

A customer asks to be forgotten. Has the agent forgotten them?



Dr Peter McCann Strain

CTO, DPhil/PhD in AI from Oxford University

Swipe >>

— THE FRAMEWORK GAP

A forget request does not land on one object. It lands on a system of derivatives (copies and summaries the system made from the original fact, for example a customer-summary cached in a sales agent's profile, an embedding in a vector store, a row in a backup). Durable agent memory is stored future authority that has to be governed.

— THE REFRAME

Governed lifecycle, not database.

THE OLD QUESTION

Can the agent recall this later?



THE QUESTION THAT HOLDS UP

Can the operator govern the fact after the system has made it useful?

— WHY THE FRAME HOLDS

High-90s vs 80-87% F1

Schema-grounded memory, validated records written against a known template, reported roughly high-nineties F1, against unstructured-recall baselines of 80 to 87 percent. The point is the design choice, not the decimal: schema first, recall second, because you can govern a contract but not a hope.

SOURCE

xmemory (Petrov et al., 2026 preprint) end-to-end F1 result, one team's benchmark not yet independently reproduced; peer-reviewed support from Modarressi et al. 'MemLLM' (TMLR 2024) on explicit read-write memory and Chen et al. 'AgentPoison' (NeurIPS 2024) on agent memory as a security surface; ICO, CNIL and EDPB erasure guidance; Mnemonic Sovereignty (Lin et al., 2026 preprint) for the lifecycle framing; Letta, Mem0, Zep and Cognee vendor documentation for memory-system mechanics.

— HOW IT WORKS

Governed forgetting is a set of receipts, not a mood.

- 01 Govern the **six-phase lifecycle**: write, store, retrieve, execute, share, forget.
- 02 Track every **derivative**: summaries, embeddings, graph edges, backups, recipients.
- 03 Demand **proof**: deletion events, mutation logs, reindex receipts, recipient acknowledgements.

— THE ARTIFACT

Memory is a governed lifecycle.

— FIGURE E13.1 · E13 - WHAT AN AGENT CANNOT FORGET

The Memory Lifecycle

Memory is a governed lifecycle, not a database. Six phases, each a separate lock on the same door, and forgetting loops back to the start.



• CONCEPTUAL MODEL

ARCHITECTING THE AI COWORKER

Memory is a governed lifecycle, not a database. Six phases, and forgetting loops back to the start.

— APPLY THE INSTRUMENT

Pick one fact your agent remembers about a real customer, employee, patient, student or user. Fill the forgetting table across every surface: where it lives, who may retrieve it, where it travelled, and the proof artefact. Name the blank cell.

— WHERE TEAMS MISREAD IT

Wiring the forget request to one delete API. Primary records disappear while summaries, embeddings, graph edges and backups quietly keep the fact alive. Inventory derivatives, propagate the delete, and collect a receipt at every surface.

— READ THE FULL FRAMEWORK

What an Agent Remembers, and Cannot Forget

Read the full essay – the argument, the sources, the figures and a reader-ready working artifact.

Substack petermccannstrain.substack.com · Medium [@peter.mccann.strain](https://medium.com/@peter.mccann.strain) · LinkedIn [peter-strain-dphil-15a607128](https://www.linkedin.com/in/peter-strain-dphil-15a607128)

New essays twice weekly, 2 June – 21 July 2026.

Next: [E14 – Agents Don't Know When To Stop](#)

— THE STACK SO FAR

E13 · Essay 13 of 22 complete · Arc III: Runtime control

YOU JUST ADDED

The memory lifecycle

STACK LAYER LIT UP

Memory

YOU CAN NOW ASK

map memory as stored future authority.

NEXT

E14 asks when an agent should stop.



Dr Peter McCann Strain

CTO, DPhil/PhD in AI from Oxford University

I build production AI systems and write about making agentic AI useful, inspectable, governable and safe enough for real work.

Follow on Substack for the full 22-essay series
petermccannstrain.substack.com