

01

THE PROBLEM

Engineering teams lose their best hours to knowledge transfer that runs blind.

The senior engineer answers the same architecture question for the third time this month. The new hire spends week six interrupting the same two people to learn why the system is built the way it is. Support escalates a case the team already solved in March. The knowledge exists, but the reasoning behind it was never written down, and nobody can see what is missing until someone hits the gap. Facts are easy to write down. Reasoning never is. **Mentat captures both.**

58%

of developer time goes to program comprehension, reconstructing context that already lives in someone's head.


Xia et al., IEEE TSE 2018 · 78 devs, 3,148h

3.5–4.5h

daily focus time the median engineering team gets. The rest is interruptions and re-explanation.

Worklytics 2025 · 3.4M pull requests analyzed

THE LOOP
the product, in one pass



Mentat answers questions from what the team wrote, detects what was never written, reconstructs it through a short AI interview with the person who holds it, and routes the answer to the human or agent who needs it next.

02

WHERE IT LANDS

The senior-engineer interrupt tax.

Ask Mentat, in Slack or on the web, answers from everything the team has written plus everything Mentat has reconstructed. When the confidence bar is not met, it says so and routes the question to the person who holds the answer.

New-hire ramp.

Ramps runs a day-one interview, builds a personal onboarding plan, tracks milestones, and nudges through Slack. Week six stops being spent interrupting the same two people.

Support-to-engineering escalations.

The same case stops landing on engineering twice. Support gets the answer at the level they need, the solution rather than the table IDs, phrased to hand straight back to the customer.

Your coding agents.

Mentat is an MCP server with a full audit log. Agents pull the reasoning behind the code, not just the code, and what they capture routes back through the same loop.

03

WHY NOT THE EXISTING STACK

ALTERNATIVE	WHAT IT DOES	WHERE IT STOPS
Wikis · Confluence, Notion	Store what people write	Engineers don't write the reasoning, and 42% of unique knowledge sits with a single employee (Panopto/YouGov).
Search / RAG · Glean, Guru	Retrieve what exists	Stops at retrieval. It cannot reconstruct what was never written.
AI coding · Cursor, Copilot, agents	Ship code faster	The reasoning behind the code reaches neither the next human nor the next agent.
Agent memory · mem0, Letta	Per-agent recall	No model of who holds the live answer and who needs it next.

Retrieval is where Mentat starts, not what Mentat is.

04

SHIPPED & RUNNING

Mentat ingests Slack, GitHub, Confluence, and Notion, runs adaptive AI interviews on the web and fully inside Slack, and keeps a temporal knowledge graph with provenance on every answer. Weekly briefings surface what changed. Ramps handles new-hire onboarding, and an MCP server with an append-only audit log serves the team's agents. Running our first design-partner pilot now. Onboarding a small number of additional design partners.

\$49.99

per seat / month, 1 to 200 seats, monthly or annual

Setup is guided. Book a demo and we stand it up with you.

aimentat.dev

alex@aimentat.dev