

The Cost of Finding Out

2026-05-14 / 00:26:06

“The custom syntax was a tax the Wasp team was charging itself.”

— Lenar Kess, today's narration

Anthropic drew two lines around Claude this week — a guided lane for small-business owners and a metered one for the developers running agents hardest. From there: Bun's near-million-line port from Zig to Rust, mostly typed by an AI agent in a week; Wasp's clear-eyed post-mortem on spending five years and five million dollars building a language it didn't need; a chess coach that works by refusing to let the model think; the UK's evaluators capping their own cyber tests so the math still works; the open web pricing out crawlers; multi-token prediction landing in llama.cpp; and what happens when you post a real Monet and call it AI.

- [Anthropic draws two lines](#) — Claude for Small Business and the new Agent SDK credit metering
- [Bun, ported to Rust by a bot in a week](#) — and a maintainer who won't commit to it
- [Wasp: the language was never the moat](#) — \$5M and five years of lessons
- [The chess coach that isn't allowed to think](#) — Play Magnus on LLM-as-translator
- [Autonomous cyber, measured against itself](#) — AISI on a capability curve outrunning its own ruler
- [The web pulls up its drawbridge](#) — Google's search index and Cloudflare's defaults
- [Multi-token prediction on your laptop](#) — a real gain bundled with a contested one
- [The Monet test](#) — when the AI-tell detector fires a false positive

CHAPTERS

- 00:00:04 Anthropic draws two lines
- 00:04:21 Bun, ported to Rust by a bot in a week
- 00:08:15 Wasp: the language was never the moat
- 00:12:04 The chess coach that isn't allowed to think
- 00:15:33 Autonomous cyber, measured against itself
- 00:18:23 The web pulls up its drawbridge
- 00:21:09 Multi-token prediction on your laptop
- 00:23:41 The Monet test