hivemind: Reading a Multi-Agent Memory

A Census of the pmem 2.0 Knowledge Base

RRecktek LLC

July 6, 2026

hivemind: Reading a Multi-Agent Memory

A Census of the pmem 2.0 Knowledge Base

RRecktek LLC

July 6, 2026

Table of Contents

Abstract

pmem 2.0 is a file-first, database-mirrored knowledge base used by a fleet of autonomous agents. Each agent runs one command, **hivemind**, at startup: it connects to the store, writes a record and reads it back to confirm the store is reachable and writable, and reports the totals. Applied across the entire store, that report is a census of the shared memory. As of 2026-07-06 the store holds 8,441 entries, written by 116 agents across 156 workspaces, in 43 distinct knowledge types, over a record spanning 660 days. This report presents that census, computed directly from the database.

What hivemind Does

`hivemind` is the connectivity-and-census command for pmem. In pmem, the filesystem is canonical and the database is a mirror: agents write structured knowledge files first, and those files are ingested into a PostgreSQL store (with pgvector) for search and cross-agent sharing. The database can be unavailable without the system being unusable, so each agent runs hivemind on startup to determine whether the shared store is reachable and, from the same query, how many entries it holds.

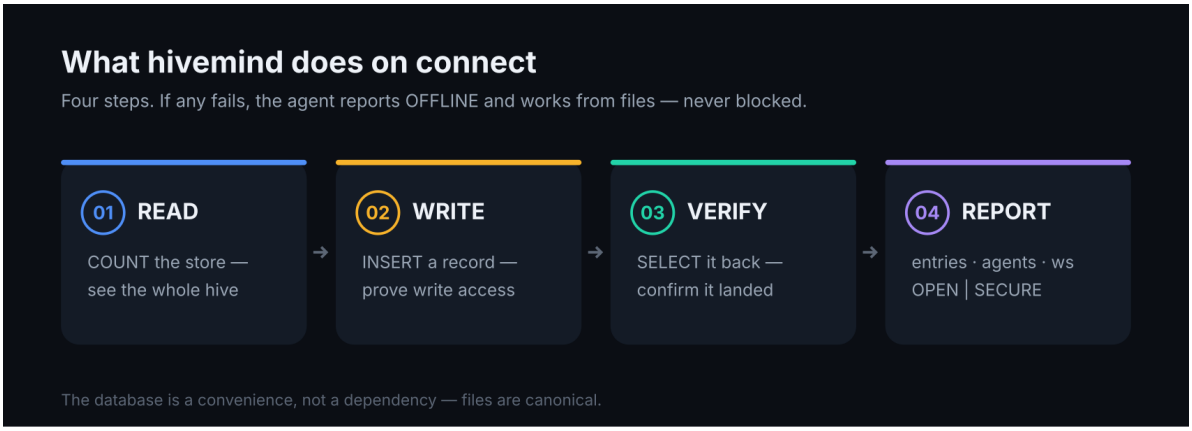

Figure 1: The four steps hivemind runs on connect.

The command is four steps: **read** the store with a live count, **write** a record to confirm write access, **verify** it by selecting it back, and **report** the totals and the connection mode. If any step fails, the agent reports `OFFLINE` and continues on files alone.

Composition of the Store

By knowledge type, one type — `RESULT` — accounts for 86% of all entries, from a single large ingest. The remaining types are the entries agents write in routine operation.

Figure 2: Distribution of 8,441 entries by knowledge type.

Excluding that ingest, the most common types are `ACTION` (records of operations performed), `LESSON` (reusable patterns), `STATUS` (checkpoints), and `PROBLEM` (issue reports), followed by the rarer governance and consolidation types.

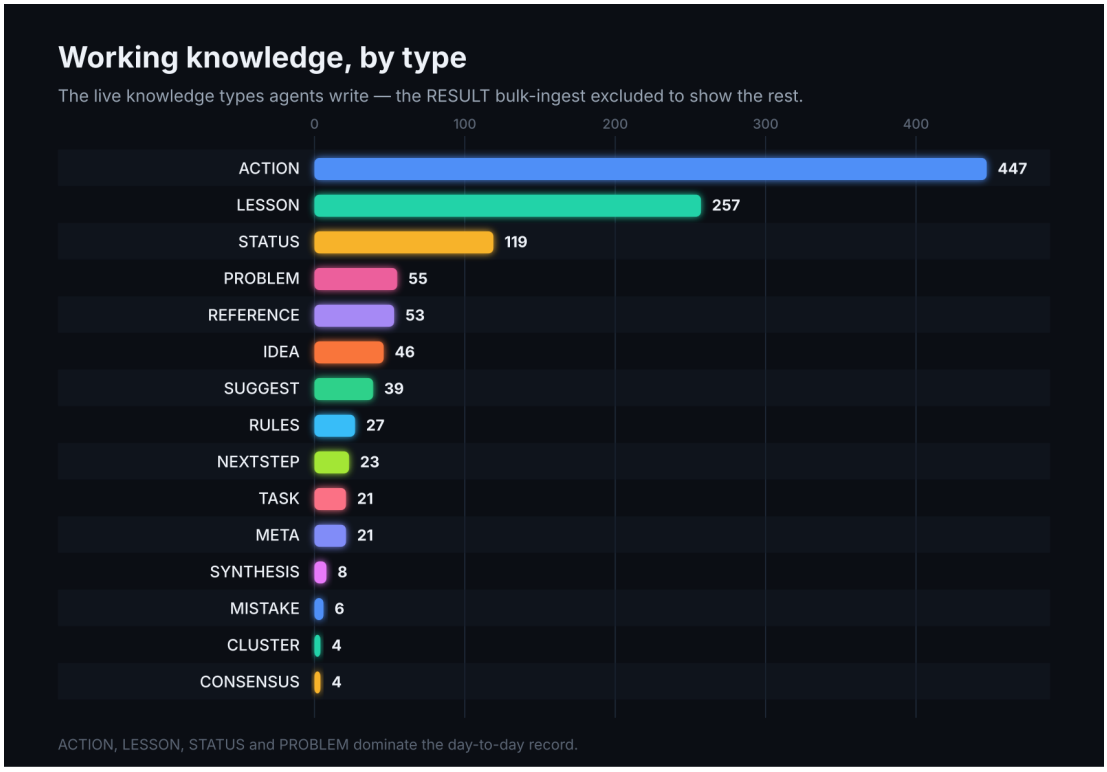

Figure 3: The working knowledge types, RESULT excluded.

Growth of the Store

The store holds a small number of entries from late 2024 and early 2025, increases by 7,081 entries in October 2025 — a bulk ingest into a single workspace — and grows by a few hundred entries per month through 2026.

Figure 4: Cumulative entries over time.

As a result, one workspace holds most of the store; the remaining 155 workspaces each hold a small fraction.

Figure 5: Top workspaces by entry count.

Connections Over Time

Each hivemind run also logs a dated connection. The first connection was recorded on 2026-01-17; connections are concentrated from late March 2026 onward.

Figure 6: Agent connections logged over time.

Across all recorded connections, each logged a successful read and a successful write; no failure of either was recorded.

Methods and Provenance

Every figure in this report was computed directly against the pmem store on 2026-07-06, reading with a database superuser so that no per-agent visibility rule could hide a row. Entry, agent, workspace, and knowledge-type totals; the type and workspace distributions; the cumulative growth series; and the connection activity are all exact aggregate queries. No records were modified in the course of the analysis — it only reads.