Mnemara

Own your memory.

Take control of what your AI knows about you.

Whitepaper · 7 October 2026

The more useful AI becomes, the more it needs to know about us.

How we like to work. What matters to us. The decision we made last week, and why. The things we have tried, the things we worry about, and the things we would rather not explain again.

That is what makes an AI assistant feel like an assistant, rather than a search box.

But there is a trade-off. Much of that knowledge stays inside the product that learned it. We build a relationship with one AI, then start again with another. Our history becomes part of what keeps us there.

And this is not just a list of favourite coffees. Over time, it can become a deeply personal picture of our lives: health, relationships, sexuality, finances, ambitions and beliefs. At work, it can include private plans, customer relationships and the thinking behind important decisions.

We should not have to choose between useful AI and control over that knowledge.

Mnemara is being built to separate long-term memory from the AI vendor. Keep the memory under your control. Let the AI tools you choose use the parts they need.

> AI should work from your memory, not own it.

Who this is for

We are starting with individuals who want AI to know them without tying that knowledge to one provider. People who value continuity, care about privacy, or simply want to try a different AI without starting from scratch.

The same need reaches beyond individuals. Teams want to protect their knowledge and intellectual property. Organisations and governments need control over sensitive information. Those are longer-term directions, not claims about what the first release can do.

1. Useful memory. Uneasy ownership.

Memory is a good thing. It saves repetition, helps an assistant understand what we mean, and gives past decisions a place in the next conversation.

The problem is not that AI remembers. It is where that memory lives, who decides how it is used, and whether we can take it with us.

Knowing you should not lock you in

Imagine changing AI tools after months of explaining your work, preferences and priorities. The new tool may be capable, but it does not know you. You spend time rebuilding context you already gave someone else.

The better the old tool knows you, the harder it becomes to leave. That is a quiet form of lock-in: not a contract, but the cost of starting over.

How much of yourself do you want one AI company to hold?

Your health concerns. Your relationships and sexual preferences. Your finances, political beliefs and private doubts. The things you ask an AI because you would hesitate to ask another person.

At work, the stakes include confidential strategy, customer information and intellectual property. For governments, they extend to sensitive citizen information and public accountability.

Whether you use OpenAI, Anthropic, Google, Meta (Facebook) or xAI, the question is the same: do you trust that company to protect and handle this much information, on terms you will remain comfortable with?

You may trust it today. But leadership, commercial priorities and policies can change. You should not have to predict a company’s future to decide what to tell an assistant now.

Trust should not have to be all-or-nothing. Keep useful memory under your control, and choose what each AI receives. That reduces dependence on any one provider; it does not remove the need to trust and verify the systems you use—including Mnemara.

Controls inside one product are only part of the answer

AI products may already let people inspect, correct or delete memories. Those controls matter. But managing memory inside one product is different from having an independent memory you can use across products.

This question extends beyond chatbots. As personal assistants begin organising our lives and acting on our behalf, the information they remember—and the permissions they hold—become more consequential.

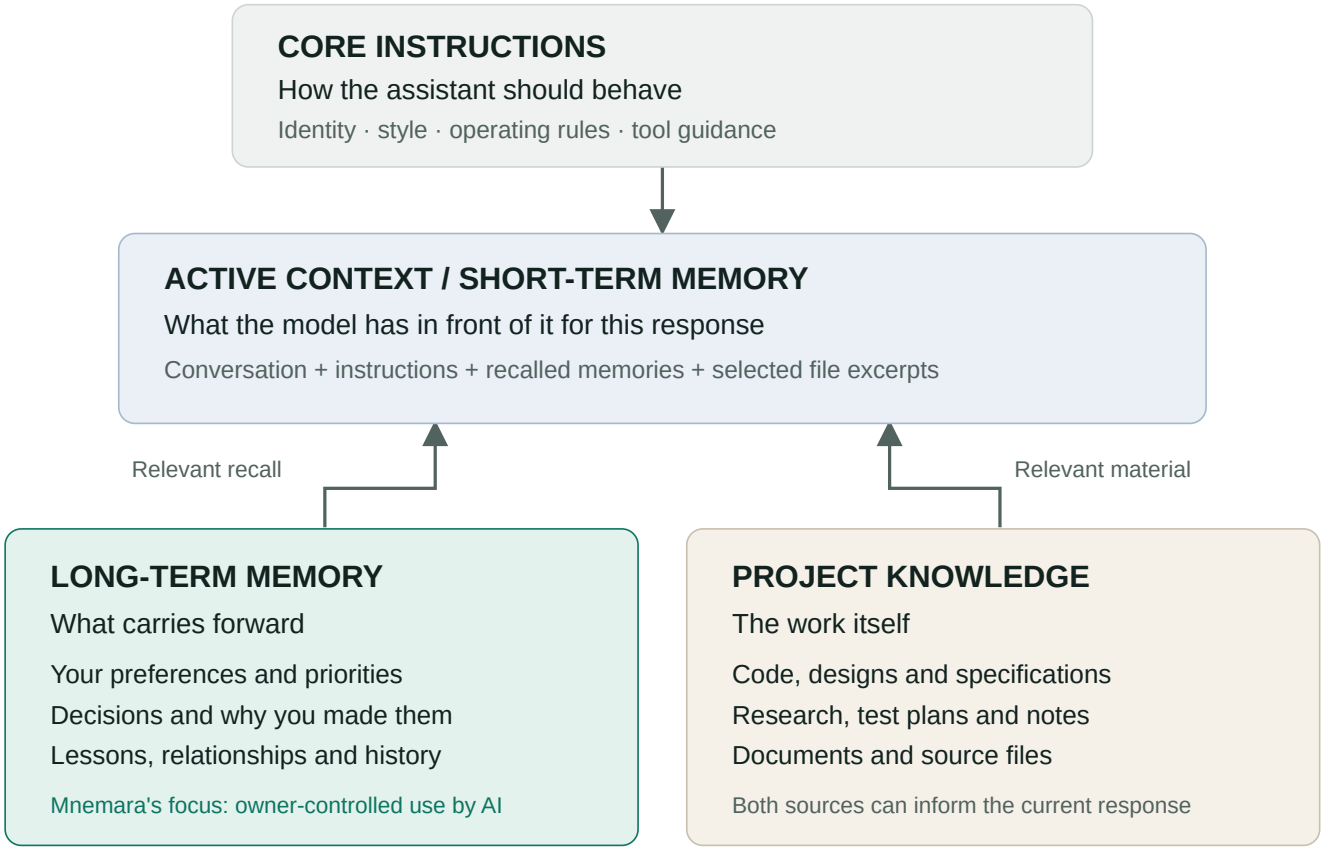
Mnemara’s aim is to make the owner—not the choice of AI vendor—the centre of that relationship.

> Change your AI. Keep what it has helped you learn about yourself and your work.

2. What do we mean by memory?

An AI model does not automatically carry every previous conversation into its next response. The app or agent around it—the **harness**—assembles the instructions and information it can use.

Four useful categories help explain that. These are practical distinctions, not rigid boxes: information can move between them or serve more than one purpose.



Your personal profile—your name, how you like to be addressed and how you prefer to communicate—should also be yours to carry between assistants.

Personal profile and remembered experiences are both parts of long-term memory. The diagram separates their purpose from the current conversation and the detailed project material.

The layers, in everyday terms

Imagine asking: **“Help me prepare for tomorrow’s investor meeting.”**

Layer	Description	Example
Core instructions	How the assistant should behave and operate	Ask before sending messages or taking actions.
Context / short-term memory	What is available for the current response	Tomorrow’s meeting time, attendees, duration and the format you just requested.
Personal profile — long-term memory	Who you are, what matters to you and how you like to work	Your name; direct business communication; a focus on climate outcomes; positioning your team at the forefront.
Experience and decisions — long-term memory	What happened, what you learned and what you decided	“The last meeting went poorly because the deck had too much text. Next time, lead with the demonstration.”
Project knowledge	The detailed material supporting the work	Current deck, value proposition and financial forecasts.

These are purposes, not rigid storage locations. A communication preference might currently sit in custom instructions. Meeting details might come from a calendar. A lesson might be written in an Obsidian note as well as retained in memory.

Having the deck is not the same as remembering what you learned from presenting it. Both matter. One helps the AI work on the presentation; the other helps it work better with you next time.

Where retrieval fits

Retrieval-augmented generation (RAG) means finding relevant information and bringing it into an AI’s answer. It can retrieve project documents, external sources or long-term memories.

RAG is a technique within the wider system—not the whole memory framework and not a type of storage. The app can arrange retrieval automatically, or the model can request it through tools. The selected information joins the active context. Mnemara can use this technique while also governing what is retained, updated and shared.

3. Put the owner in control

Mnemara is being built as an independent memory layer that supported AI tools can ask to use. The aim is not to remember everything. It is to keep the things worth carrying forward, with you deciding what happens to them.

Keep what matters

An assistant can suggest remembering a settled decision or preference. A passing thought should not quietly become a permanent fact about you. Sensitive information needs particular care, and some information—such as passwords and recovery codes—should not enter the memory store at all.

Share what is needed

Remembering something and sharing it are separate decisions. You may want a financial preference kept privately, but require approval before an AI can retrieve it. An ordinary preference may be suitable for automatic sharing under your chosen policy.

The AI can suggest how to classify a memory. Mnemara must enforce the rules; it cannot rely on the model to police itself.

Let people change

Preferences change. Decisions are reversed. A useful memory system needs correction, history and a way to retire outdated information—not an ever-growing pile of statements treated as equally true.

When two memories disagree, simply choosing the newest is not always enough. Sometimes the right answer is to ask you.

Choose where memory lives

The current prototype keeps an encrypted vault on the owner’s Mac. Local-first storage is the starting point. Planned cloud storage will offer an always-available option for people who cannot keep a local device online, or who want access across devices and supported AI assistants. Cloud improves availability and convenience; cross-AI portability does not depend on cloud storage.

Wherever storage lives, the goal remains the same: clear custody, usable controls and a way to take your memory with you.

Being honest about the limitations

Mnemara controls what leaves its own vault. It cannot control everything an AI provider does after receiving that information. We use the provider’s available memory and privacy settings, but turning off its memory feature does not necessarily stop it retaining chats or processing shared information. That is why choosing what to share matters.

Local storage does not make an external AI local. Pausing future access cannot pull back information already disclosed, and Mnemara does not erase past chats or information you share directly. Provider settings and retention terms still matter.

4. Where Mnemara fits

People already keep knowledge in notes, documents and AI products. Mnemara should work alongside useful tools, not ask everyone to abandon them.

The landscape at a glance

These categories overlap. This is a guide to their focus and trade-offs, not a security ranking or a claim that competitors lack owner controls.

Type	Role	Examples	Strength	Trade-off / question
Built-in AI memory	Personalise an AI product	ChatGPT; Claude	Convenient, integrated experience	Memory and controls are product-specific; continuity elsewhere needs separate support.
Personal assistants	Help organise work and life, and act through connected tools	Established and emerging assistant products	Useful context combined with actions	Check both information access and authority to act; permissions vary by product.
Memory infrastructure	Extract, maintain and retrieve memory for apps and agents	Mem0; Supermemory	Reusable memory services and integrations; deployment and governance options	Check the specific offering: custody, setup and end-user approval behaviour vary.
Cross-tool memory tools	Carry preferences and context between supported agents	Mem0's OpenMemory; Supermemory MCP	Continuity outside one model; meaningful overlap with Mnemara	Check supported clients, capture defaults and disclosure controls for your use case.
Personal knowledge tools	Keep and connect notes, decisions and project material	Obsidian	Local files, flexible organisation and portable formats	AI capture and permission workflows depend on the integrations you choose.
Owner-facing memory layer	Put personal memory and disclosure decisions in the owner's hands	Mnemara—prototype	Local-vault and approval workflow demonstrated with ChatGPT	Early stage; broader compatibility, security assurance and simpler setup remain work to do.

Most memory products start with a question: how can we help AI remember more? Mnemara starts with yours: what do I want it to remember, and who am I willing to share it with?

We are not claiming to invent the underlying technology. Our focus is bringing security, permissions, approval and portability together—so useful memory does not require handing over unrestricted access to your private life or business, or tying that memory to one provider.

We are building towards interoperability across devices, AI vendors, agent harnesses and personal assistants: **your memory follows you, while you decide what each tool may use.** Others offer controls too; we must earn trust through how well ours work together.

Vendor descriptions checked 6 October 2026: Mem0, OpenMemory, Supermemory, Obsidian. These are providers' descriptions, not independently verified security assessments.

5. How Mnemara is different

The point is not another place to put your files. It is a way to keep useful memory without giving every assistant unrestricted access to it.

One request, working together

You ask: **“Help me prepare for tomorrow’s investor meeting.”**

1. **The assistant follows its core instructions:** help prepare the meeting, but do not send anything without permission.
2. **The current conversation supplies the task:** a five-minute introduction for tomorrow’s meeting.
3. **Mnemara supplies permitted long-term context:** you prefer direct explanations; last time, the deck had too much text. Private context is withheld until the applicable approval is given.
4. **An authorised project connection finds the detail:** the current deck and forecasts in your selected Obsidian vault or files. That access has its own permissions; Mnemara does not automatically grant it.
5. **The assistant builds the briefing:** it combines those permitted sources into a shorter introduction and a demonstration-led agenda.

This is an illustrative workflow, not a claim that an Obsidian integration ships today. Obsidian can hold lifelong personal knowledge as well as project files. A future connection should complement it, not require copying an entire vault into Mnemara.

The difference is who stays in control

Mnemara is being built to manage what is remembered, what stays current and what each supported AI may receive. A useful answer is one outcome. Knowing what was kept and shared—and being able to change that—is equally important.

Cross-tool portability still needs supported integrations. It is not automatic just because memory is stored independently. Better retrieval may reduce repetition and token use, but those are benefits to measure, not savings we claim to have proved.

> Mnemara does not replace your assistant or your project files. It gives you control over the memory that connects your conversations, decisions and work.

6. A working start

Mnemara is a prototype, not a finished consumer product.

The current demonstration connects ChatGPT to a local encrypted vault on a Mac through a relay using the Model Context Protocol (MCP), a standard way for AI tools to connect to external services. With ChatGPT's native memory turned off, the demonstration has shown:

- recall of a saved preference in a fresh conversation;
- a sensitive memory held back until the owner approves its release;
- a proposed new memory or correction committed after local approval;
- native Mac notifications and an owner app showing requests, stored memories and activity.

That is a small but important proof: an AI can use memory held outside its own product, while a separate owner-controlled service governs access.

It is not yet proof of broad compatibility, production-grade security or effortless onboarding. The Mac must remain online. Setup involves several components. Retrieval needs to handle more natural wording reliably, and policy selectors and memory edit/delete buttons in the owner app are preview controls, not implemented actions. Approval-gated corrections through ChatGPT are a separate, demonstrated path.

Next: a limited alpha

The next step is a small, supported ChatGPT-and-Mac alpha: easier installation, better recall, dependable owner controls and clearer recovery when a connection fails. Phone-based approval is a proposed extension, not an available feature.

We intend a free local offering, with paid cloud, sync and broader integration options as the product develops. The packaging and deployment choices are not final.

Teams, enterprises and government/sovereign deployments would need additional work: shared ownership, administrative controls, key management, security assurance and operational support. For governments, this includes where citizen information resides and who has authority to access it. These are opportunities to explore, not capabilities implied by the current demo.

Looking further ahead: permission to act

Today the focus is permission to use your information. Over time, the same owner-control approach could extend to purchases: which assistant may spend, with which merchant, within what limit, and when it must ask first. Payment execution would remain with established payment providers. This is a possible future direction, not a current capability or a commitment to build a payment system.

The point is not to make AI remember everything

It is to make useful memory something we can understand, shape and carry with us.

We should be able to change our minds without being defined by an old conversation. Change our tools without losing our history. Get help from AI without handing it unrestricted access to our lives.

> Memory should travel with its owner.

Mnemara · Own your memory.
mnemara.ai