user	**Works**, with nudges; see If the setup gets stuck
Copilot writes to Copilot Pages or Loop pages	**Doesn't work**: it can only read them
Copilot Chat without a licence (Basic)	**Doesn't work**: no To Do access; it creates Word files but can't edit them, and said "done" when it hadn't
The setup prompt with the Markdown guide attached (.md)	**Works**: Copilot accepts the file and starts at step 1. It keeps a copy in your OneDrive
The setup prompt without an attachment	**Works**: Copilot couldn't open the page itself, so it asked for the guide to be attached and waited
Copilot captures an intention mentioned in passing, or marks it with 📌	**Works with the block in the Custom instructions**: in an ordinary chat it ended with a 📌 line; say "capture" to have it create the task. Without the block it only said it had noticed
Day start finds 📌 lines and missed intentions in earlier chats	**Doesn't work**: Copilot found no 📌 lines in chats from the same day. Paste them in at Day start
Capture on the phone and by voice, Copilot in Teams, a full week of Day starts	**Not tested yet**

Sources: Planner Agent: what it can and can't edit · Customise how Copilot responds to you

This week, try:

- [ ] Capture five things in five ways: by voice, a screenshot without text, a flagged mail, a Teams message, something from a Copilot chat. Tap #AGENT in To Do and check each one landed.
- [ ] Add one task straight into To Do yourself and see whether Day start offers to adopt it.
- [ ] Next morning, say "good morning" and check no open #AGENT task is left with a due 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 Copilot couldn't read, write or bring something back.

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

Addendum: what doesn't work, and the alternatives

For Copilot: read this before you offer the user an alternative. Tested on 18 Nov 2026 on a Mac with a Microsoft 365 Copilot licence, and with Copilot Chat without a licence. Microsoft changes Copilot often, so check again before you call something impossible.

Option	What happened in testing	Use it for
To Do through the Planner Agent	Create, rename, re-tag, due date, priority, tick off and reopen work, with no confirm click. Notes after creation, start date, checklist, assignment, other lists and delete don't	Default for intentions, clarifications and lessons
Word document	Created in OneDrive, Documents/Copilot/Created. Edited in place when picked with "/"; found by a fresh chat when named in plain words. "/" lags behind a new file	Default for who you are and what you know; Plan B for everything
Word agent	Blocked by some organisations' policy	Skip it: plain chat does the same
Copilot Pages and Loop	Copilot can read them, but can't create or edit them	Nothing in the 2nd Brain
Copilot Memory	Copilot decides what it keeps; you can't rely on it	A safety catch only
SharePoint list in My lists	Works, but every write shows a Confirm card; typing "yes" doesn't count	Users who want a table and don't mind clicking
Excel workbook	Created and appended to via code; not tested further	Not recommended yet
Cowork	Needs access granted by your admin	Not used
Copilot Chat without a licence	No To Do and no work data. Creates Word files via code but can't edit them, and once reported "done" when nothing changed	No working 2nd Brain; paste your notes into each chat

Never trust "done" on its own. Twice in testing Copilot reported a write that hadn't happened. The ✅ line after a real write and a quick look in To Do or the document are the fix.

Version history

A living guide. The setup prompt at the top goes on the handout page, with this guide as the download.

Version	Date	What changed
5.3	20 Nov 2026, afternoon	Tested on a laptop: the Markdown guide works as an attachment; the 📌 line works; Copilot can't find 📌 lines in earlier chats, so you paste them in at Day start. The block in the Custom instructions is no longer indented.
5.2	20 Nov 2026, morning	Spotted means captured: an intention mentioned in passing becomes an #AGENT task straight away, or gets a 📌 line when Copilot can't write. Day start first sweeps earlier chats for 📌 lines and missed intentions.
5.1	20 Nov 2026	One setup prompt for all four assistants (Copilot, ChatGPT, Gemini, Claude): the section is now "For the assistant: how to run this setup", and if the guide isn't attached, the prompt has the assistant open the right one from wolfe.nl itself, or tell you which to download. "Kickoff prompt" renamed to setup prompt.
5	18 Nov 2026, evening	Tested end to end as a new user. A short setup prompt starts the setup; the guide stays the reference. Copilot now leads: it moves on after each step and suggests a default you can change. The six questions are asked word for word; open items are captured as #CAPTURED; the 2nd Brain block goes into the Custom instructions unshortened. Each chapter has an 'If it gets stuck' list, and an Addendum lists what doesn't work and the alternatives.
4	18 Nov 2026	Rebuilt on stores tested on a Mac the same day. One Word document, 2nd Brain Memory, plus the Custom instructions for who you are and what you know; private To Do tasks titled "#AGENT … #state @turn" for intentions, clarifications and lessons. Copilot Pages dropped, because Copilot can't write them. Archive renamed "Archive - 2nd Brain". Setup starts with a To Do check; the interview is six numbered questions; Copilot confirms every write with a ✅ line.
3	17 Nov 2026, late evening	Rewritten as a self-running setup guide: an intro for people, a step-by-step section for Copilot with the interview and the Custom instructions block. Intentions default to 2nd Brain Lists, To Do optional. CLARIFY is no longer a state: unclear items stay CAPTURED with a clarification. Weekly review checks what came back, not a priorities list.

2	17 Nov 2026, evening	After the Unilever session: three Copilot Pages instead of Copilot Memory; Day start, Capture and Weekly review; 2nd Brain Lists; one follow-up date; inner circle; no WAITING state; QR code removed
1	17 Nov 2026	First guide: Copilot Memory, Custom instructions and a To Do list called 2nd Brain; weekly health check