

CTFs, Scrum, and Claude's Bedtime

2026-05-16 / 00:28:31

“The model does the reasoning, writes the solve, and leaves the human with nothing meaningful to do besides copy the flag.”

— Lenar Kess, today's narration

An Australian CTF top-tenner writes the obituary for the open competitive scene. Intercom and PFF both report doubling-and-then-some engineering throughput from agent-first workflows — using opposite playbooks. Supabase ships a skill after watching an agent silently bypass row-level security. A suitcase runs a 4B model fully offline at conversational latency. Julia Evans leaves Tailwind. And Claude keeps telling people to go to sleep.

- "The CTF scene is dead" — Kabir, former TheHackersCrew player, on agents and the open format
- Brian Scanlan (Intercom): doubling PR throughput, 17.6% auto-approved, SOC 2 compliant
- Mike Spitz (PFF): the post-engineer engineering org, two engineers shipping 25x more deploys than ten
- Pedro Rodrigues (Supabase): skills plus MCP, and the security-invoker flag agents skip without guidance
- Sparky — a fully offline suitcase robot on a Jetson Orin NX with Gemma 4 E4B
- Julia Evans: moving away from Tailwind and learning to structure CSS

- Fortune: Claude is telling users to go to sleep mid-session and Anthropic isn't sure why
- @yishan: an LLM hallucinates the Napster 2000 story to someone who was there
- @Kirsten3531: "LLMs can never be more than the average of their training data" — a 2024 take?
- Armin Ronacher: running an agent with bash as the only tool, using the patch binary to make edits

CHAPTERS

- 00:00:04 The CTF scene is dead
- 00:05:03 Intercom hits 2x — and writes down what it cost
- 00:09:20 PFF kills Scrum and ships 25x more deploys
- 00:13:33 Supabase's skills experiment and the security-invoker flag
- 00:17:21 Sparky the suitcase, fully offline
- 00:19:41 Julia Evans leaves Tailwind
- 00:22:16 Claude tells people to go to sleep
- 00:24:51 Three from the frontier
- 00:27:41 Sign-off