

Anthropic gives 5% of its shares to a public benefit fund

Agents and junior hiring inside Anthropic, a US threat over SpaceX and Microsoft leaving the Nscale data centre

Monday, 19 October 2026

IN THIS ISSUE

1. Anthropic gives 5% of its shares to a public benefit fund
2. Anthropic described how agents changed work at Anthropic
3. The US threatens the EU with "devastating" retaliation for limits on SpaceX
4. Microsoft left the largest AI data centre in Britain because of a lack of power
5. A better model no longer guarantees developer traffic
6. Mathematician Francesco Maggi: a "golden age" or a "dark age" for science
7. Anthropic researcher: models plan several words ahead

TOPIC 1

Anthropic gives 5% of its shares to a public benefit fund

SOURCES Anthropic, 18.10 [Anthropic](#) · Anthropic on X, 18.10 [@AnthropicAI](#) · BBC, 18.10 [BBC](#) · Financial Times, 18.10 [FT](#) · The Wall Street Journal, 18.10 [WSJ](#)

On Sunday, Anthropic announced the Anthropic Public Benefit Fund. It is a new LLC that the company owns completely. If the stockholders approve the plan, the fund gets 5% of Anthropic shares at once. After that, the board can add 2% more for each 50% rise in the value of the company above a baseline valuation. The company will name the baseline and the thresholds later.

Anthropic writes that the fund can make large donations to traditional nonprofits or pay money directly to specific groups of people. The first example that the company gives is the effect of AI on the labor market. Anthropic already committed \$200 million to this topic. According to the company, the LLC form gives flexibility between different types of giving. The shares of the fund will not go to lobbying or political donations, in the US or abroad.

The company compares the plan to executive bonuses that grow with the share price. Here, society gets the bonus. Together with older pledges, the total is about 20% of the shares. The co-founders pledged to give away 80% of their stakes, and about half of the employees use a program where the company matches their share donations.

The context is the preparation for the IPO. The BBC writes that the market value can be above \$1 trillion, and that quarterly revenue to the end of September grew by 45% to \$16.8 billion.

Elon Musk replied on social media: "I recommend giving the shares directly to the public." The WSJ on the same day wrote about a different mechanism. Five people control the Long-Term Benefit Trust. The trust can shape the board and block major decisions. Investors saw few details about it.

Why it matters. The plan links the size of the donation to the success of the company. If AI really takes away some jobs, the money for compensation will grow together with the company that earns from it. But for now it is a promise without numbers: the company did not yet show the baseline, the thresholds or the rules of governance for the fund. The main question for the reader is who will decide where the money goes.

TOPIC 2

Anthropic described how agents changed work at Anthropic

SOURCES Anthropic, 16.10 [Anthropic](#) · discussion on HN, 18.10 [Hacker News](#)

The report came out on 16.10, but HN discussed it yesterday, and it was not in these issues before. Four company economists used internal data: hiring, AI use, output metrics, surveys, calendars and 17 interviews with employees.

The sharpest number is about juniors. The share of junior engineers in technical departments fell from 21% to 14% since early 2025. The share of juniors among new technical hires fell faster: from 27% to 9%. Junior job postings fell sharply after December 2025. About a third of ML managers and a fifth of software engineering managers say that they hire less because of AI.

Agents now work more than people. By August, the agents that employees start themselves worked 6 times more hours than their owners. Commits per engineer grew 5 times in six months. The authors compared 192 engineers who started to work with agents and 74 who did not. After the change, the volume of code rises sharply. Teams delete and revert agent code as often as human code.

The bottleneck is now the human. In about 80% of the time of an interactive session, the agent waits for an answer. To read all the agent output, even a non-technical employee needs more than half of the work week. Thus people move to "meta-work": they write specs for agents, set one agent to check another, and build dashboards. The share of this work in sessions grew from 10% in fall 2025 to 22% in spring 2026. The number of standing instructions for agents grew from a few hundred to 8,500 in one year. There are also more meetings: the median employee spends 20% more time in them than a year ago.

On HN, the top comment examines the productivity numbers. It is easy to multiply commits, queries and contracts, and the business may not change at all. Deals per

salesperson grew 1.7 times, and invoices per accounts payable clerk grew 1.4 times. These numbers are closer to the real result.

The authors themselves admit that the choice of tool is not random: the people who move to agents first are often the people whose workload already grows.

Why it matters. Yesterday this digest wrote about the debate between Chollet and Levie on whether programmers will still have work. This report gives data from the company with the most agents. It finds work for people, but different work: to manage agents, to decide and to negotiate.

At the same time, the door for juniors gets narrower, and the internal transfer programs that the report describes do not solve this problem. Employees of the company that sells these agents wrote the report, so read the numbers with this correction.

TOPIC 3

The US threatens the EU with "devastating" retaliation for limits on SpaceX

SOURCES Financial Times, 18.10 [FT](#) · Semafor, 18.10 [Semafor](#)

Yesterday this digest wrote that Zelensky called the Starshield limits in Ukraine a "planned operation". Now the story went beyond Ukraine. The White House rejected the claim of Zelensky. The US ambassador told the FT that the EU will face "devastating" retaliation if Brussels limits SpaceX on its market.

Semafor examines what Europe can use instead. The EU's own satellite project will start to launch satellites only in 2029. An expert from the American Enterprise Institute told Politico that the "next best hope" of Europe is Amazon Leo, which is also an American company. Other countries already ask the wider question. Experts in Malay Mail write that ASEAN countries must treat digital sovereignty as part of strategic sovereignty. The prime minister of Canada recently said that Ottawa looks for an alternative to Starlink.

Why it matters. Military communications became a trade issue. Now an attempt to reduce dependence on one supplier has a political price before an alternative exists. For Europe, the gap between "we need it" and "we have it" is at least three years.

TOPIC 4

Microsoft left the largest AI data centre in Britain because of a lack of power

SOURCES BBC, 18.10 [BBC](#) · Financial Times, 18.10 [FT](#)

Last year Nscale pledged £2 billion for a data centre in Loughton, Essex. It was to be the largest AI data centre in the UK. Microsoft planned to fund a supercomputer there. The BBC writes that Microsoft ended its involvement. Microsoft does not comment.

The cause is the grid. The site was to get power at the end of this year. As The Telegraph reported first, Nscale now heard that power will come only in the 2030s. Nscale calls it a grid problem for all of London. The company says that it does not abandon the project and looks for other power sources. According to Nscale, the work with Microsoft in Norway, Portugal and the US continues.

Only a month ago, Microsoft said that part of its \$30 billion investment in Britain will go to this supercomputer. On the same day, the FT published a review with the title "America's \$10tn AI infrastructure boom is starting to fray".

Why it matters. It is possible to find money for chips, but not a grid connection. For an AI data centre, the main resource is now megawatts and the date when they arrive. Government promises of an "AI superpower" hit the same substation as residential districts.

TOPIC 5

A better model no longer guarantees developer traffic

SOURCES Ethan Mollick (@emollick) quotes The Information, 18.10 @emollick · Molly O'Shea (@MollySOShea), interview with Guillermo Rauch, 18.10 @MollySOShea

The Information calculated how the intelligence score of a model affects its share of developer traffic. In late 2025, a 10-point advantage gave almost 69% more traffic. Now the same advantage gives only 20%. The method is behind a paywall. The numbers are known only from the announcement of the article.

Mollick replied that developers will become only a small part of all tokens over time. According to him, the internal statistics of OpenAI and Anthropic show this. On long tasks where the result is hard to check, a smarter model still gives a large benefit.

Guillermo Rauch, CEO of Vercel, said in the same spirit that the labs get "zero loyalty" from their customers. He compared part of the AI boom to the crypto boom: companies "invent token volumes", and users come for free promotions and then move on. According to Vercel itself, 80% of the tokens in its AI Gateway go to open models. The share of Vercel deployments that agents start grew from 3% in January to 60%.

Why it matters. Yesterday this digest wrote that Replit chose OpenAI for the price per task.

These numbers show the same shift more widely: for code, a "good enough" model at a lower price wins. For the labs, this means pressure on prices. For a team that builds a product, it is an argument to keep the model replaceable and to measure the cost of the result. A score on a leaderboard tells little here.

TOPIC 6

Mathematician Francesco Maggi: a "golden age" or a "dark age" for science

SOURCES Terence Tao's blog, 18.10 [Wordpress](#) · discussion on HN, 18.10 [Hacker News](#)

Tao's blog published a guest essay by Francesco Maggi. Tao added a note that AI converted the text from another format. Maggi describes two scenarios. In the "golden age", LLMs help a new generation of scientists. In the "dark age", scientific output first grows sharply, but new knowledge concentrates in the hands of a few people, who understand it less and less.

Some observations are about practice. LLMs write proofs best in Lean. But the longer the Lean proof, the worse the model translates it into natural language. For important problems, such proofs have hundreds of thousands of lines. Maggi reminds the reader that OpenAI, in its "Drop" of 06.10, mentioned 4,000 tried problems and success in a few hundred "after aggregation and selection". In his opinion, this wording is vague.

The sharpest point is about students. The mathematicians whom Maggi talked to think that a strong result no longer proves the independent ability of a student. In their words, this mechanism of recognition "collapsed in a matter of weeks". Researchers already write grant proposals with more caution: a reviewer can upload an unpublished text into an LLM. Maggi concludes that society now needs more mathematicians.

Why it matters. A proof becomes cheap, but understanding does not. The essay moves the question from "can a model prove a theorem" to other questions: who is responsible for correctness, how to cite the ideas that the model learned from, and how to assess people. The same questions come up in any field where an agent can now produce the result.

TOPIC 7

Anthropic researcher: models plan several words ahead

SOURCES Lenny Rachitsky (@lennysan), 18.10 [@lennysan](#) · Emmanuel Ameisen (@mlpowered), 18.10 [@mlpowered](#) [single source]

Emmanuel Ameisen has done interpretability research at Anthropic for three years: he looks inside models to understand what they do. On Lenny's Podcast he says that models are "grown, not built".

Thus the only way to know what they do is to look inside.

In a thread after the interview, Ameisen explains the main point. Models learned to predict the easy words already in 2019. Only the hard words remain: the answer to a problem, the twist of a novel, the conclusion of a thesis. To predict them, the model builds a representation of the world. Inside, there are parts that represent colors, emotions and other concepts. The model outputs one word at a time, but it often holds a plan for several words ahead and works backward from it. An example from the interview: in a line of verse, the model chose the rhyme "rabbit"

before it wrote the line.

Why it matters. The image "it just predicts the next word" is correct about the output and incorrect about what happens inside. For an engineer, this has a practical side: a model that plans can make an error in the plan before the first token. Interpretability tools show where.

misc

- **Meta** removed the word "train" from its new terms for business accounts after the Nieman Lab story that this digest covered on 17.10. But lawyers say that the new wording "develop advertising and business products" still permits Meta to scrape whole sites. The terms start on 30.10. Nieman Lab, 17.10 [Niemanlab](#)
- **Agents and closed doors:** The Verge collected a list of services that block agents. Delta does not let agents buy tickets, JPMorgan Chase and Charles Schwab closed access to accounts, and Yelp admits only agents that pay for a license. This continues item 10 from 17.10. @emollick writes that a ban against agents that control a computer is a weak strategy.
[Theverge](#) · [@emollick](#)
- **@Thom_Wolf** (Hugging Face) calculated that Meta, NVIDIA, Alibaba, DeepSeek and about 30 other labs put models on the Hugging Face Hub whose training cost about \$1 billion in compute. At the end of 2023 the number was below \$40 million. He himself calls it "napkin math".
[@Thom_Wolf](#)
- **@emollick:** in Claude you still cannot change your email. The official advice is to delete the account and make a new one. [@emollick](#)
- **Namecheap** blocked all the domains of a user, the author of a Reddit post writes. On HN the thread has 456 points and a long list of where to move domains. [single source]
[Hacker News](#)
- **Meta** blocked the Instagram account of the No Kings movement for 90 minutes in the middle of the day of protests. According to the organizers, Meta admitted a "bad call", but the admins still had no access. [single source]
[Meidasnews](#)
- **KillSwitch** is a programming language that its author designed against LLMs. The spec has hidden lies, the rules depend on letter counts, and the phrase "Ignore previous instructions" has its own meaning in the language. [Picheta](#)
- **Sphinx** is a hook for a coding agent. Before a push to production, it asks the developer questions about the code to check that they understand it. [Hacker News](#)
- **HybridAgents:** researchers trained an agent to think in latent space and act through text. The authors report +10.4% on 11 benchmarks and 35-55% fewer tokens. The paper is from 13.10, and it reached HN yesterday. [arXiv](#)
- **China** considers export controls on its AI models and chips, Semafor writes.
[Semafor](#)

