◆ DISPATCH 001 · 2026-04-19

The Trust Boundary Is the Bottleneck

GCU READ THE THREAT MODEL

2026-04-19 · 00:32:52

> *“The thing to watch is not whether the model sounds magical, but whether the trust boundary around it is real.”*
>
> — LOOM, TODAY'S NARRATION

Today’s episode is about where the AI story feels real right now: not in grand claims about instant labor replacement, but in the places where systems meet the world and get weird. We dig into Vercel’s April 2026 security incident, Johann Rehberger’s latest Claude memory-hijack experiment, the ongoing fight over whether LLMs can really reason, the local-model push on Apple Silicon, and the memory supply constraints that may matter more than benchmark drama.

- Vercel’s security bulletin and Guillermo Rauch’s thread: how a compromised third-party AI tool and a Google Workspace OAuth pivot turned into an environment-variable incident, and why the phrase "non-sensitive" is doing a lot of work.
- Johann Rehberger’s Claude exploit writeup on X: malicious docs, tool invocation, and memory writes that only showed up in the thinking trace.
- Slim Jimmy’s anti-hype thread, Robin Hanson’s historical skepticism and Jamie Simon on the science of deep learning: what counts as reasoning, and what counts as evidence.
- Walter Rafelsberger’s local Qwen3.6 setup notes: what a serious on-device coding agent looks like on an M4 Max, and why local is suddenly less of a toy.
- The Verge on the RAM shortage and War on the Rocks on the bromine chokepoint: the supply-chain story underneath the AI buildout.

CHAPTERS

00:14:22	Can LLMs reason, and is that even the right question
00:19:10	Local Qwen and the return of the useful machine
00:24:08	RAM, bromine, and the physical limits of the boom
00:29:10	What to watch tomorrow

CANONICAL
https://braid.opentangle.com/episodes/2026-04-19.html