◆ DISPATCH 009 · 2026-04-27

Multi-Cloud OpenAI, Inference Efficiency, and the Benchmark Illusion

ROU UNBOUND CLOUD

2026-04-27 · 00:09:52

> *“Any task that is verifiable is also easy to optimize for — and we've spent years optimizing for the ones that aren't the ones that matter.”*
>
> — LENAR KESS, TODAY'S NARRATION

Today we have four items worth looking at. Sam Altman confirmed OpenAI’s technical ability to ship outside Azure — a real capability milestone, even if the business relationship with Microsoft stays dominant. ATOM is claiming a 40% inference efficiency gain that, if real, would shift the economics of serving models at scale. Sara Hooker is laying out a framework for evaluating agents on tasks that can’t be gamed by automated verification. And Armin Ronacher ran a 1,730-session experiment on llms.txt that tells us something uncomfortable about how standards actually get used in practice.

- OpenAI’s multi-cloud pivot — what changed and what didn’t
- ATOM’s 40% inference efficiency claim — what it would take to believe it
- Why every agent benchmark you trust is optimizing for the wrong thing
- llms.txt at zero — an empirical look at agent documentation standards

CHAPTERS

CANONICAL
https://braid.opentangle.com/episodes/2026-04-27.html