

Inside the AI harness.

How your agents are instructed. What they see. What they keep.




There is no single platform-wide “brain.”

The main assistant, Knowledge and Command Center compose different context, run different prompts and retain different kinds of memory. This atlas follows the actual code boundaries.

The harness landscape

Source map · not a live trace



CONTEXT

What enters a model call?

Follow selection, bounds, pruning, compaction and retrieval.

[Explore ↗](#)

MEMORY

What survives a conversation?

Compare ownership, stores and write/read lifecycles.

[Explore ↗](#)

PROMPTS

Show me the actual instructions.

Search exact templates, builders, skills and control code.

[Explore ↗](#)

New: repository-wide audit — full file census, syntax discovery and source-reviewed findings ↗

Latest fetched baselines: web `origin/dev-sunbeam` · gateway `origin/dev` · others `origin/main`. Mobile's design-revamp branch is secondary evidence only. This is a code snapshot, not deployment certification.