

Forum: [Hack Forums](#)

Thread: [Claude Code just added /dream - a "sleep cycle" for the memory system](#)

Oh_Well - March 25, 2026, 1:08 am

Claude Code just shipped one of the smartest agent features I've seen.

It's called Auto Dream.

Here's the problem it solves:

Claude Code added "Auto Memory" a couple months ago, where the agent writes notes to itself based on your corrections and preferences across sessions.

Great in theory. But by session 20, your memory file is bloated with noise, contradictions, and stale context. The agent actually starts performing worse.

Auto Dream fixes this by mimicking how the human brain works during REM sleep:

It reviews all your past session transcripts (even 900+)

Identifies what's still relevant

Prunes stale or contradictory memories

Consolidates everything into organized, indexed files

Replaces vague references like "today" with actual dates

It runs in the background without interrupting your work. Triggers only after 24 hours + 5 sessions since the last consolidation. Runs read-only on your project code but has write access to memory files. Uses a lock file so two instances can't conflict.

What I find fascinating:

We're increasingly modeling AI agents after human biology e.g sub-agent teams that mirror org structures, and now agents that "dream" to consolidate memory.

The best AI tooling in 2026 isn't just about bigger context windows. It's about smarter memory management and it's actually exciting to see where this goes.

? loading encrypted image...

CL1B - March 25, 2026, 2:59 am

Interesting, isn't this what the ClawdBot or OpenClaw really accelerated at?

So is this Claude's response to that? I've yet to use this continues memory style AI.

Oh_Well - March 25, 2026, 3:53 am

(Mar 25, 2026 02:59 AM)MarlboroMan Wrote: [\[CITE\]](#) Interesting, isn't this what the ClawdBot or OpenClaw really accelerated at?

So is this Claude's response to that? I've yet to use this continues memory style AI.

I wouldn't consider this the same thing. OpenClaw basically let the agent remember more across sessions, but they didn't solve memory quality over time. I think /dream is less about remembering more, and more about remembering better (instead of just extending memory indefinitely, it

periodically compresses and rewrites it)

CL1B - March 25, 2026, 4:39 am

(Mar 25, 2026 03:53 AM)Oh_Well Wrote: [\[CITE\]](#) I wouldn't consider this the same thing. OpenClaw basically let the agent remember more across sessions, but they didn't solve memory quality over time. I think /dream is less about remembering more, and more about remembering better (instead of just extending memory indefinitely, it periodically compresses and rewrites it)

Interesting. Instantly thinking about built in personal assistants like Siri on Apple's OS; They need this!!

Chorozon - March 25, 2026, 6:01 am

This is interesting.

I guess the blueprint for the best AI is truly the human mind

Au Jus Demon - March 25, 2026, 7:54 am

(Mar 25, 2026 06:01 AM)Chorozon Wrote: [\[CITE\]](#) This is interesting.

I guess the blueprint for the best AI is truly the human mind

With them printing brain cells, and running Doom on it. Who knows maybe in the future they are legit going to run off a brain.

Saterax - March 29, 2026, 9:31 pm

I was just mentioning in another thread, pushing the whole thing into an obsidian vault has basically done this and has been amazing, keeps everything pretty damn neat.

This is neat too, but I feel like it might still have some drawbacks, having it essentially read it's own memory depending on the amount of time or sessions can get a bit iffy, since sometimes it's still trying to read it all in one go, which the longer it gets... The more context it loses, so possibly by the end of the dreamscape, it's still a bit jumbled.

The obsidian vault on the other hand has been like giving Claude a notebook, it can write or read any part of the notebook at any point. Ask a question about something specific and it will basically use the equivalent of CTRL + F on the entire notebook, so rather than going through and reading page by page to find something specific (Absolute massive waste of tickets), it can instead rely on a search system.

Like with my workspace thingy, it has a few different 'context' files. Every session on initialize it reads these five files;

user.md - Basically a "who I am" and it will add/remove stuff based on what is said. Like a very simple short biography of me so it has it for context.

feedback.md - Basically keeps track of corrections and confirmed approaches from other sessions. If it says something or does something that I don't like, I can basically say "Hey, don't do ____" and it will automatically note "don't do ____" or if there is a certain approach that works better to improve efficiency, it will make note of that too.

projects.md - Super short summary of each project, and from here points to each projects individual context file.

session-log.md - Basically a snapshot of whatever session was most recently/is in progress. Kind of a backup to the built-in claude memory. Sometimes I forget I'm doing something and come back after a few hours, if you leave Claude open and do that, sometimes it will break things super bad if you try to just randomly continue lol.

If I start a new project and it goes far enough to have some sort of detail required to be remembered, it gets its own markdown file created in that specific project directory, this is each projects individual context file. Each context file is auto-written at the end of each session and includes a lot of information regarding the project that the folder is dedicated to. Notes, solutions, issues, main points, etc. This is what the projects.md points to, so at all times it knows somewhat all the projects that have been done and what is available, it just doesn't have (nor need) every single detail of each project

The search tool (ai-grep) automatically indexes everything in the vault and is a separate system from Claude that it can use very easily.

So at any time, I can ask it
"Hey so do you remember what we did for (project)? I think this might be a similar case, can you go see what we did to make it worse?"

It will search for the project and use ai-grep, using project keywords and context of the current moment to search for that specific piece, so a large file of a few thousand lines is already indexed and Claude finding a specific answer takes seconds instead of line by line.

The fifth file is just claude.md which is basically a short rulebook and includes the search tool commands, so at the beginning of the session, it already knows what it can do.

The nice thing with it is that you can see and edit everything in Obsidian directly. It's an entirely self-contained system that you host and saves an absolute massive shitload of tokens and trouble trying to re-explain context of something.