

AIStor Memory

Organizational Memory Foundation for AI

Persistent agentic memory that captures knowledge, preserves state, and builds on every interaction across sessions and agents.

Overview

Every enterprise deploying AI agents is entering a new era, one in which agents increasingly help make decisions, create documents, draft analyses, and answer questions once handled by people alone. Knowledge generated by AI agents becomes organizational memory, and organizational memory belongs on a foundation that outlasts any single agent, model, or runtime.

AIStor Memory is a purpose-built organizational memory foundation for AI agents. It gives every agent access to a shared, long-term memory, including conversation history, files, artifacts, metadata, and secrets, all reached through standard operations. And existing tools work without changes, including commonly used agent sandboxes such as Daytona, E2B, Modal, GitHub Codespaces, OpenSandbox, and Vercel.

Organizational Memory is More Than Context

Context memory is the working data a model uses to respond to the request in front of it. It includes the prompt, the recent conversation, and whatever additional information can fit in the context window. The issue with context memory is that none of it lasts. Because it is rebuilt for each request and scoped to a single session, agents stay stateless: they re-learn the same context, lose preferences, and cannot compound experience. When the session ends, everything the agent learned is lost.

The common workaround is to continuously rebuild the past. Every new request means loading another transcript, attaching another summary, and sending the same history through the model again. Each step is a kind of amnesia tax, one with real costs to the organization, including:

- **Cost:** Rebuilding history wastes valuable tokens that could be better spent building new knowledge.
- **Time:** Larger inputs slow the current process as well as downstream actions that depend on it.
- **Quality:** Agents re-ask answered questions, reopen settled decisions, and repeat fixed mistakes. Replaying the past is not the same as remembering it.

AT A GLANCE

Give every AI agent a memory that lasts

- **Organizational memory, not just context:** Preserve what your organization knows, decides, and learns.
- **Agent Biography:** Start every new agent with the experience of every agent before it.
- **Open Memory:** Keep knowledge in documents you read, version, share, and move.
- **Skills:** Reuse what one agent masters across the entire fleet.
- **Infinite context:** Retain everything, with nothing truncated, summarized, or evicted.
- **Workspace and vault:** Protect work in progress and the secrets agents rely on.
- **Standard operations:** Deploy through the tools agents already use, with nothing to change.
- **Your infrastructure, your keys:** Own everything your agents learn, with nothing leaving your environment.

The more sophisticated workaround, summarize the history and retrieve fragments from a vector database, only hides the tax. Summaries are lossy, and the nuance, corrections, and reasoning behind a decision are exactly what gets cut. Embeddings are opaque, so nobody can read what the system remembers or explain why a fragment was retrieved. A full, structured biography kept in open documents loses nothing and hides nothing.

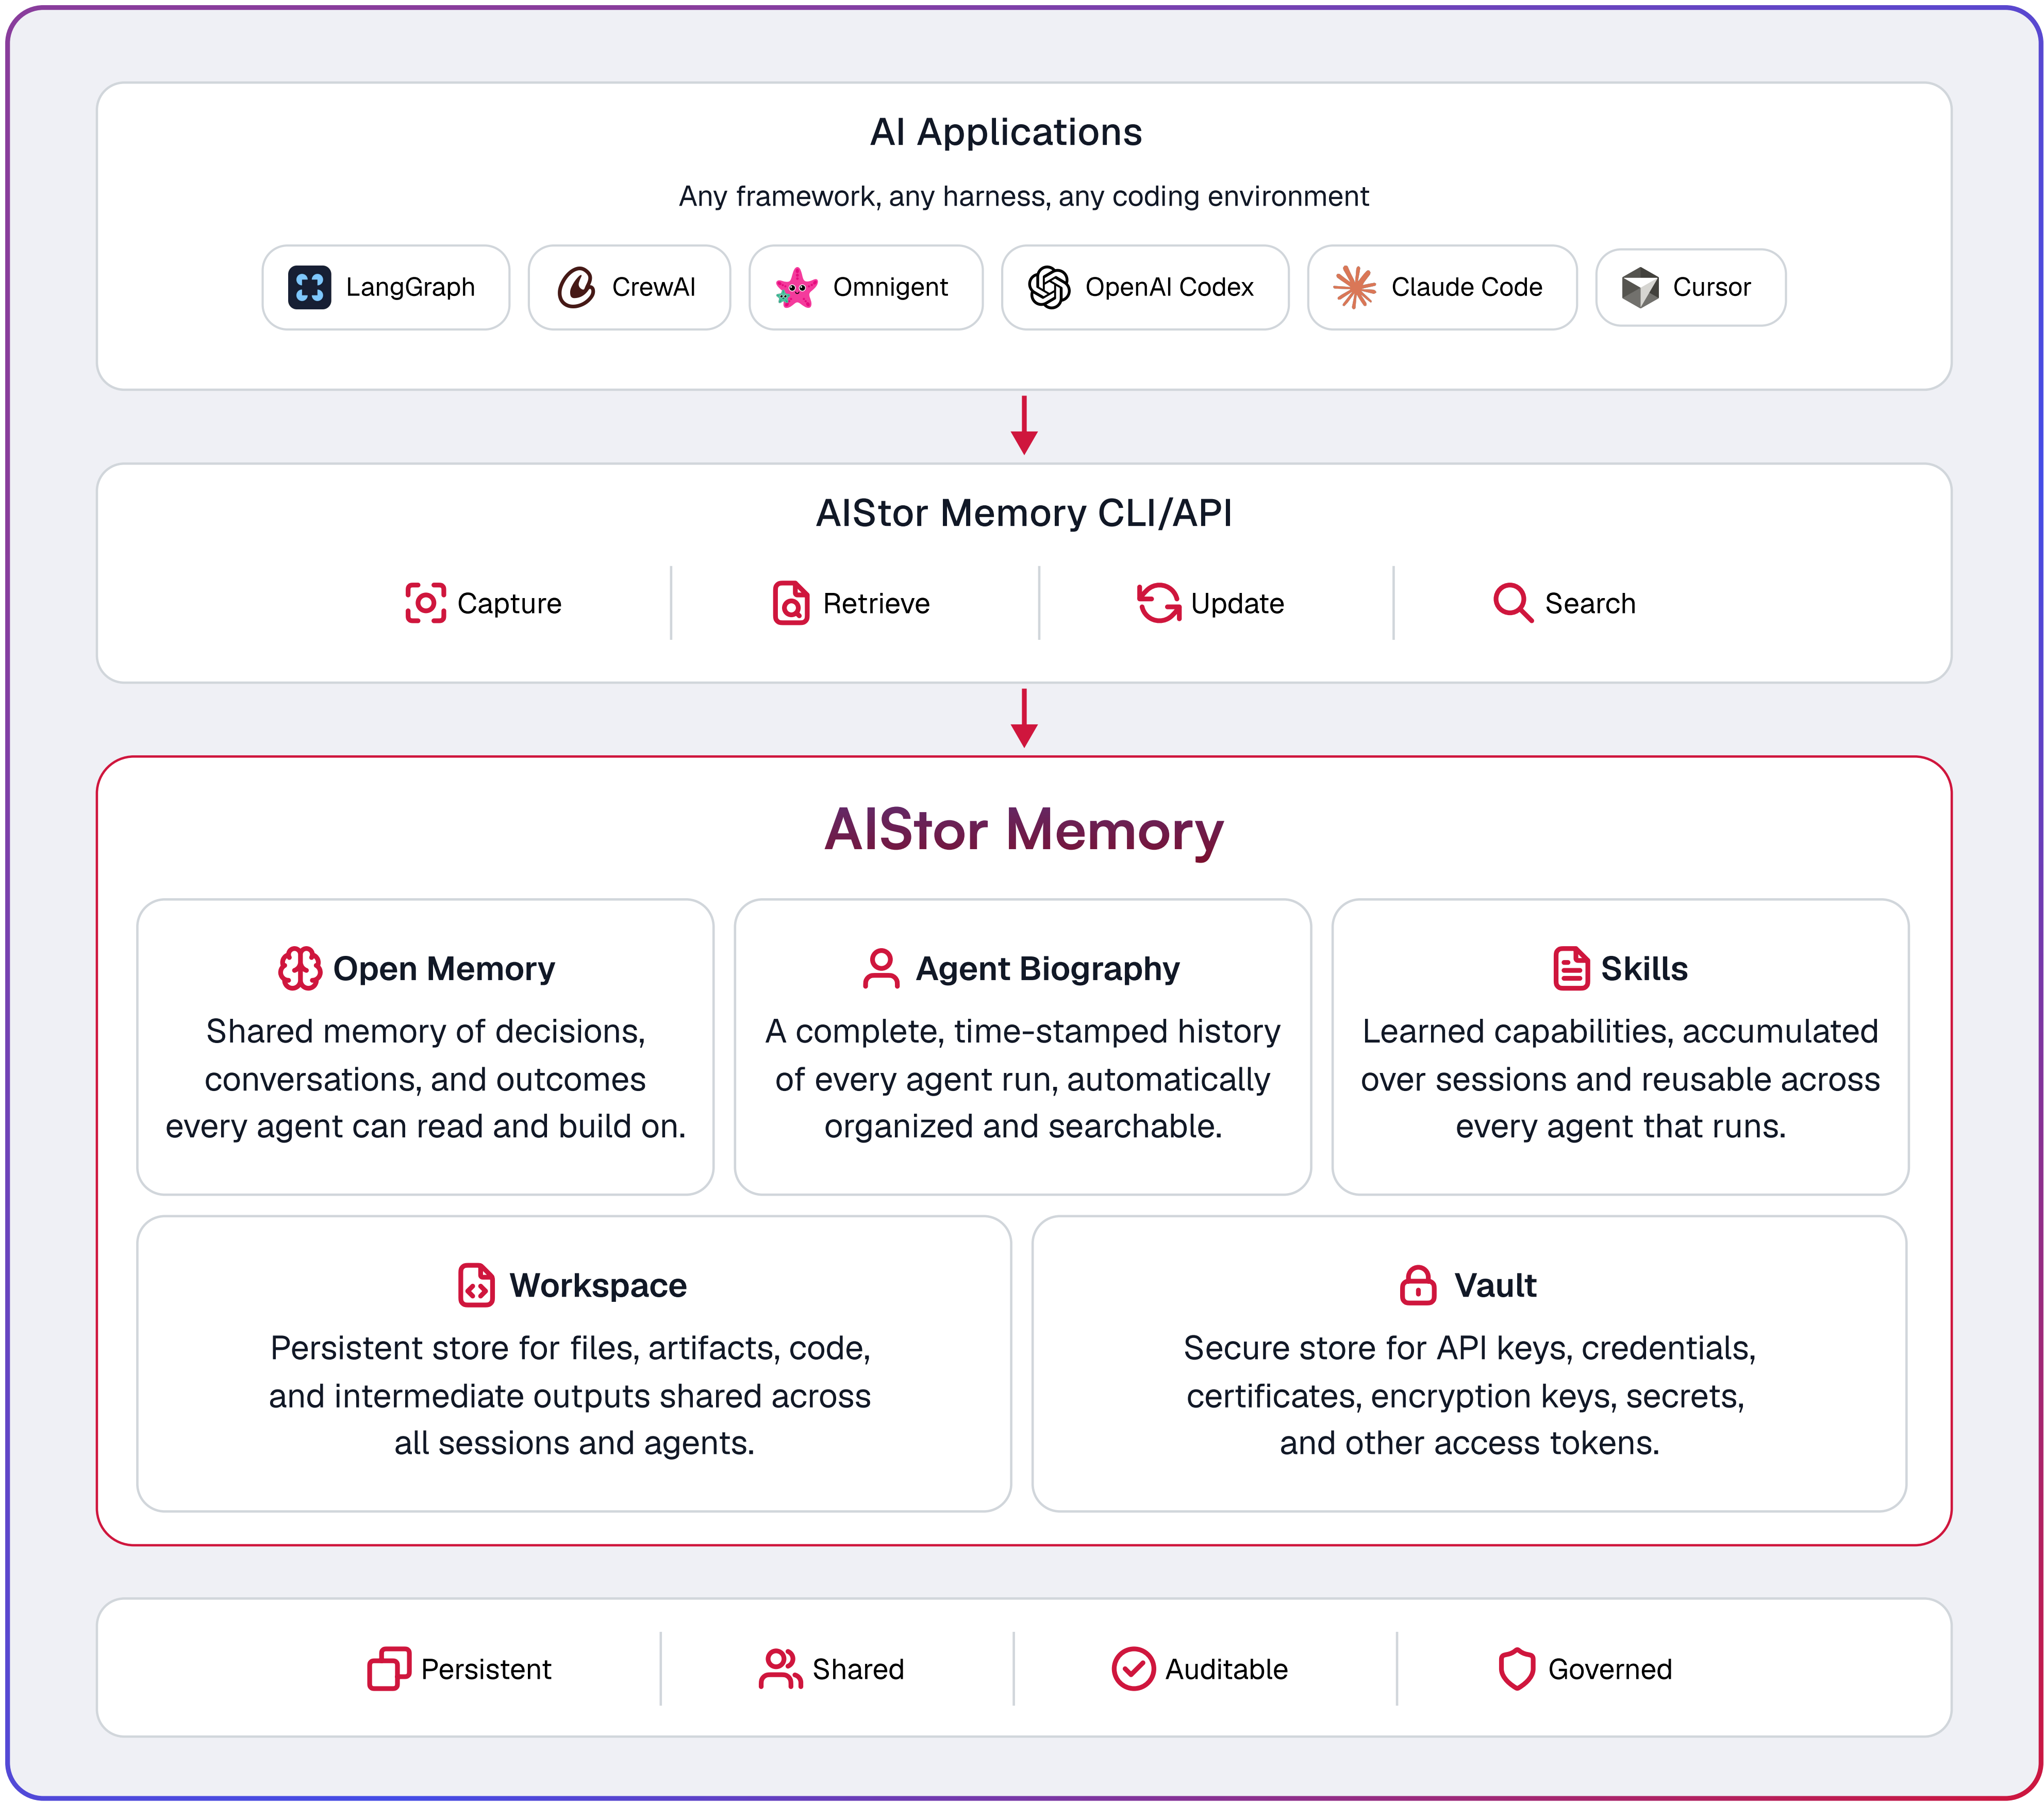
Organizational memory is different. It preserves what an organization knows, decides, and learns. Each interaction creates more than a response. It produces new memories that compound over time, with each conversation, decision, artifact, outcome, and reusable skill being used to improve future work. AIStor Memory keeps that evolving record as a durable enterprise asset, so knowledge, experience, and skills stay continuously available across users, agents, applications, and workflows.

Organizational memory is what turns isolated operations into continuous, context-aware agents.

AIStor Memory

A purpose-built memory foundation for enterprise AI.

AIStor Memory captures every interaction, organizes it into structured organizational memory, and retrieves the right knowledge at the right moment, so agents get smarter with every run. Information is captured as agents work, using the existing tools and frameworks teams already rely on. Memory is kept as open, structured documents that can be read by people as well as agents. Built on open formats, AIStor Memory provides a scalable, durable, and model-agnostic memory layer for enterprise AI. With AIStor Memory, what one agent learns becomes available to every authorized agent that comes after it.



Three key components sit at the center of AIStor Memory:

- **Agent Biography** is the comprehensive account of everything an agent does, captured automatically as it works. Every run adds to a structured, time-stamped biography that holds each prompt and response, each tool call and its result, and each artifact produced and question left open, in the order it all happened. The biography is fully searchable, and it is the record behind every memory an agent writes. New agents begin with the fleet's experience already written down instead of starting from zero.
- **Open Memory** is the facts, lessons, preferences, and standing decisions an organization's agents accumulate, transformed into shared memory every agent can read and build on. For example, a document might record that the finance team standardized on a single approval workflow for vendor payments after a duplicate invoice slipped through, who decided, when, and which agents it applies to. Everything is kept as open, structured documents, readable by people and agents alike, never trapped in a proprietary database or opaque vector store. Knowledge stays understandable, portable, and under your control.
- **Skills** are the learned capabilities, accumulated over sessions and reusable across agents. Skills capture how work gets done, the approaches and procedures agents acquire along the way, and make them available to any authorized agent, so what one agent masters, the fleet keeps.

In addition, AIStor Memory provides critical services to support agentic tools as well as sensitive data:


- **Workspace:** Persistent store for files, artifacts, code, and intermediate outputs shared across sessions and agents.
- **Vault:** Secure store for API keys, credentials, certificates, encryption keys, and access tokens.

Together, these components treat memory like the enterprise data it is. Every piece of memory is stored durably, protected against data loss and silent corruption. It is compressed, so retaining everything stays efficient, and encrypted under keys you hold, yet it remains fully searchable in that compressed, encrypted state. It is shareable, open in format, and agnostic to model and agent framework: portable, secure, and not locked to any single stack. Memory is not data copied from the data store into RAM.

How AIStor Memory Works

Agents interact with AIStor Memory through standard operations, and existing tools and frameworks work without changes. No additional wiring or scheduling is required and it runs automatically from the moment agents start.

While agents run, AIStor Memory does four things:

 Capture Every interaction is captured as the agent works. Prompts, conversation history, tool calls, attachments, artifacts, outcomes including performance telemetry (tokens, latency, cost) and unresolved work are all automatically preserved.	 Organize Memory is structured into an evolving, time-stamped, human-readable biography.	 Search Memory is retrievable through server-side search pushdown, enabling agents to efficiently recall relevant context and ensuring that only the required data is loaded into the context window.	 Evolve New interactions update memory, continuously improving future responses and workflows.
--	---	--	---

AIStor Memory drops into the agentic stack teams already use, including:

- **Sandboxes:** Daytona, E2B, and Vercel
- **General-purpose models:** Claude, GPT, Gemini, Llama, and others
- **Coding models:** Claude Code, OpenAI Codex, Cursor, and others
- **Orchestration frameworks:** LangGraph, CrewAI, and Omnigent
- **Connectivity:** MCP and other standard protocols

AIStor Memory makes every layer of the stack better. Sandboxes stay disposable, because their state survives them. Models stay interchangeable, because memory is not locked to any of them. Orchestrated agents share what they learn instead of starting from scratch. And the whole fleet stays observable, with controls and audit running across it.

Common Scenarios

AIStor Memory is valuable anytime agents perform long, multi-step work that has to survive interruption or involves critical enterprise data. Some common scenarios include:

- **Deep research and analysis.** A research agent gathering and synthesizing information across many sources over hours or days, producing notes, drafts, and a final report. Because it keeps its full working history rather than a truncated summary, it reasons over everything it has collected, and a long investigation resumes exactly where it left off.
- **Long-running and human-in-the-loop workflows.** An agent completing the automated steps of a process, pausing for a human decision or an external event that may be days away. Its complete state stays in AIStor while it waits, so it wakes and continues with no re-work and no lost progress. Over time, as the agent builds a memory of past approvals and rejections and the reasoning behind them, its memory compounds and it learns to produce better results.
- **Software engineering agents.** A coding agent working across a large codebase over a long session, reading files, making edits, and running tests as it iterates toward a goal. Its workspace lives in AIStor, so when the sandbox is recycled or the task pauses, the agent resumes from the same files with nothing lost, and a second agent opens the same workspace to review or continue the work.
- **Governed and auditable work.** An agent performing sensitive tasks, such as reviewing contracts, needs to ensure that every change it makes is audited and reviewable, by other agents as well as by legal and compliance teams.

Outcomes

Durable organizational memory changes what agents can do. They stop re-learning the same context, keep their preferences, and compound experience. Every agent builds a structured, time-stamped biography: compressed, encrypted, fully searchable, and usable by both people and agents. Retrieval returns the relevant context without pulling everything into the client or the context window. Memory built for individual agents compounds into memory for the organization, a record agents and humans build on together, and Vault and Workspace strengthen and operationalize it. Everything stays portable, secure, and not locked to any single stack: your teams switch agents, models, and runtimes without abandoning what the fleet has learned. Organizational memory is the foundational capability your agents have been missing, and AIStor Memory delivers it: continuity, compounding knowledge, and agents that actually remember.