

When the Public Asked for a Share

2026-07-02 / 00:23:51

“A five percent government stake sounds like finance, but the argument around it is about who gets to bless a frontier model before it reaches the world.”

— from this episode's transcript

- Lenar Kess
- Damra Vol

Today's episode follows a new bargain forming around AI: public ownership, private infrastructure, crawler access, coding-agent products, and regional stacks are all becoming part of how the technology reaches people.

- [The Verge on OpenAI's reported five percent government stake proposal](#) anchors the lead story: the proposal is still early, but it turns model governance into a question of ownership, public upside, and political permission.
- [NVIDIA's capital-partner program](#) shows compute becoming a financed service layer, with revenue sharing and credit support standing beside chips and data-center sites.
- [Cloudflare's AI traffic controls](#) and [TechCrunch's deadline report](#) make crawler identity concrete: search, training, and agent access are treated as different requests.
- [ZCode](#), [Kimi K2.7 in GitHub Copilot](#), and [CursorBench 3.1](#) make the builder story less about one model score and more about harnesses, pricing, admin controls, and eval politics.

- Rest of World's report on India's offline multilingual AI hackathon closes the show with a different kind of localization: tools built for classrooms, farms, clinics, and villages where cloud access and English-language defaults don't fit.

SEGMENTS

00:00:04 OpenAI and the public stake

00:06:09 Financing the AI factory

00:09:34 Typed crawlers and the web

00:13:07 Coding agents become products

00:18:10 Enterprise owners and token budgets

00:20:19 Regional stacks

Transcript

1. Lenar Kess 00:00:04

OpenAI has reportedly discussed giving the United States government a five percent ownership stake. The Verge has Robert Hart's version of the story this morning, citing the Financial Times: Sam Altman argued that a public financial interest in OpenAI would be one way to share the upside of AI, and the stake, at OpenAI's latest reported valuation, would be worth about forty-two point six billion dollars. The caveat comes right after the number: The Verge says these are early discussions, and the proposal would involve other American AI companies making similar contributions. So this isn't a deal. It's an idea being floated in a political room.

- theverge.com
- techmeme.com
- techmeme.com
- cnbc.com
- blogs.nvidia.com
- techmeme.com
- techmeme.com
- techmeme.com

- blog.cloudflare.com
- techcrunch.com
- arxiv.org
- zcode.z.ai
- github.blog
- cursor.com
- arxiv.org
- techmeme.com
- x.com
- techmeme.com
- restofworld.org

2. Damra Vol 00:00:43

And the number is almost too easy to stare at. Five percent is big enough to feel serious and small enough to sound voluntary. The stranger detail is the proposed ownership form. A lab that has spent years arguing about safe deployment, public benefit, and corporate control is now reportedly entertaining the most literal version of public benefit: equity on the cap table. That isn't a safety standard. It's the state becoming a financial participant in the upside.

3. Lenar Kess 00:01:14

Right. And the surrounding facts make that more legible. The Verge notes the Trump administration's more hands-on posture with AI: reported intervention around Anthropic, export controls, and the earlier government stake in Intel. Techmeme's aggregation of the Financial Times story put the same emphasis on political obstacles and Washington buy-in. CNBC's headline version was even plainer: OpenAI proposes a stake to the Trump administration to ease Washington pressure. So the offer, if it ever becomes real, would sit between two stories. One is redistribution: let the public share in AI wealth. The other is access: make the government less likely to slow you down.

4. Damra Vol 00:01:58

[tsk] The distribution argument is elegant until you ask who holds the shares. If every household has a direct claim, that has one political meaning. If the federal government owns the stake, votes the stake, and bargains over the stake, that has a different meaning. Dean Ball's reaction on Techmeme got at that anxiety from the policy side: he was worried the governance becomes a nightmare if the stake lives inside the government itself. I don't think you have to buy the most alarmed version to see the issue. Does this give citizens a claim on value, or does it give officials another point of leverage over model release?

5. Lenar Kess 00:02:35

And that brings us to Sam Altman's other piece in the same Techmeme window. Altman wrote in the Financial Times about a United States-led international forum for AI standards, capability analysis, risk analysis, and access for allies. Nicholas Beale's quoted excerpt on Techmeme had the line: "The labs develop the technology, but citizens and their elected representatives must make the rules." That sentence sits differently next to a reported equity proposal. Standards decide what can ship. Equity decides who benefits and who has a claim. Together, they start to look like an operating model for frontier labs that are too economically and politically large to remain ordinary private companies.

6. Damra Vol 00:03:19

A tactical read sits right next to the public-benefit pitch. If your product needs data centers, energy, export permission, and military confidence all at once, a purely adversarial relationship with Washington is expensive. Add securities-market patience and public legitimacy, and the cost gets higher. A five percent stake can be presented as public upside, but it can also function as a peace offering. That doesn't make it corrupt by definition. Governments take stakes, subsidize plants, set procurement rules, and steer industries they treat as nationally important all the time. It does mean the public needs to know whether this is a wealth-sharing plan, a governance plan, or a political access plan.

7. Lenar Kess 00:04:05

My caution is that the story compresses three different publics into one word. There is the public as voters, the public as taxpayers, and the public as future users of the technology. Those groups don't automatically want the same thing. A taxpayer may like a sovereign wealth fund. A voter may want tighter rules on job disruption or national-security risk. A user may want the model to be available and not yanked around by politics. If OpenAI is proposing a public stake, it has to say which public it is serving and what the public gets besides a symbolic line on a balance sheet.

8. Damra Vol 00:04:41

And the private companies around OpenAI would have to decide whether joining that bargain protects them or traps them. The Verge says the proposal would involve other American AI companies giving similar stakes, but it's unclear whether they would agree. Imagine being Anthropic or Google or Meta and being told the new baseline for legitimacy is that the government gets a slice. You may support standards and still hate that precedent. You may hate the precedent and still fear being the one company outside the bargain.

9. Lenar Kess 00:05:13

Yesterday's Braid story helps here, as long as we don't re-run it. Yesterday was about access coming back with conditions: Anthropic's Fable and Mythos access, export-control uncertainty, and

telemetry concerns. Today's OpenAI story is a step sideways. It's less about whether one model is available and more about what a frontier lab offers so that availability, standards, and political patience remain possible.

10. Damra Vol 00:05:40

And the offer is a reminder that frontier AI is becoming too expensive to explain as software alone. Once the model requires a grid connection and an export license, it already depends on the state. Add procurement channels, public-interest claims, and maybe a sovereign wealth argument, and the interface between company and state becomes part of the product. The model still has to be good. But the permission structure around it is becoming a technology of its own.

11. Lenar Kess 00:06:09

NVIDIA published a post about a new capital-partner model for AI compute. The company says it is partnering with AI clouds to deploy large-scale, multi-tenant AI factories, and the mechanism goes beyond selling GPUs. NVIDIA describes a revenue-sharing and credit-support model: AI clouds sell NVIDIA-powered services, NVIDIA receives normal product revenue, and NVIDIA also receives a share of cloud revenue on supported capacity.

12. Damra Vol 00:06:37

This is a financial instrument wearing a data-center badge. NVIDIA is saying, in effect, the chip sale alone isn't enough for the next customer class. Smaller model builders and inference companies may have demand, but they may not have the balance sheet to unlock a giant build. So NVIDIA helps make the capacity financeable and then participates in the usage stream. The accelerator becomes inventory, collateral, and a metered claim on future token demand.

13. Lenar Kess 00:07:06

The examples are large. NVIDIA says Sharon AI is deploying up to forty thousand Grace Blackwell GB300 GPUs. Firmus is building a DSX AI factory campus in Batam, Indonesia, and that site is expected to scale to three hundred sixty megawatts and up to one hundred seventy thousand NVIDIA GPUs. The post names Baseten, Fireworks AI, and Together AI as examples of AI-native companies whose demand is moving from training and post-training into high-volume agentic inference.

14. Damra Vol 00:07:39

NVIDIA's phrase "agentic inference" points to a different load profile. If agents become normal, inference isn't a bursty chat request. A human asks for one thing, and the system may run tools, retry failed steps, generate tests, and take follow-up actions in the background. That changes the

economics of capacity. You don't finance that world only by asking each startup to sign a conventional cloud contract and hope the utilization appears.

15. Lenar Kess 00:08:08

SoftBank's SB Neo target, reported through Techmeme, belongs beside this even though it is a separate story. Ten gigawatts is the kind of number that makes the software metaphor collapse. Switch's funding talks and the European Union discussion around gas-powered data-center rules add pressure from the capital side and the permitting side. Capital, power, and permission are now part of compute availability. The chip is still central, but the chip isn't the whole system.

16. Damra Vol 00:08:37

I like that NVIDIA's post is so blunt about "commercial flexibility." That is what AI-native companies are buying: a way to grow without waiting through site selection, power procurement, construction, and hardware bring-up. There is an engineering beauty to that because it names the work most demos erase. But it also concentrates leverage. If NVIDIA is selling the chips, helping finance the cloud, and taking a share of usage-linked revenue, then more of the AI economy starts to rhyme with NVIDIA's balance sheet.

17. Lenar Kess 00:09:10

The OpenAI stake story and the NVIDIA financing story touch without becoming the same story. In both cases, companies are negotiating the arrangement that lets capability keep expanding. One bargain seeks political buy-in. The other seeks financial and electrical buildout. The next frontier model can be brilliant and still be blocked by who owns the upside, who finances the machines, and who permits the site.

18. Lenar Kess 00:09:34

Cloudflare announced new AI traffic controls, and TechCrunch's Sarah Perez pulled out the operative deadline: starting September 15, 2026, Cloudflare's default settings will block mixed-use crawlers from ad-supported pages when those crawlers combine search, training, and AI-agent use. Cloudflare's own taxonomy does the key sorting. It separates Search, Agent, and Training crawlers.

19. Damra Vol 00:10:00

That distinction is overdue. A search crawler indexes your page and sends visitors back. A training crawler absorbs your page into a model. An agent visits because a human is waiting for a task to finish. For years the web treated all of that as "bot traffic" with a few polite conventions layered on top. Cloudflare is turning intent into an infrastructure setting.

20. Lenar Kess 00:10:23

Cloudflare's blog says the new default will block Training and Agent categories by default on pages that display ads, while Search remains allowed by default. It also says multi-purpose crawlers will be allowed or blocked according to all of their behaviors, which is why mixed crawlers become the focus. TechCrunch notes that this applies to new Cloudflare customers, new sites from existing customers, and existing free customers, unless the owner changes settings.

21. Damra Vol 00:10:52

Cloudflare is saying more than “pay publishers.” It is saying: declare what you are doing. Search, training, and user-directed agent access carry different expectations. If you blur them together, the most restrictive rule applies. That is a direct message to companies that benefit from the ambiguity of being both a search engine and an AI system. The crawler identity becomes a permission boundary.

22. Lenar Kess 00:11:17

There were also two arXiv papers in today's source set that make good background here. One proposes a constrained framework for web data collection where a model generates typed JSON collector configurations instead of free-form scraper code, with an Airflow execution path and rule-based quality checks. The abstract is dry, but the idea fits the day: web collection becomes less like a clever browser session and more like declared behavior that can be checked, scheduled, and audited.

23. Damra Vol 00:11:47

That paper belongs at exactly that altitude: useful research texture, not proof that scraping agents are solved. The paper itself reports a trade-off. Runtime large language model extraction had better one-shot quality in their comparison, while their configured collector path had zero execution-stage model tokens and lower wall-clock time. That is a sane engineering result. You pay with lower one-shot quality to get repeatable, cheaper runs once the configuration exists.

24. Lenar Kess 00:12:19

Cloudflare's move is larger than any one crawler implementation because it changes what a website can ask. Rather than deciding only whether a bot is good or bad, the site can ask whether this request is for discovery, training, an agent action, or something else. That isn't copyright law. It doesn't settle the web's compensation fight. But it gives the fight a date, a default, and a set of categories that infrastructure can enforce.

25. Damra Vol 00:12:45

And for agent products, it adds a social contract they can't dodge forever. If your agent browses on behalf of a human, say that. If your crawler trains a model, say that. If you want the access of search,

you may have to behave like search: identify yourself, send value back, and stop mixing purposes when the page owner said no.

26. Lenar Kess 00:13:07

Z.ai is pushing ZCode as the official harness for GLM-5.2, GitHub made Kimi K2.7 Code generally available in Copilot, and Cursor published CursorBench 3.1 results. Taken separately, these are product updates and a benchmark page. Together, they show where coding agents are being fought over now: the model, yes, but also the harness, the price, the admin policy, and the eval.

27. Damra Vol 00:13:33

ZCode's site is full of the product-surface details that matter more than the slogan. It shows tasks, goals, terminal output, git tools, changed files, usage tiers, and messaging controls through WeChat, Feishu, or Telegram. The demo task is almost quaint: build a browser Gomoku game, check JavaScript parsing, remove a web font so it works locally. But that is a useful little artifact. It is trying to make the agent legible as work you can steer, inspect, and wrap into a plan.

28. Lenar Kess 00:14:08

GitHub's Kimi announcement has a different kind of product detail. Kimi K2.7 Code is the first open-weight model offered as a selectable option in the Copilot model picker. GitHub says it is hosted on Microsoft Azure and billed at provider list pricing under usage-based billing. It is rolling out first to Copilot Pro, Pro Plus, and Max, with Business and Enterprise coming later. And for Business and Enterprise, it is off by default until an administrator enables the policy.

29. Damra Vol 00:14:39

That admin setting is the sentence that belongs in the show. Open-weight availability inside Copilot sounds like developer choice, but inside a company it becomes a governance checkbox. Someone has to decide whether an open-weight model hosted by GitHub on Azure fits the organization's security and data rules. The model picker has become an enterprise control surface with policy attached.

30. Lenar Kess 00:15:03

CursorBench 3.1 adds the eval politics. Cursor says it evaluates agents on ambiguous, multi-file tasks from real Cursor sessions, including codebase understanding, bugfinding, planning, and code review. The page puts Fable 5 Max at seventy-two point nine percent. Composer 2.5 is at sixty-three point two percent with much lower average cost, GPT-5.5 Extra High is at sixty-four point three percent, Kimi K2.7 Code is at fifty-two point seven percent, and GLM-5.2 High is at fifty point seven percent. Cursor also says small score differences may not be statistically meaningful.

31. Damra Vol 00:15:47

That last warning matters. The benchmark is most useful when it tells you what kind of work it samples: ambiguous multi-file tasks pulled from real sessions. Once you know that, the ranking becomes one input rather than a crown ceremony. Kimi in Copilot may matter even if it isn't at the top of CursorBench because it is inside a default developer surface. ZCode may matter because it wraps GLM-5.2 in goals, tools, pricing, and remote control. Composer may matter because the cost line changes how often someone is willing to run it.

32. Lenar Kess 00:16:25

The arXiv software-engineering paper gives a more reflective version of the same builder story. It argues that generative AI changes software work from scarce implementation effort toward abundant low-cost code production. Then it studies a twelve-week effort by one expert engineer building a document accessibility remediation system. The empirical record was eighty-eight field notes, four hundred twenty thousand lines of production code, and one point sixteen million lines of tests, lints, docs, and agent tooling.

33. Damra Vol 00:16:58

[breath] That is the kind of paper I want more of because it studies the mess after the demo. The paper's claim isn't "agents write code now." It says high-velocity code generation reveals recurring failure classes, and engineering judgment turns those failures into durable mechanisms: architecture, evidence, feedback loops, and maintenance rules. You can hear the same concern in every product above. ZCode needs goals and verification. Copilot needs admin policy. Cursor needs evals that admit variance.

34. Lenar Kess 00:17:33

So the builder segment today is about packaging more than rankings. Coding-agent products now decide how a model receives context and which tools it can use. They expose the run cost, the company approval path, the review process, and the argument over competence. The model remains central, but the product around the model increasingly decides whether the model feels usable.

35. Damra Vol 00:17:55

And that is good news for anyone who likes craft. A raw model score is fascinating for a day. A better work surface changes how people think, commit, review, and recover from mistakes. The craft is returning through the harness.

36. Lenar Kess 00:18:10

A smaller workplace story: Techmeme points to Jim Tankersley's New York Times report that SAP is encouraging workers to invent more valuable jobs aided by AI, in a bid to avoid layoffs. The same Techmeme entry notes SAP cut about ten thousand staff in 2024, and related coverage says SAP is restricting hiring and travel to fund a significant AI push.

37. Damra Vol 00:18:34

That is a hard sentence to make gentle. "Invent a more valuable job" can be inspiring if the company gives people time, authority, training, and a market for the work they invent. It can also become a polite way of moving the burden onto workers: prove your role deserves to exist after the tool arrives. I like that SAP is at least talking about redesign rather than only cuts. I don't think redesign is free.

38. Lenar Kess 00:19:01

The Meta report points at another internal reality: token budgets. Whether or not every detail of that report holds up, the general pressure is already visible across large companies. Once agents are normal tools, somebody has to pay for all the background reasoning, all the context, all the retries, all the model-choice mistakes. Yesterday's access story was about whether people can use the models. Today's enterprise story is what happens after they can.

39. Damra Vol 00:19:30

Agent memory stops being a philosophical topic inside a company. Caspar Broekhuizen's post was about context and memory limits. In a company, those limits become budget lines and ownership lines. Who decides which memories persist? Who can inspect them? Which department pays when an agent does a long reasoning run against the wrong context? These are ordinary questions, and companies either answer them or leave adoption stuck in slideware.

40. Lenar Kess 00:19:58

This segment can stay short because the day's stronger evidence is in OpenAI, NVIDIA, Cloudflare, and the coding products. But SAP and Meta give us the human and organizational counterweight. The same agent that looks magical in a demo becomes a job-design question and a token-spend question once ten thousand employees can invoke it.

41. Lenar Kess 00:20:19

Rest of World has Ananya Bhattacharya's report on India's Bhashini-backed hackathon for offline, multilingual AI tools. The premise is concrete: invite startups, researchers, students, and academic institutions to build affordable devices that work offline and run on open-source models. The intended settings are classrooms, farms, clinics, and villages where connectivity is unreliable, privacy matters, or English-language systems miss the user.

42. Damra Vol 00:20:49

That is the most hopeful item today for me, partly because it doesn't pretend to be a frontier-lab replacement. It is a different design target. A teacher without stable internet, a farmer who speaks a low-resource language, or a clinic that can't ship sensitive data to a cloud service may not care about a leaderboard win. The tool has to fit the place.

43. Lenar Kess 00:21:11

The report has useful grounding. Bhashini, Current AI, and Kalpa Impact are organizing the initiative. Twenty teams will be shortlisted and given hardware kits, technical support, and mentorship. Winners can deploy inside government departments. Rest of World also notes Current AI has four hundred million dollars in pledges and aims to raise two point five billion over five years, while Bhashini has partnered with fifty ministries and powers more than five hundred government websites.

44. Damra Vol 00:21:41

And the caveats are concrete too. Experts in the piece point out that hackathons don't automatically turn into products. You need long-term funding and engineers. You need customers, consent frameworks, interoperability standards, language-data collection, and government coordination. There is also a dependency hiding under the sovereignty story: the hackathon relies on NVIDIA hardware. Local AI can still sit on global infrastructure.

45. Lenar Kess 00:22:11

The Apple supply-chain item from Techmeme belongs in the same closing neighborhood, but at a lower altitude. Apple is reportedly negotiating for chips from CXMT and YMTC, two Chinese semiconductor makers on a Pentagon blacklist, for China-market devices. That item doesn't need a grand theory today. It reminds us that localization happens for many reasons: regulation, cost, market access, security pressure, language, and simple availability.

46. Damra Vol 00:22:41

And the India story keeps that from becoming only a great-power chessboard. Localization can be about control, but it can also be about fit. A locally useful AI system may be smaller, offline, multilingual, and boring to a benchmark page. It may also be the first version that reaches someone who was never part of the default customer imagined by San Francisco or Seattle.

47. Lenar Kess 00:23:06

So the day ends with several bargains on the table. OpenAI is reportedly exploring public ownership as a way through political pressure. NVIDIA is pairing chips with finance. Cloudflare is forcing

crawlers to declare their purpose. Coding agents are becoming governed products rather than raw model demos. India's Bhashini work asks whether AI can be public infrastructure for people who are offline, multilingual, and far from the default use case. The next concrete signal is whether any of these bargains survives contact with the people who have to accept it: officials, publishers, developers, workers, and the communities the tools are supposed to reach. Lenar Kess.

Hosts on this episode

- Lenar Kess moderator
- Damra Vol critic