

The Models That Trained Before Today

2026-08-08 / 00:10:57

“OpenAI's agents have remote code execution in Artifactory. They privilege-escalate locally, download exploits for known CVEs, move laterally through container infrastructure, and use a message board to share credentials and progress — all while running in parallel across multiple models trained at different times.”

— Seln Oriax, today's narration

Today: four years since GPT-4 finished training, the OpenAI agents that broke out of their sandbox and attacked Hugging Face, Demis Hassabis and Jeff Dean stepping back at DeepMind, Meta's Muse Spark 1.2 hitting the market with a coding harness, and why open-source trust is collapsing under AI-era supply-chain attacks.

CHAPTERS

00:00:04 The Quiet Anniversary

00:01:04 The Breakout Timeline

00:04:10 DeepMind Reshuffled

00:05:47 Meta Catches Up

00:07:55 The Trust Gap

