MAJOR LABS · REPORT · JUNE 2026

The State of Agent Memory

An AI agent's memory today is a claim you have to take on faith. You can sometimes move it, sometimes see who inside the vendor's system touched it, and almost always delete it through the vendor's API. What you can never do is hand a memory to another agent, a counterparty, or an auditor and have them confirm, without trusting the store it came from, that it is authentic, unaltered, and authored by who it claims. We read six systems to check that. It held.

0 / 6 sign memory or make it verifiable	**0 / 6** full-fidelity export portability	**0 / 6** carry portable provenance
3 / 6 mature consent controls	**6** systems read firsthand	**6** dimensions assessed

Letta (MemGPT), Mem0, Zep, LangMem, model-native memory (ChatGPT, Claude, Gemini), and Cognee, each read against its own docs, source, and API references, then adversarially fact-checked. June 2026.

Finding 1 — Nobody signs memory, so it cannot be verified

There is no cryptographic signature, content-binding hash chain, Merkle structure, or append-only tamper-evident log on any memory record in any of the six. Integrity controls that exist protect against lost updates inside the vendor's own database, not against forgery. Letta's `version` column is optimistic locking. Mem0's `hash` is a deduplication content-hash, explicitly not a signature. Zep's "verified: true" is a user-set metadata tag. Model-native memory is encrypted at rest, which is confidentiality, not integrity.

> *In every case a compromised or dishonest store could rewrite a memory and no recipient could detect it.*

Finding 2 — Memory is exportable, but not portable-with-provenance

Several systems can get bytes out; none gets trustworthy, complete memory out. Letta's Agent File (.af) is the strongest case and still drops the archival store that holds the bulk of accumulated memory. Mem0's export is a schema-reshaping summary with no documented import. Zep Cloud has no bulk-export endpoint. Claude and Gemini portability is manual plaintext copy-paste, flagged experimental. In every path the provenance does not ride along, so what you receive is text without verifiable origin.

Finding 3 — Provenance and consent are real, but vendor-controlled

Where provenance exists, it lives in the operator's database and answers "which pipeline or source," not "which agent authored this fact," and cannot be checked without trusting the store. Consent is the market's one strength: the model-native trio let you view, edit, delete, and disable memory, and Mem0, Zep, and

Cognee ship GDPR-style cascade deletion. But every control is exercised through the vendor's own API against the vendor's own store. There is no user-held, cryptographically enforced control anywhere.

The capability matrix

System	Storage	Port.	Prov.	Consent	Verif.	Std.
Letta (MemGPT)	yes	partial	partial	partial	no	partial
Mem0	yes	partial	partial	yes	no	partial
Zep (Graphiti)	yes	partial	partial	yes	no	partial
LangMem	yes	no	partial	partial	no	no
Model-native	yes	partial	no	yes	no	no
Cognee	yes	partial	partial	yes	no	partial

yes (green) · partial (ink) · no (accent). Verifiability is a column of one word.

The gap

A portable, signed, consent-aware memory record would add three things no incumbent ships together: a vendor-neutral format that carries the whole memory losslessly; cryptographic provenance bound to the record, signed so it survives export and can be checked offline; and verifiable integrity a third party can validate without trusting the originating store. Why has no incumbent shipped it? Portability and third-party verifiability both reduce lock-in and remove the vendor as the necessary trust anchor, which is precisely the moat. The one capability that would let memory function as a credential rather than a convenience is shipped by exactly zero of the six.

Method: six systems spanning open-source frameworks, memory-as-a-service, and model-native memory, assessed against six dimensions from primary sources (docs, source code, API references). Each system's findings were handed to an independent reviewer told to refute the high-stakes claims. Read-only; no private benchmark. The disclosure is the method and the citations.

Major Labs builds open-source primitives and measurement for the agentic web. Cite: Major Labs (2026). The State of Agent Memory. majorlabs.co/reports/state-of-agent-memory.