

Meta will ship a harness and open models - Alexandr Wang, YC SUS

The weekend handed out positions: YC Startup School filled a stadium, and two statements came out of it that continue yesterday's open weights thread. Plus two security stories worth reading as instructions.

Monday, 27 July 2026

IN THIS ISSUE

1. Meta will ship a harness and open models - Alexandr Wang, YC SUS
2. Jensen on the same stage - a verbatim quote on open source
3. The relay market: other people's Anthropic tokens at 2% of the price
4. Opus was down for an hour and forty, and nobody noticed
5. GrapheneOS: the case where an OS is being made a crime
6. Stanford measured what AI did to the labor market
7. @levelsio: "vibe-coded 100% of my SaaS subscriptions"
8. Ruff turns on 413 rules instead of 59
9. Quiet evidence in the open weights dispute: the attacks come from subscriptions
10. Mollick: a guide to models went stale two days after it was written

1. Meta will ship a harness and open models - Alexandr Wang, YC SUS A continuation of yesterday's open weights story, this time an announcement.

Quote from the stage: **"Meta will ship a harness and open-source models"**. The thread itself verified it: under the post, **@alexandr_wang** replied **"yep"**. He also handed everyone in the room **\$1000 in API credits** for the new Muse Spark.

His framing of competition, from the same stage: **"with agents, a startup against an incumbent is no longer David against Goliath, it's mecha-Goliath against Goliath"**.

Why it matters. The word "harness" is the key one here. A harness is what gets written by hand around the model: the prompt, the tools, memory, the scheduler. When Meta starts giving away the harness along with the weights, that competes directly with the build-it-yourself approach. The skepticism in the thread is healthy (the hardest line: "an LLM without a harness is worth less than Chinese open weights"), but the direction is worth noting. When it ships, look at the harness and ask what in it beats your own wrapper.

2. Jensen on the same stage - a verbatim quote on open source Yesterday it was his point about steering agents. Today there is a second one, and it lands right in the argument from item 1 and from the letter two days ago.

Verbatim: **"Without open source, the mobile cloud industry would never have happened. And more importantly, PyTorch... remember the early versions, Caffe, Torch, Theano, all of it was open. Without all that, where would modern AI have come from?"**

Why it matters: this is the strongest form of the argument in the weights dispute: the whole stack everything runs on is already open.

3. The relay market: other people's Anthropic tokens at 2% of the price A Vectoral investigation, 172 points. The most practical security story of the day.

Numbers from the source:

- a package worth **\$3,333 of official Anthropic credit for 425 yuan (~\$60)**, which is **\$0.13 of value per \$1**
- the median discount among top relays is **97.8%** (Now Coding), and it holds across all 49 on their list
- the scheme has four layers: card merchants with virtual cards for US/EU billing → account pools (hundreds of accounts, rotating limits) → relays with clean Chinese-language UIs → end buyers, including those chasing distillation

Why it matters: the economics of these relays rest on **mass account registration and stolen or virtual cards**. That is why labs keep tightening anti-fraud, and why cheap "Claude API keys" from the internet are always either stolen or pooled. When someone offers "the same Opus, cheaper", it helps to know how that works underneath.

4. Opus was down for an hour and forty, and nobody noticed 92 points on HN, Anthropic's status page. Verbatim from the source: **"Elevated errors for Opus 5"**, affecting `claude.ai`, Console, the API, **Claude Code and Claude Cowork**.

Timeline: investigating **09:17 UTC** → identified 09:45 → fix 10:34 → resolved **10:44 UTC**. That is **1 hour 27 minutes** of trouble, with the active phase running 12:17-13:44 Kyiv time on Sunday.

Why it matters: any pipeline resting on one model and one provider with no fallback simply goes quiet during hours like these, and people blame a hang. A fallback to another model is extra complexity for ninety minutes a month. The cheaper insurance: a script that polls the provider's status page every 15 minutes and posts to chat when an incident touches the API. No LLM, no separate infrastructure.

5. GrapheneOS: the case where an OS is being made a crime 378 points, top of the night. US prosecutors charged Atlanta resident Sam Tunick under a federal statute on **destruction of property to prevent its seizure**: his GrapheneOS phone wiped itself during an attempted airport search.

Details from the source: he was detained at Hartsfield-Jackson after returning from the Dominican Republic. Agents already had his name and photo flagged "suspected terrorist activity" over ties to the anti-Cop City movement. The defense says the CSAM questioning was a **pretext**, that he asked for a lawyer **four times** and was refused each time, that no warrant was shown and no rights were read. An EFF expert and a surveillance specialist say they have not seen anything like it: **"this signals that GrapheneOS is criminal by default"**.

Alongside it, at 399 points, a technical writeup from GrapheneOS themselves. Rate limiting in the secure element: **4 hours after 10 attempts, 41 days after 15**, with a maximum of 20. Secure element firmware cannot be updated without owner authentication, so a government cannot force the manufacturer to lift the limit. Plus a Motorola partnership from 2027, so not only Pixel.

Why it matters: the "wipe on duress password" feature is legally toxic exactly where it is needed most, at the border. For people who fly, this stopped being hypothetical. The practical conclusion is boring and it works: at the border, a **clean phone with nothing extra on it wins**.

6. Stanford measured what AI did to the labor market 250 points, 320 comments. A SIEPR policy brief, peer-reviewed economics.

Key findings verbatim: **the effect of AI on aggregate employment right now is most likely small**; the hard market **for graduates** is partly due to AI; the effect on worker productivity is **mixed but positive overall**; adoption inside firms has accelerated, but **very unevenly**. A caveat from the authors themselves: early empirical work is not the final word.

Separately, the authors cite Dario Amodei's forecast (AI wipes out half of white-collar jobs, 20% unemployment) as an example of **AI leaders themselves amplifying the fear**.

Why it matters: item 7 below should be checked against this. @levelsio, at 347K views, writes that the industry is being washed away; Stanford says the aggregate effect is still small. Both are right about different things: the indie niche really is first in line, while the labor market overall holds. For a CTO that difference matters more than either headline.

7. @levelsio: "vibe-coded 100% of my SaaS subscriptions" The loudest conversation of the day in the first newsletter: Patrick Haede's post "the indie hacker product is dying, and coding agents are why" pulled 225K views, and levelsio's reply pulled 347K.

His claim, unvarnished: **"Cancelled and vibe-coded 100% of my SaaS subscriptions. The only things I pay for now are domains, hosting, storage and AI APIs. The irony is that the replacement is coming for him too, and he knows it"**. Plus an observation in numbers: his own traffic and revenue **are falling**, and so are those of indie people he knows; Yongfook's Google referrals **halved** over the year.

The soberest counterpoint in the same thread (@onlinedopamine): complex products sold to **non-technical IRL businesses** - plumbers, cleaning services, pool companies - are holding, because those people will never sit down and vibe-code a replacement.

Why it matters: the frame is simple. What an agent can recreate in an evening is exposed; distribution, data and a real non-technical customer are protected. This is the first time it has been said by someone who made money that way and is showing his own falling numbers.

8. Ruff turns on 413 rules instead of 59 336 points. Astral raised the linter's default rule set from **59 to 413 rules** (Ruff has 968 in total), a default untouched since v0.1.0. The logic:

plenty of rules catch syntax and runtime errors but were sitting disabled. Rollback is one line: `select = ["E4", "E7", "E9", "F"]`.

Correction: **the release itself is from 23.07**, and what fell into the window is the HN thread.

Why it matters: any Python codebase without Ruff will dump a pile of findings on the first run under the new default. That is what an updated default does.

9. Quiet evidence in the open weights dispute: the attacks come from subscriptions A fine detail highlighted by @amasad: a former Anthropic employee (@NoahLebovic) writes that in practice **hackers prefer the heavily subsidized lab subscriptions**, while open weights stay on the sidelines.

Why it matters: if that holds, the "open weights = a weapon" argument sags from an unexpected side. In practice the weapon turns out to be cheap frontier access, which is exactly what the relays in item 3 sell. Two stories of the day converged. Honesty tag: this is one person's word, not a measurement, unlike yesterday's AISI numbers.

10. Mollick: a guide to models went stale two days after it was written The guide on "which model for what" was written on Thursday, and already had to be rewritten for **Opus 5 and Codex voice mode**, both of which shipped on Friday. The conclusion: **"keeping up is hard even when you follow this closely"**.

In a day, three city-builder games were made from one prompt (Cezanne, Piranesi, and a GPT-5.6 Sol version) and **the sources were published**, an illustration that the cost of that kind of experiment is now zero.

Why it matters: an argument for Boris Cherny's advice from yesterday's item 4: wrappers age faster than they can be written.

Misc • Sam Altman closing YC SUS: **"The next six months will deliver as much model progress as the last two years"** - a claim worth checking in January • Garry Tan from the same stage: **"build a personal AGI - a brain nobody curates is a dumpster with good search"**; and separately, "don't LARP, be sincere"

- Naval, 282K views: "if open weights had backdoors, the closed labs would definitely have found and shown them" - the most elegant argument for openness all week
- Chelsea Finn (Physical Intelligence) at YC: RL for robots runs into **the cost of physical rollouts** - 1M trajectories of a minute-long task is about **700 robot-days**
- Google disclosed **\$94.1B in SpaceX shares - 6% of the company** (313 points)

- Htmx 4.0 shipped **as a Game Boy cartridge** (387) - the first JS library to release that way
- "Kill The Cookie Banner" - top of the day on HN, **873 points**: a European campaign to abolish cookie banners
- Terence Tao posted his "Mathematics in the age of AI" slides from ICM 2026 (113)

- PGSimCity - PostgreSQL explained as SimCity (258), and an ESP32 desktop aircraft radar (260)

