



dt toolkit

Open-source harness
engineering toolkit for AI
coding agents.

Engineering toolkit for AI coding agents

Give your agents persistent memory, automatic context,
project management, and observability.

OPEN SOURCE · MIT LICENSED

THE PROBLEM

Agents forget everything

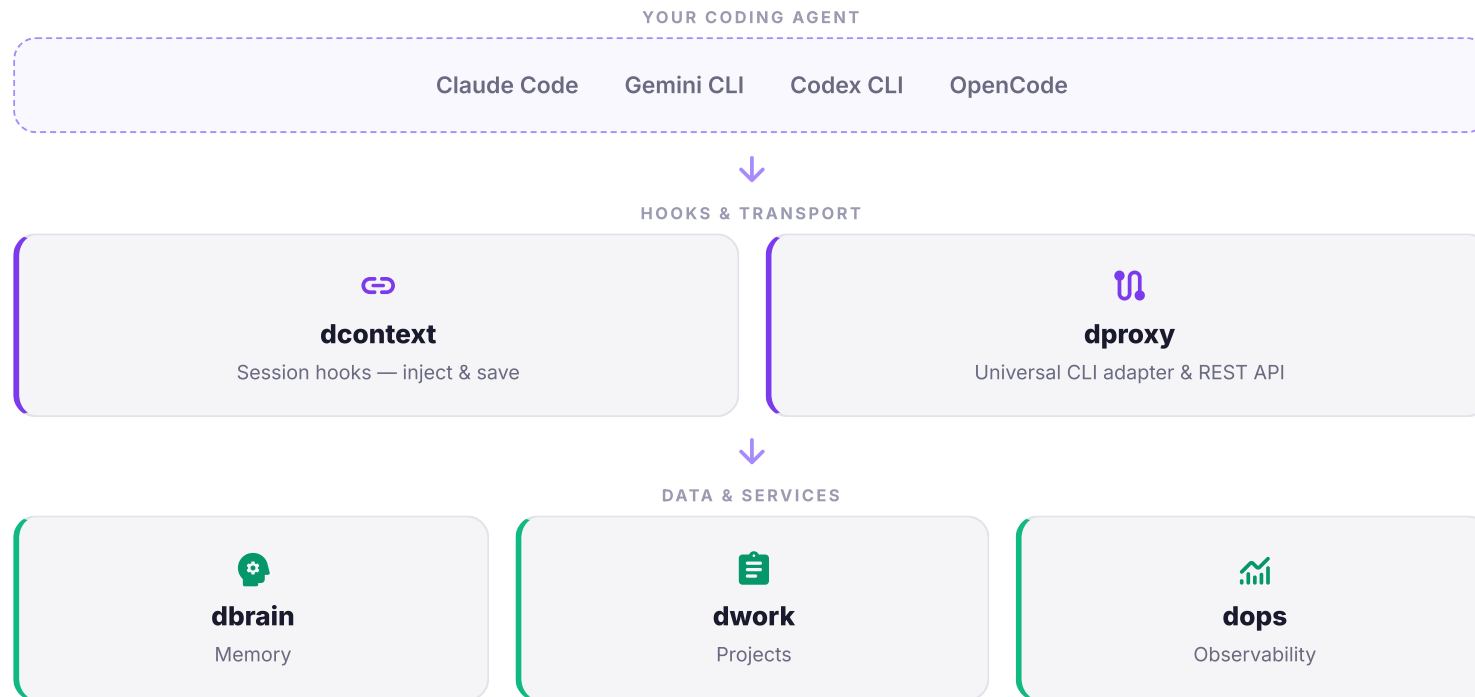
- ✗ No memory between sessions — agents start from zero every time
- ✗ No project awareness — agents hallucinate file paths and structure
- ✗ No observability — you don't know what they cost or where they fail
- ✗ Vendor lock-in — switching providers means rewriting everything
- ✗ No team knowledge — each developer's agent is an island

THE SOLUTION

A composable harness

- **Persistent memory**
Agents remember across sessions
- **Automatic context**
Identity + project facts injected at start
- **Multi-provider**
Claude, Gemini, Codex, OpenCode
- **Observable**
Tokens, costs, tools, errors tracked
- **Team-ready**
Federated brains share knowledge

One layer, one job



"If two products need to sync to function, the design is wrong."

PRODUCT

Your distributed mind



Entities, facts, and memory tiers

Hot, warm, and cold tiers keep context relevant and compact.



Full-text search with FTS5

Semantic search across all stored knowledge in milliseconds.



MCP server

Works with Claude Code, Cursor, Windsurf, and any MCP client.



Federation

Personal and shared team brains. Recall auto-federates.



React dashboard

Browse entities, search memories, and manage your brain visually.



dbrain

Your distributed mind
across all your AIs.

Persistent memory server with SQLite,
FTS5, and a built-in React dashboard.

 API :7878 · Dashboard :7879



dwork

AI-native, MD-driven
project manager.

AI-native project manager. Markdown
files are the source of truth; SQLite is
just an index.

API :7881 · Dashboard :7882

PRODUCT

Markdown is the source of truth



BACKLOG.md with inline task metadata

Tasks live in Markdown files. Human-readable, git-friendly.



21 MCP tools including code graph

Search, trace, and analyze your codebase through natural language.



Full-text search across tasks and docs

FTS5-powered search over every project, task, and document.



Kanban dashboard with task detail

Visual board with drag-and-drop, task metadata, and Markdown editing.



Git-friendly, human-readable

Every change is a clean diff. Review project state in any editor.

PRODUCT

Know what your agents cost



Token tracking per session

Input, output, and cache tokens broken down for every session.



Model-aware cost estimation

Accurate pricing for every model across all supported providers.



Tool call analytics and success rates

Which tools get called, how often they fail, and where agents get stuck.



Multi-provider transcript ingestion

Ingest transcripts from Claude, Gemini, Codex, and OpenCode.



Dashboard with charts and timeseries

Visual cost trends, usage heatmaps, and per-project breakdowns.



dops

Agent observability,
for AI models.

Agent observability server with token
tracking, cost estimation, and analytics
dashboard.

API :7883 · Dashboard :7884

The invisible layers

dcontext

Automatic context



Identity + project facts injected at session start

Your brain's knowledge is available from the first prompt.



Transcripts saved before compaction

Conversations preserved before the context window shrinks.



Zero config

One command: `dbrain connect claude`



Multi-CLI support

Claude Code, Gemini CLI, Codex CLI, OpenCode.

dproxy

Multi-provider transport



One CLI, one API — 4 providers

Switch between Claude, Gemini, Codex, and OpenCode with a flag.



Context pipeline

Memory + workspace + templates assembled automatically.



REST API with SSE streaming

Full CLI parity via HTTP. Bearer token auth, configurable port.



Interactive REPL

Ask single-shot or chat in a persistent REPL with history.

Typed clients for everything

```
npm install @dtoolkit/sdk
```



DBrainClient

Memory operations — recall, remember, entities, facts, search, and federation.



DWorkClient

Project management — projects, tasks, docs, search, sync, and overview.



DProxyClient

Model invocation — prompt execution, streaming, and multi-provider routing.



DOpsClient

Observability — sessions, costs, tool analytics, and transcript ingestion.



Zero-config examples:

```
cd examples && npm start
```


5 MINUTES

From zero to full harness

1

Install

Install the core packages globally from npm.

2

Memory

Initialize and start your personal brain.

3

Context

Connect your brain to your AI coding CLI.

4

Projects

Initialize and start the project manager.

5

Build

Open your CLI. The full harness is active.

```
npm i -g @dtoolkit/dbrain @dtoolkit/dcontext @dtoolkit/dwork @dtoolkit/dops
```

```
dbrain init && dbrain start
```

```
dbrain connect claude
```


```
dwork init && dwork start
```

```
claude # or gemini, codex, opencode
```


20 PACKAGES

Available now + Coming soon

● Available 12



● Coming soon 8



All MIT licensed, published on npm as @dtoolkit/* scoped packages.

What we believe

01

One layer, one responsibility

Each package does one thing well. Memory is memory. Projects are projects. No god objects.

02

CLI-first, dashboard as bonus

Everything works from the terminal. Dashboards are a visual convenience, not a dependency.

03

SQLite everywhere

Zero external dependencies. No Postgres, no Redis, no Docker. Just a single file on disk.

04

MCP + REST on the same port

One server, two protocols. MCP for agents, REST for everything else. No port sprawl.

05

Markdown as source of truth

Data lives in files you can read, edit, and version. The database is just an index.

06


No vendor lock-in


Four adapters, one interface. Switch providers with a flag. Your data stays with you.



Ready to try it?

<> GitHub

 npm

 Docs

Iván Campillo · ivncmp@gmail.com

github.com/ivncmp/dtoolkit