Opt-Out Goes Universal, Space TPUs Boot Up and Agents Go Open-Source

Four labs agree on a single robots.txt signal for AI training, Google runs Gemma 4 on TPUs in orbit, OpenAI promises a public Community Safety Report, and Hermes Agent is the open-source project everyone is starring.

Themes: Governance & Safety · Infrastructure & Economics

Note: search tools returned mostly aggregator digests dated October 30-31 and no primary company posts; specifics are provisional and should be checked against the original announcements.

1) One robots.txt Signal to Opt Out of AI Training

Google, Microsoft, OpenAI and Anthropic said they will honour a single shared signal in a site's robots.txt that tells AI crawlers not to use its pages for training, replacing the patchwork of one rule per company. The signal comes from an IETF draft co-written by Kevin Kelley of Anthropic. Google and Microsoft are reported to begin honouring it in 2027, and the IETF has not endorsed it, so it has no formal standing yet.

The concept, explained simply: robots.txt is a plain text file at the root of a website that tells automated crawlers which pages they may fetch. An opt-out signal is one extra line saying 'do not use this content to train models', which compliant crawlers read before collecting data.

Why it matters: Publishers have had to track a growing list of crawler names to block training use. One signal that four of the largest labs respect lowers the cost of saying no. It is voluntary, though: it only works for companies that choose to comply and does not cover data already collected.

2) Google Boots Up TPUs in Orbit

Google says it successfully booted the Tensor Processing Units aboard its satellite, and that the on-board Gemma 4 model responded correctly to a prompt. Details come from digests rather than a full technical post, so power, cooling and performance figures are unknown.

The concept, explained simply: A TPU is Google's custom chip for the matrix maths behind neural networks. Running a small open model such as Gemma 4 in orbit is 'edge inference': the model runs where the data is produced instead of sending raw data back to Earth.

Why it matters: Space-based compute is attractive for processing satellite imagery on board and, longer term, for cheap solar power. For now this is a proof of concept that chips and a small model survive launch and radiation, not a competitor to data centres on the ground.

3) OpenAI Plans a Community Safety Report

OpenAI announced a Community Safety Report, with the first edition planned for early 2027. It also disclosed more than 200 referrals to law enforcement in the third quarter of 2026. This is from a digest, not checked against OpenAI's own post.

The concept, explained simply: A transparency report is a regular public summary of how a company handled misuse, such as how many cases it flagged and how many it passed to authorities. It lets outsiders see trends instead of trusting a one-off statement.

Why it matters: As chatbots and agents become part of daily life, regulators and parents want numbers on harm and response. Published figures set a baseline other labs may be pushed to match, and they also raise privacy questions about when a lab reports users.

VIRAL / OPEN-SOURCE SPOTLIGHT

Hermes Agent (Nous Research, MIT license): a roundup of GitHub trends reports it gained nearly 20,000 stars in a single week, among the fastest-growing AI projects. It is an open agent that plans tasks, calls tools and keeps working over many steps, built to run on open models. It is taking off because a free, permissively licensed agent lets anyone try autonomous workflows without a paid subscription. I could not verify today's exact star count, so treat this as a press-roundup candidate. A runner-up is Google AI Edge Gallery, which runs open models including

Gemma 4 offline on phones, a neat companion to today's edge-inference story. Run any agent with least privilege, as in earlier days.

MARKET SIGNAL

Cooperation is appearing where the cost of conflict is rising: rival labs agreeing on one opt-out standard, and Google probing whether compute can move off-planet, both reflect pressure over data supply and power. Meanwhile open-source agents attract developer attention faster than any closed product. Figures here come from aggregators and are unverified.

PRACTICAL TAKEAWAYS

1) Add the AI training opt-out to your site.

If you publish content, watch for the final robots.txt syntax and add it once the draft stabilises, while keeping existing per-crawler rules.

2) Treat the signal as a request, not a lock.

It only binds compliant crawlers, so sensitive material still needs authentication or paywalls.

3) Consider edge inference for bandwidth-heavy data.

Where data is large or remote, running a small model next to the source can beat shipping everything to the cloud.

4) Read transparency reports critically.

Look for definitions, trends over time and what is excluded, not just headline counts.

5) Sandbox open-source agents before trusting them.

Try Hermes Agent in a container with limited credentials and review what tools it can call.

SELF-TEST

Q1. What does the shared robots.txt signal do?

A1. It tells compliant AI crawlers not to use a site's content for training, in one rule that several labs honour.

Q2. Why is the signal not a guarantee?

A2. It is voluntary, not endorsed by the IETF, and does not remove data already collected.

Q3. What is edge inference?

A3. Running a model where the data is produced, such as a satellite, instead of sending raw data elsewhere.

Q4. What is a transparency report?

A4. A regular public summary of misuse handling, such as flagged cases and referrals to authorities.

TOMORROW

What We're Watching: the exact robots.txt directive name and whether other labs and crawlers such as Meta join the standard. What We're Watching: Google's technical write-up of the orbital TPU test. What We're Watching: OpenAI's own post on the Community Safety Report. What We're Watching: whether Hermes Agent's growth holds or fades.

TOPICS COVERED SO FAR (RECAP)

Recent days covered Gemini 4 Carbon and staged releases, DeepSeek V4.1-Flash and MIT licensing, agent harnesses, the AI payback debate, Mistral Large 4 and open weights, Anthropic's IPO filing, Claude Opus 5.5, Nvidia's Vera Rubin POD, GPT-6.1 Astra, multi-agent design and offensive use of autonomous agents. Today added the shared AI training opt-out, orbital TPUs and edge inference, transparency reporting and the Hermes Agent surge.

Sources: thenextweb.com, walterstein.eu, guruchahal.substack.com and github.com/diclogic/ai-daily-digest on the opt-out standard and other October 31 items; besthub.dev and trendshift.io on trending open-source projects.