

Welcome to the AI that remembers you.

Sage is a personal AI that knows your work, your people, and what you're trying to get done — so you stop re-explaining yourself every time. This guide is your first week, laid out.

[Start your free trial](#)

14 days of Pro. No card required.




Fig. 1 — your world, as
Sage holds it

In this guide

A note before you start	01	Your first ten minutes	02	The six rooms of Sage	03
Seven modes, one Sage	04	Your memory is yours	05	Bringing your team	06
Plans, plainly	07	What Sage will never do	08	Your first week	09

Every other AI starts from zero.

You've used the brilliant ones. ChatGPT, Claude — remarkable minds that forget you the moment you close the tab. Every conversation begins with the same ritual: who you are, what you're building, who's involved, what happened last week. You are the context, and you carry it by hand, every single time.

Sage is built on one idea: the memory is the relationship. Tell it about your world once, and every conversation after that adds to a single private memory of your work. Ask a question on Tuesday and it already knows what you decided on Monday. Mention a name and it knows who that is.

That's why Sage compounds. Most tools are as useful on day 100 as they were on day one. Sage on day 100 has a hundred days of your decisions, your people, and your open loops behind every answer — and you've stopped explaining yourself entirely.

This guide walks you through the whole house: the rooms, the modes, the privacy promise, and what to do in your first week. It's a short read. The rest, Sage will remember for you.

— *The Sage team*

Ten minutes, once. Then it's Sage's job.

There's no import, no integration checklist, no manual to finish. Setup is a short conversation — and everything after it is memory doing the work.

01

Tell it about your world

A two-minute setup: who you are, what you're working on, who matters. Talk the way you'd brief a sharp new chief of staff — plain sentences are enough.

~2 minutes

02

Have one real conversation

Bring an actual problem — the decision you're circling, the email you're avoiding. Watch for the moment Sage uses something you told it earlier. That's the memory working.

From memory · 3

03

Come back tomorrow morning

Open *Today* and your morning is already sorted: what needs a decision, the follow-ups you'd have forgotten, one clear focus. This is the daily loop that makes Sage stick.

The payoff

Six rooms, one memory.

Everything in Sage draws on the same private memory. Each room is just a different way to put it to work.



Today / your morning briefing

Open Sage and the day is already laid out: what needs a decision, what moved while you were away, one clear focus. It reads across everything Sage knows — not just today's inbox.

TRY Just open it tomorrow morning. It's already written.



Chat / ask from anywhere — ⌘ J

One question, any screen, and it already knows you. No pasting context, no re-explaining who's who. It answers like someone who's been in the room all along.

TRY “What should I not drop this week?”



Read / documents, decoded

Drop in a contract, a deck, or a long thread. You get the gist, the risks, and the questions you should be asking — then interrogate it like you'd grill a sharp colleague. Answers stay grounded in the document.

TRY That contract you've been meaning to actually read.



Draft / writing in your voice

Sage learns how you actually write, then drafts emails, messages, and posts in your tone — not a robot's. You stop staring at a blank page and start editing something that's already most of the way there.

TRY “Draft the follow-up to this morning's meeting.”



Map / see what Sage knows

Your people, companies, projects, and goals, drawn as a living map of how they connect. It's proof Sage actually understands your world. Tap any node to ask about it.

TRY Open it after week one — watch it fill in.



Memory / the vault, open to you

Everything Sage remembers, browsable. When an answer draws on your history, you'll see a small From memory · 4 chip — tap it to see exactly what was recalled. Edit or delete any piece, anytime.

TRY Skim it after a few days. It reads like your work, organized.

One Sage, seven registers.

Sage isn't seven different assistants — it's one relationship that shifts register with the work, the way a great chief of staff does. Switch modes from the chat header, or press  J and just start; Sage follows your lead.



All Systems

Full context across your work. The default — start here.



Executive

Strategy and high-leverage decisions.



Investor

Capital allocation and portfolio.



Builder

Systems, architecture, technical work.



Mentor

Teaching, coaching, developing people.



Ministry

Faith, service, congregational care.



Personal

Reflection, life rhythms, presence.

Modes are configuration, never a paywall — every plan gets all seven. And whichever mode you're in, it's the same memory and the same calm, direct voice underneath.

Most AI resets every conversation. Sage gets more useful every day — the more you use it, the more it knows, and the less you have to explain.

WHY SAGE COMPOUNDS

Private by design, not by setting.

A memory this personal only works if it's unambiguously yours. So the privacy promise is structural — built into how Sage stores things, not a toggle you have to find.



Export everything, anytime

Your whole memory, downloadable whenever you want it. It's your record of your work — you can always take it with you.



Delete any piece — or all of it

Remove a single memory or your entire account. Deletion is a hard delete, all the way down — not a hidden flag.



Never sold, never trained on

Your data is never sold and never used to train models. And answers show their work — the *From memory* chip tells you what was recalled, every time.

A shared brain, on your terms.

On the Team plan, you invite people into a shared workspace — and Sage becomes the colleague who's read everything the team chose to share.

Here's the part that matters: **everything you tell Sage is private to you by default.** Nothing enters the shared workspace unless you explicitly share it — one memory at a time, or a whole conversation at once. Unshare anything and it's gone from the workspace immediately, no announcement made.

When Sage draws on a teammate's shared context, it says so plainly — *“per Alice's note”* — so you always know where an answer came from. Ask *“who on the team knows about the Meridian contract?”* and Sage answers with names and evidence, or an honest *“no one's shared anything on that.”*

- Invite by email from Settings; seats are included in your plan
- Private by default — sharing is always your explicit choice
- Share one memory or a whole conversation; unshare instantly
- Every shared fact carries its source, by name
- In group threads, Sage stays quiet unless you @mention it

Start free. Upgrade when Sage is running your day.

Every trial is 14 days of Pro, no card required. And one thing no plan will ever gate: **memory**. The relationship is structurally the same on every tier — higher plans buy sharper models, more channels, and more room, never “Sage remembers you more.”

	Solo	Pro	Team	Sovereign
	\$29 / month	\$149 / month	\$499 / month	\$2,000 / month
Made for	Getting your world into one place	Running your day on Sage	Sharing context across a small team	When Sage is mission-critical
Messages / month	500	5,000	20,000	Unlimited
Voice minutes / month	30	300	1,200	Unlimited
Channels	Web	Web, voice, and Telegram	Everything in Pro	Everything in Team
The rest	Persistent memory that compounds	Natural ElevenLabs voice · all coaching styles	Shared workspace with seats	Your own custom voice · priority support

Cancel anytime from billing settings — you keep access through the period you paid for. If you hit a limit, Sage tells you and you can upgrade in a click; it never silently bills you for overages.

Judged by what it refuses.

Most AI products are engineered to maximize your time in the app. Sage is engineered to give it back. These aren't missing features — they're standing refusals, reviewed and reaffirmed as product policy.

No ~~streaks~~

No “Day 47 with Sage!”, no streak anxiety. If you've been away two weeks, Sage mentions it only when it's actually relevant — like a colleague would.

No ~~nagging notifications~~

Sage never pings you to drive a return visit. Notifications fire only when you asked for one — “remind me Friday about Joshua” — and that's it.

No ~~productivity scores~~

You will never be “73% productive this week.” Sage reflects patterns back without judgment; it doesn't grade you.

No ~~memory paywall~~

Sage will never remember you less because you're on a cheaper plan. Memory is the relationship, and the relationship isn't for sale by the tier.

No ~~persona dress-up~~

One Sage, one consistent voice. No avatar picker, no character library — identity is memory and language, not pixels.

No ~~bot pointed at your customers~~

Sage works for you, never as a widget you aim at someone else. The relationship is the product, and it's yours.

A week that compounds.

You don't need to do all of this. But people who do rarely go back to explaining themselves to a blank chat box.

- | | |
|--|-------|
| ☐ Finish the two-minute setup — your world, in your words | DAY 1 |
| ☐ Bring Sage one real problem and talk it through | DAY 1 |
| ☐ Open <i>Today</i> with your coffee — your first briefing | DAY 2 |
| ☐ Drop a real document into <i>Read</i> and grill it | DAY 3 |
| ☐ Have Sage draft one email you'd normally put off | DAY 4 |
| ☐ Try a voice conversation, or connect Telegram | DAY 5 |
| ☐ Open the <i>Map</i> and see a week of your world, connected | DAY 7 |

Meet the AI that remembers you.

Fourteen days of Pro, no card required. Two minutes to set up. By next week, it knows your world.

[Start free](#)

SAGE — A PERPETUAL CORE STUDIO

15% OF ALL REVENUE SUPPORTS THE INSTITUTE FOR HUMAN ADVANCEMENT

OWNER'S GUIDE • VOL. 1