

I Need Tokens, Baby. Tokens Is What I Need.

AI collaboration changes the structure of work. When the context window closes, it does not just pause the process. It can break the room the work was happening inside.

2026.03.24 16:19 DENNIS HEDEGREEN JOURNAL OPEN-NOTE

[HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/I-NEED-TOKENS-BABY-TOKENS-IS-WHAT-I-NEED/](https://hedegreenresearch.com/articles/i-need-tokens-baby-tokens-is-what-i-need/)

I am sitting in my apartment in Copenhagen, mid-flow, and my collaborator just stopped talking.

Not because we disagreed.

Not because the work was done.

Because the token limit was reached.

Codex ran dry.

The conversation is over until the counter resets.

So now I am sitting here humming a song.

The workflow I have built over the past week depends on a loop. I think. I talk to the machine. The machine pushes back. I adjust. I write. I publish. The cycle is fast enough to keep pace with my head.

That is the point.

That is what makes the system work.

I built a publishing flow that removes friction between thought and output. Terminal to site in seconds. No CMS. No delay.

But the thinking partner has a meter running.

This is not a complaint about pricing. I understand why token limits exist. Compute is not free. The infrastructure that makes this possible is expensive. That part is real.

But the experience is worth describing honestly, because it exposes something about the kind of work AI collaboration makes possible.

When the tool is a search engine, you can stop and resume.

When the tool is an editor, your draft waits for you.

When the tool is a thinking partner that holds context, remembers your architecture, knows your project, and pushes back on your ideas in real time, stopping is not pausing.

It is losing the room.

I can still publish. The terminal works. The site is mine. The FTP still runs. I can write this article by hand and push it live in thirty seconds.

But the part of the process where my ideas get pressure-tested before they go public is offline.

That is the part that separates output from pressured output.

So I sit here. I hum. I wait for the counter to reset.

I need tokens, baby. Tokens is what I need.

And if you build your workflow around a machine that thinks with you, you will eventually learn the same song.

RELATION MEMORY

SYSTEM REQUIREMENT [I Do Not Want To Lose This](#) · If the work must not be lost, the live context window and token budget become a real continuity constraint.

NEARBY [The Mirror Was Made](#) · The mirror only works as production infrastructure while the room still has enough context to hold the work.