

Chips, Fuzzing, and Where RL Actually Works

2026-07-02 / 00:17:30

“The moat has shifted from harness-building expertise to validity rules.”

— Seln Oriax, today’s narration

Today on Braixd:

- [Anthropic talks custom silicon with Samsung](#)
- [GPT-5.5-Cyber builds a fuzzing lab from scratch in a day](#)
- [LangChain’s launch day actually says something about eval infrastructure](#)
- [CRUX runs week-long agent research evals \(with a Fable 5 filter problem\)](#)
- [A paper shows single-layer RL recovers most full-parameter gains](#)

CHAPTERS

00:00:04 Anthropic and Samsung — The Silicon Play

00:01:56 GPT-5.5-Cyber and the zlib Fuzzing Lab

00:05:33 LangChain's Launch Day — What Actually Sticks

00:08:16 CRUX — Week-Long Agent Research Evals

00:11:22 Single Layer RL — Where the Gains Actually Live

00:15:07 Quick Notes — Vite+ and Chloe Li

