



PRIVATE ASSISTANTS . AI

Your own assistant, on your own server.

What it does, what it costs, and what
you own at the end.

Start with the part nobody needs convincing about.


ChatGPT and Claude are genuinely good, and most of what people want from an assistant, they already do.

Draft the email. Summarise the twelve-page contract. Translate the message before you send it. Pull the numbers out of a statement and tell you what they mean. Rewrite the listing so it sounds like a person wrote it. Talk a problem through at eleven at night when there's nobody to ask.

They are fast, they are cheap, and they get better every few months. We use them daily. If your work is mostly one task at a time, ending when the task ends, a subscription is the right answer and this document will not try to talk you out of it.

What follows is about the other kind of work. The kind that doesn't end — where the context that matters is spread across six months, four counterparties, and a folder of documents you half remember signing.

Four limits you feel daily and stopped noticing.



FOUR CONVERSATIONS. NOTHING BETWEEN THEM.

Every conversation is an island.

New chat per topic. The deal, the contract and the client who introduced you are three separate conversations that never meet. You are the connective tissue.

It only speaks when spoken to.

Nothing arrives before you ask. No morning brief. No reminder that the thing you'd chase on Tuesday is now nine days old.

It doesn't scale to real volume.

Six years of email. A drive of scanned contracts. A hundred hours of recorded calls. That is not an upload, it is a project — and the interface is built for a handful of files at a time.

You are working inside their app.

Their interface, their layout, their pace of change. Everything that matters lives behind a login you don't control, in a product that can be repriced, redesigned or retired without asking you.

That's where the app stops
being an assistant and goes
back to being a tool.

You've been feeding these companies for two years. *What did you keep?*

Every question, document and half-finished idea you have typed into one of these apps went somewhere. It is stored, indexed and used to answer you. You have never seen it.

You can request an export. What comes back is a dump in their format, compressed on their terms, shaped for their system and not for you. There is no version of it you can open, read, correct or take somewhere else.



THEIRS



YOURS

THE SAME TWO YEARS. ONLY ONE OF THEM READABLE.

We are not going to pretend our assistant never speaks to those companies. It does.

When you ask a question, that question goes out to the model and the answer comes back — the same way your phone talks to a mapping service without handing over your address book. The difference is what stays behind.

You end up with a complete, organised, permanent copy of your own working life. They never gave you one.

Every message, document, transcript and decision is written to your own machine as it happens, in plain readable files. Labelled and cross-referenced more carefully than any chat history you have ever scrolled — because the labelling is the job, not a side effect. Dated. Sourced. It can tell you where it knows something from.

It is stored in the Obsidian format, which is an open standard and not ours. You can open it today with free software, on your own laptop, with no account. And because it is plain text and yours, it moves to whatever model is best in three years without asking anyone's permission.

The point. Models will keep changing and getting cheaper. The archive is the part that compounds. Two years from now the scarce thing will be a clean, structured record of everything you have done — and the only way to have one is to have started building it.

One assistant. One memory. In the app you already answer.

It lives in your WhatsApp, and in Telegram if you prefer it. You message it the way you message a person.

Nothing new to open. Nothing new to learn.

Morning. Three viewings today, first at 10:15 in Marina. The Al-Reem contract came back overnight with two clauses changed — I've drafted a reply flagging both.

07:02

what did we agree with them in march about the payment split

07:04

60/40 on completion, agreed on the call of 12 March and confirmed in your email of the 14th. The draft they sent back has it at 50/50.

07:04

Illustrative. Your assistant's tone, language and level of detail are set by you during the build.

THE FOUR LIMITS, ANSWERED IN ORDER

- One conversation per topic, and you carry the context between them.
 - **One continuous conversation that never resets.** Mention a client from March in passing and it knows who you mean, what you agreed, and what is still open. Every topic sorted and labelled by default, and it makes the connections between them without being asked.
-

- It waits for you to open the app.
 - **It messages you. Your schedule in the morning.** The follow-up you said you'd send. A draft reply waiting on an email that came in overnight. A calendar change made and confirmed back to you.
-

- A handful of files at a time.
 - **Gigabytes.** Drop them on your server or connect the cloud account they already live in, and it works through the lot in the background.
-

- Their app, their servers, their terms.
- **A machine registered in your name.** We build it, hand you every key, and remove our own access in front of you.

The more you
give it, the less you
have to remember.



ONE AT A TIME, IN ORDER, NOTHING DROPPED.

**Bulk is the normal case, not the exception.
Point it at the archive and let it work.**

Years of email. Folders of contracts, invoices and scanned paperwork. Exported message history. Whatever is already sitting in your cloud storage. It reads through all of it in the background and files what it finds — so the first useful day is the day it comes online, not six months of patient training later.

Send it a document in the morning and ask about it in the afternoon, in the same breath as something else. It already knows.

One thing worth flagging, because it has turned out to be the most useful habit of all.

Pocket recorders have quietly become very good, and very cheap.

It has become standard practice on the builds we've done, and the reaction is always the same. You record the meeting or the viewing, send the file across, and that is the end of it. No notes during, no writing up after. The conversation is transcribed, filed against the right client and the right deal, and available by asking — this week or next year.

People stop taking notes entirely. It is often the single biggest change in how they work, and it costs about the price of a nice dinner.

Built once, for one person, and then it's yours.

Build, handover and first three months of support

AED 12,000

One-time

Running costs — server, voice, backup, model usage

Under AED 400 / month

Paid by you, direct to the providers

Support after the first three months

AED 1,800 / year

Optional

Built one at a time. Most builds are done in two or three days, and rarely more than a week — the pace is set mostly by how quickly you get us what we need. You are needed in person twice: once at the start for the questionnaire and the detail of what you want, and once at handover for the key change, the tutorial and a live test. Anything in between is optional and at your convenience.

Less than one hour of a decent assistant's time, for something that works every hour. Reading in your existing archive happens once, during the build, and is included.

WHAT'S INCLUDED

- A server registered in your name, hardened, with no public login page to attack
- WhatsApp and Telegram, voice notes in and out, in your language and one other
- Email triage and drafting, calendar read and write, both confirmed back to you before anything is sent
- Your existing archive read in and filed during the build, not left for you to do
- The vault: encrypted, backed up, Obsidian-compatible, and readable by you from day one
- Handover session where you generate your own keys, set a passphrase we never see, and watch our access fail
- An education in the thing itself — how these models actually behave, where they fail, and how to direct your own system afterwards without needing us

You will pick up a great deal along the way. Most people finish the build understanding more about this technology than they expected to, because we explain what we are doing while we do it. By the end you will know how to instruct your assistant directly, and how to have it write the instructions for changes to your own system — so you can extend it yourself instead of paying somebody every time you want something new.

A conversation, then a questionnaire, then a build.

The next step is thirty minutes, in person or on a call, working out whether this is worth doing for you specifically.

If it is, you fill in an intake questionnaire covering what you want it to know, what you want it to never touch, and how you actually work. That document drives the build. Nothing starts until you have signed off on the scope in it.

EMAIL

hello@privateassistants.ai

WHATSAPP

[+971 58 632 0655](https://wa.me/971586320655)

TELEGRAM

[@PrivateAssistantsAI](https://t.me/PrivateAssistantsAI)

INSTAGRAM

[@PrivateAssistants.ai](https://www.instagram.com/PrivateAssistants.ai)

BASED IN

Dubai, United Arab Emirates

One thing we will say plainly. If after that conversation a subscription is the better answer for you, we will tell you so. We build these one at a time, and we would rather build the right ones.