



PRODUCT BRIEF · LATENT SIGNAL

A private memory layer over your email.

Memento turns your whole email archive into living, source-attributed documents — and gets richer the more you use it.



01

The problem.

Most adults carry a decade or more of email — often 100,000 to 500,000 messages. Beneath the unread count sits a real record of a life: the projects you built, the people who mattered, the decisions you made, the topics you followed.

But email tools are built around what's urgent right now — the to-do, the message demanding immediate attention. They're good in the present and poor at the past. Email clients optimize for a chronological stream of messages, and search only helps when you already remember the right keyword — and even then it just hands you back a list of messages to reread. The long-term value, which lives *across* years of threads, people, and topics, stays buried.

Memento takes a different approach to finding things. Instead of matching keywords, it runs an *agentic search* that combines full-text search, vector search, and social-graph navigation — gathering and cross-referencing evidence across the archive before assembling a result.

The answers are far richer and more relevant to what you actually asked —
grounded in who and what is genuinely connected to your query, rather than which words happened to match.

02

What Memento does.

Memento reads your email archive and organizes it around the things worth remembering, not inbox chronology. It maintains living documents across four dimensions, plus a Home surface that ties them together.

● People

Relationship wikis for meaningful contacts: who they are to you, how the relationship evolved, and the details worth remembering. Like a CRM already populated from your email — including a social graph built from real interactions — that you can refine and update.

● Projects

Bounded narratives for work and life that normally scatter across dozens of threads and people over months or years — a renovation, a family trip, a product launch — reconstructed as a timeline of phases and key decisions. Ask it to extract attributes like total cost or an inventory of resources.

● Concepts

User-declared evergreen topics you follow, backed by real sources from the archive. Like projects, but without a timeline — security updates, or a technical area like front-end development. Memento keeps these concept wikis up to date as new email arrives.

● Newsletters

Coverage summaries, recurring themes, and recent items from the broadcast sources you subscribe to. Newsletter content is structurally different from personal correspondence — Memento separates it out and surfaces the items you actually care about.

A Home overview surfaces your priorities, active work, closest contacts, and things you might otherwise miss — dormant relationships or emerging projects.

Every factual claim traces back to the real email it came from. Follow any citation to the source message and open it in your email client — so you know the narrative is accurate, not invented.

03

Ask Memento.

Ask Memento is where you explore the archive rather than read pre-built pages. You ask open-ended questions — "*why did this project stall?*", "*who introduced me to that team?*" — and discover things you didn't know were there. It's powered by the same agentic search: Memento works like a researcher, running searches, reading the emails, resolving who's who, and assembling the evidence before it answers.

Conversations are saved as sessions you can return to and continue, and any session can become the seed of a new Project or Concept — so an exploratory question can turn directly into a durable, curated memory.

How it works.

Building a well-curated personal knowledge wiki from scratch is labor-intensive and technically daunting — out of reach for the average person. Because Memento works on your email archive, it populates that wiki with no manual effort, already resolving social graphs and canonical contacts across multiple work and personal accounts.

Deterministic extraction first

Before any model call, Memento resolves canonical contacts, maps email aliases, detects newsletter sources, and builds communication rollups directly from the archive. Agents start from a structured view of your history instead of guessing from raw search results — so the model spends its budget on synthesis, not rediscovery.

Multi-dimensional agentic search

Each dimension has its own agent workflow and prompt. A project agent looks for collaborators, decisions, timelines, and missing periods; a person agent looks for roles, relationship signals, and shared projects; a concept agent gathers evidence around an evergreen topic. The same archive answers different questions because each agent searches with the shape of the target memory in mind.

Purpose-built tools

Agents draw on thread summaries to condense long chains, gap detection to flag missing periods, compact message batching to assemble evidence efficiently, and context-status checks before expanding a search. The result is the *right* context delivered to the model — not a pile of loosely related emails. Agent runs are durable and stream their work live, so you can watch and audit exactly how a result was assembled.

Memory that compounds.

Narratives are not generated once and frozen. You can create, edit, and curate every document, and add notes with knowledge that isn't in the archive at all. Those edits and notes are treated as authoritative — agents never overwrite them with a guess.

More importantly, new narratives are built on top of existing ones. When Memento generates a person's page, it draws not only on raw emails but on the notes you've added and the corrections you've made elsewhere. A note on a contact can change how that person appears in related projects, concepts, and search. A project narrative you edited six months ago can shape the next update. And as new messages arrive, Memento keeps the narratives current.

*The effect is a memory surface that gets more personal and more accurate over time. Every note, correction, and curated project becomes part of the memory Memento builds **with** you — compounding rather than resetting.*

The personal graph.

Inboxes group messages by sender address, so the same person shows up many times across work, personal, and alias addresses. Memento resolves all of that into one person, automatically, using a deterministic graph algorithm over your email — no manual tagging. When it isn't certain two identities are the same, it asks rather than deciding for you.

From that graph it surfaces relationships you wouldn't think to search for: your closest collaborators, dormant relationships you've lost touch with, the groups who consistently appear together, and the *bridges* — the people who connect otherwise separate parts of your network. The graph is also what gives each narrative its relational context, so a person's page reflects not just their messages but their place among everyone else.

Local-first, on your terms.

Email is among the most sensitive data a person holds — private details about real people who never signed up to be processed by a cloud service. Memento runs locally by default, so that data needn't leave your machine. The app is a single binary that serves its interface and API on localhost only, treats the underlying archive as read-only, and writes solely to its own data.

It's also flexible. Run it on your own laptop or self-host it on a cloud provider of your choice, and point it at whichever LLM you prefer — any OpenAI-compatible provider works, including self-hosted local models. And because it's open source, how your data is handled can be inspected rather than trusted on faith. Memento is built on the open-source [msgvault](#) project, which handles mailbox acquisition, archive storage, search, and sync.

Why this isn't just RAG.

Memento uses retrieval techniques — vector search among them — but it is not a RAG chatbot. A RAG system retrieves a few matching snippets for a single prompt, answers, and throws the work away; the only thing it "remembers" is a raw index of text. Memento instead understands the structure of your archive first — people, aliases, threads, relationships — and runs multi-step agentic retrieval over that structure rather than a single similarity lookup.

The output is not a disposable answer but a durable, cited document that persists, absorbs your edits, and informs the next generation. **The valuable artifact is the living memory itself, not the one-time reply.**

09 — THE OPPORTUNITY

An emerging category, on a defensible footing.

Personal AI memory is an emerging category, but most entrants are cloud services that centralize sensitive user data. **Memento's strategic advantage is local-first, source-attributed memory over an asset users already own but cannot use — their email history — with a privacy posture and ownership model cloud-first products struggle to match.**

The compounding-memory dynamic doubles as retention: the product gets more valuable the longer it's used. The core product can be extended with closed-source connectors to other data sources — CRM, Slack, or a user's photo archive — and managed hosting opens Memento to a less technical audience.

Latent Signal

George Chiramattel · Ann Catherine Jose

georgeck@gmail.com · ann.jose@gmail.com

latentsignal.org/projects/memento
github.com/latentsignal-org/memento