



EXCLUSIVE · AI

OpenAI's GPT-7 leaks reveal a 12-trillion parameter monster — and a new approach to memory

Internal documents obtained by ByteDispatch show the next OpenAI flagship pairs a sparse mixture-of-experts core with a separate persistent-memory layer, marking the company's biggest architectural shift since GPT-4.

by Mira Chen and Jonas Park · April 26, 2026 — 6:30 AM PT · 8 min read

[Read the full story →](#)

LATEST

AI

Anthropic closes \$20B round at a \$300B valuation

Lightspeed, Sequoia, and a sovereign fund lead the company's third mega-round in 14 months.

by Priya Singh · 2 hours ago · 4 min

HARDWARE

Apple's M5 Ultra benchmarks leak — and they crush every competitor

The unreleased Mac Pro chip posts a 38% multi-core lead over AMD's Threadripper 7995WX in early Geekbench results.

by Tomás Reyes · 5 hours ago · 6 min

SOFTWARE

Linux 7.0 ships with a redesigned memory model and 64K page

STARTUPS

Y Combinator's W26 batch is 47% AI agent companies — here's the

support by default

Linus Torvalds calls it the largest VM rewrite in two decades; cloud providers begin rolling it out next quarter.

by **Aoife Donnelly** · Yesterday · 7 min

REVIEWS

We tested 14 robot vacuums on stairs; only 3 didn't tumble

Cliff sensors are still inconsistent in 2026. Roborock and Dreame top our chart; one \$1,400 model went over the edge twice.

by **Elena Markova** · 2 days ago · 11 min

HARDWARE

NVIDIA RTX 6090 Founders Edition review: too much GPU for any one task

Three-slot, 600W, and faster than the 5090 by 70% in raster. We ran it for two weeks across games, ML training, and 8K video editing.

by **Tomás Reyes** · 3 days ago · 14 min

STARTUPS

Stripe launches embedded banking — will it eat Mercury and Ramp?

Stripe Banking ships with FDIC-insured accounts, programmable cards, and a treasury layer; pricing starts at \$0/month.

by **Priya Singh** · 4 days ago · 8 min

full list

Vertical agents for legal, accounting, and supply-chain operations dominate, with consumer agents taking a backseat for the first time since 2023.

by **Marcus Bauer** · Yesterday · 12 min

AI

Google's Gemini 3 Ultra hits 92% on the ARC-AGI-2 benchmark

DeepMind's latest model becomes the first to clear the human-baseline mark on the famously hard reasoning suite.

by **Mira Chen** · 2 days ago · 5 min

SOFTWARE

Rust 1.99 stabilizes async traits and ships incremental compilation v3

Five-year-old async-trait shim is finally redundant; cold builds drop 22% on average across the top 1000 crates.

by **Aoife Donnelly** · 3 days ago · 6 min

REVIEWS

Vision Pro 2 review: still magic, still \$3,500, still confusing

Lighter, sharper, and finally useful for productivity — but Apple still hasn't answered the basic question of who this is for.

by **Elena Markova** · 5 days ago · 16 min

- 1 GPT-7 leaks: 12T parameters
- 2 M5 Ultra crushes Threadripper
- 3 Stripe vs. Mercury, finally
- 4 Robot vacuums on stairs
- 5 Rust 1.99 ships async traits

THE DAILY DISPATCH

One email, every morning, with the five tech stories that matter.

Subscribe

MOST READ THIS WEEK

- 1 YC W26: 47% AI agents
- 2 Gemini 3 cracks ARC-AGI-2
- 3 RTX 6090 reviewed
- 4 Vision Pro 2: \$3,500 magic
- 5 Linux 7.0 lands



ByteDispatch

[About](#)

[Contact](#)

[Privacy](#)

[Terms](#)

[RSS](#)

[Twitter](#)

[GitHub](#)

[LinkedIn](#)

[Bluesky](#)