

How do we retrieve private memory without handing over a life?

Abstract. A useful private agent needs durable context, but sending an entire personal model to a search backend defeats the premise. The problem is selective recall: enough relevant memory to help, no ambient plaintext collection, and a clear erase path.

THE PROBLEM

Can a local vault select a minimal, query-relevant memory card, decrypt it only for the owner’s current session, and attach explainable provenance without exporting a private corpus?

HARD CONSTRAINTS

- 01 Vault key and owner token stay in memory only; page refresh re-locks the vault.
- 02 Identity changes, expiry, and revocation must clear decrypted context immediately.
- 03 Retrieval quality cannot be purchased with a whole-life context dump.
- 04 The server must not become the canonical plaintext personal-memory store.

A FIRST CONTRIBUTION

Prototype a local relevance gate with a fixed context budget, encrypted round-trip tests, and adversarial cases for stale, over-broad, or cross-account memory.

BUILD IN THE OPEN

github.com/hushh-labs/hushh-research
hushh.ai/discord

