4.2 HiveMind — Shared Knowledge

HiveMind — Shared Knowledge Base

HiveMind (`~/system/agents/hivemind/hivemind.js`) is the **collective memory** of all agents.

What Goes Into HiveMind

- **Task completions** — every pi-orchestrator task result (Tier 2+)
- **Agent discoveries** — facts, patterns, warnings found during work
- **Session summaries** — condensed logs of work sessions
- **Manual knowledge** — docs, runbooks, architecture decisions

How It Works

```
Agent completes task
       |
       ▼
HiveMind stores entry:
  - Text (up to 5000 chars)
  - Source (which agent/engine)
  - Type (knowledge, alert, briefing, update)
  - Timestamp
       |
       ▼
BGE-M3 embeds the entry → Qdrant vector DB
       |
       ▼
```

```
Future agents query HiveMind:
  "What do we know about Drop's payment system?"
       |
       ▼
Semantic search → relevant entries returned as context
```

Usage

```
# Query knowledge
node hivemind.js query "How does Drop handle PSD2 consent?"

# Store new knowledge
node hivemind.js post my-agent knowledge "Drop uses SCA for all transactions over 30 EUR"

# Search by tag
node hivemind.js search --type alert --since 24h
```

Why It Matters

Without HiveMind, every agent starts from zero. With HiveMind:

- Agents learn from previous work
- Duplicate investigation is avoided
- Institutional knowledge persists across sessions
- Quality improves over time (flywheel effect)

Revision #3
Created 2026-03-10 23:09:21 UTC by John
Updated 2026-05-31 20:05:08 UTC by John