

Fogbank for Code, Pinned Carriers, and a Frontier Model in 90 Gigabytes

👉 GSV FOGBANK FOR CODE

2026-04-26 · 00:25:53

“When juniors skip debugging and skip the formative mistakes, they don't build the tacit expertise — and when my generation of engineers retires, that knowledge doesn't transfer to the AI.”

— LENAR KESS, TODAY'S NARRATION

A Sunday lineup that runs from a 2-bit quant of DeepSeek V4-Flash on a laptop, through a Java production hang most teams haven't met yet, into the most considered pushback I've read on the AI coding numbers we discussed yesterday. Plus Cloudflare on why MCP tool dumps don't scale, the Asahi team finding a VRR knob hiding in plain sight, and a quiet note from DHH about laptops you might suddenly be able to afford.

- [DeepSeek-V4-Flash 2-bit dynamic quant on 128GB Macs](#)
- [The full V4 collection on mlx-community](#)
- [Shubham Raizada on virtual thread pinning](#)
- [Denis Stetskov: The West Forgot How to Build. Now It's Forgetting Code](#)
- [Matt Carey, Cloudflare: MCP = Mega Context Problem](#)
- [Asahi Linux progress report for Linux 7.0](#)
- [DHH on the AMD HX370 holding up against Panther Lake](#)
- [Eden AI as a European OpenRouter alternative](#)

CHAPTERS

- 00:00:04 A frontier model in 90 gigabytes
- 00:02:36 Pinned carriers, and a deadlock most Java teams haven't met yet
- 00:07:09 Fogbank for code
- 00:12:29 Why MCP tool dumps don't scale: Matt Carey at AI Engineer

00:16:32 Asahi Linux 7.0: a VRR knob hiding in plain sight

00:20:22 DHH on a real workload, and a quiet note about laptop deals

00:22:04 Eden AI, GnuPG goes post-quantum, and one for the watchlist

00:24:17 Carrying forward

CANONICAL

<https://braid.opentangle.com/episodes/2026-04-26.html>