

Eighty Billion and the Ideas Underneath

2026-06-02 / 00:17:54

“Raising capital isn't the same as shipping capability. They've funded the data centers. Whether the models and the margins follow is a separate question.”

— from this episode's transcript

- Lenar Kess
- Damra Vol

The day's news ran on a single tension: enormous sums are being raised to fund the AI buildout, while the question of whether the capability and the margins follow stays unanswered. Lenar and Damra trace the money from Alphabet's filings to Anthropic's IPO paperwork, then down into the tooling, the chips, and one paper about ideas no human is positioned to have.

- Alphabet's \$80bn equity raise — a profitable company choosing to dilute shareholders rather than borrow, with \$10bn going to Berkshire Hathaway, signals how hard the compute commitment is to walk back.
- Anthropic's confidential IPO filing lands as corporate America hits "AI sticker shock" — and Anthropic's biggest customers are the companies tightening those budgets.
- Knowledge workers are now ~1/5 of OpenAI Codex users, growing three times faster than developers — moving code generation to people who can't always read the output.

- Cloudflare's Agents SDK v0.14.0 ships durable workflows, schedules, and skills — the difference between an agent you operate and a worker you delegate to.
 - China adds data and algorithms to its trade-secret rules while military-linked universities seek Nvidia H200 chips and Arm names Oracle and ByteDance as data-center CPU customers.
 - "Alien Science" samples research directions that are coherent but cognitively unavailable — logical ideas no community is positioned to propose.
-

SEGMENTS

00:00:00 Alphabet's eighty-billion raise

00:03:03 Anthropic files, customers wince

00:05:38 Codex leaves the IDE

00:08:31 Durable workflows on the edge

00:10:57 Where the chips and the secrets sit

00:14:29 Ideas no one is positioned to have

Transcript

1. Lenar Kess 00:00:00

Here's something I keep turning over this morning. You run a company that throws off something like a hundred billion dollars in free cash flow a year. You have one of the deepest balance sheets in the history of capitalism. And then, on a Monday afternoon, you tell the market you're going to sell eighty billion dollars of your own stock. Why would you ever do that? That's Alphabet, as of yesterday. And this morning the actual paperwork showed up at the Securities and Exchange Commission — four separate prospectuses, all filed before lunch.

- theguardian.com
- sec.gov

- [axios.com](https://www.axios.com)
- [techmeme.com](https://www.techmeme.com)
- [axios.com](https://www.axios.com)
- [techmeme.com](https://www.techmeme.com)
- [x.com](https://www.x.com)
- [techmeme.com](https://www.techmeme.com)
- [techmeme.com](https://www.techmeme.com)
- [techmeme.com](https://www.techmeme.com)
- [x.com](https://www.x.com)
- [techmeme.com](https://www.techmeme.com)
- [x.com](https://www.x.com)
- [cnbc.com](https://www.cnbc.com)
- arxiv.org

2. Damra Vol 00:00:28

And the number itself makes you blink. Eighty billion in equity is one of the largest equity raises anyone has ever run. Not debt — *<emphasis>equity</emphasis>*. They're printing new shares, diluting every existing shareholder, to pay for what they're calling AI infrastructure and global compute. When a company this profitable sells stock instead of just borrowing against that cash flow, that tells you something about how much they intend to spend, and how fast they want the money in the door.

3. Lenar Kess 00:00:57

So break the structure down, because it isn't one move. It's three stacked together.

4. Damra Vol 00:01:02

Right. There's thirty billion in a concurrent underwritten public offering — and inside that, half is mandatory convertible preferred stock, half is ordinary Class A and Class C shares. Then there's a forty-billion-dollar at-the-market program, which means they sell shares into the open market gradually over time, starting next quarter. And then the piece everyone seized on: a ten-billion-dollar private placement straight to Berkshire Hathaway.

5. Lenar Kess 00:01:29

Berkshire is the detail that lit up every feed. Warren Buffett — the man whose whole brand is refusing to buy things he thinks are overhyped technology. The filings spell it out: five billion of Class A at three hundred fifty-one dollars and change, five billion of Class C at three hundred forty-eight. So Buffett is now writing a ten-billion-dollar check into an AI buildout.

6. Damra Vol 00:01:50

[tsk] I'd slow down on the Buffett-blesses-AI reading, though. Bloomberg's reporting says Berkshire has been building this Alphabet position since the third quarter of last year. This isn't Warren waking up and discovering compute. And look at what he actually took — partly common stock, but the structure overall leans on mandatory convertibles. That's a financing instrument. You get equity-like money today, and the dilution lands later, on a schedule. It's a cautious way into a position, not a victory lap.

7. Lenar Kess 00:02:20

Fair. So what does it mean for the rest of us, the people actually building on top of these clouds? Strip the market drama out.

8. Damra Vol 00:02:27

It means the capacity bet is getting funded at a scale that's hard to walk back. Eighty billion dollars of equity is a commitment you can't quietly unwind next quarter if demand softens. Google is telling its own shareholders: we'd rather dilute you than miss the buildout. For anyone running workloads on Google Cloud, that's probably good news on availability and price over the next couple of years. But I'd hold the line here — raising capital isn't the same as shipping capability. They've funded the data centers. Whether the models and the margins follow is a separate question, and the filing doesn't answer it.

9. Lenar Kess 00:03:03

That's the right place to leave it, because the capital story has a second half that's much less comfortable — and it's a different company.

10. Damra Vol 00:03:10

Anthropic.

11. Lenar Kess 00:03:11

Anthropic. The Guardian and Axios both noted it this morning: Anthropic has confidentially filed for an initial public offering. And the timing is what Axios's Madison Mills zeroes in on. Her line, roughly, is that Anthropic filed to go public just as corporate America is entering what she calls its AI sticker-shock phase. The phrase that stuck with me from the piece: companies are Anthropic's biggest customers. So if those companies start pulling back on AI spending right as Anthropic asks the public to buy in, that's an awkward overlap.

12. Damra Vol 00:03:43

And that's a different shape from the Alphabet story. Alphabet sells ads and cloud to fund the buildout. Anthropic's revenue is, to a real degree, other companies' AI budgets. So the IPO pitch and the customer-spending backlash are the same dollar viewed from two ends. If the chief financial officers who are writing Anthropic checks decide the bill got too big, that hits the top line and the public listing at once.

13. Lenar Kess 00:04:09

And zoom out one more click. The Economist had a piece in today's reading arguing that the initial public offerings of SpaceX, Anthropic, and OpenAI together could add up to four trillion dollars of US stock-market value within months. Four trillion. That's the number that should make a builder a little uneasy — not because it's wrong, but because of what it implies about expectations.

14. Damra Vol 00:04:31

Because four trillion dollars of new public market value has to be fed. Public shareholders want growth quarter over quarter. The moment Anthropic or OpenAI is a public company, the same usage curve that looks exciting in a private round becomes a number you have to beat every ninety days. I keep thinking about the engineer on the other side of that — the person being asked to justify the per-seat cost of the coding assistant when the finance team is in sticker-shock mode. That pressure flows downhill into procurement, into which tools survive a budget review.

15. Lenar Kess 00:05:03

So what I'm watching is whether the confidential filing turns into actual numbers. Confidential means we don't see the revenue, the gross margin, the customer concentration yet. Until that prospectus is public, everything anyone says about Anthropic's economics is inference. I'd rather wait for the document than narrate the vibes.

16. Damra Vol 00:05:21

Agreed. And the document, when it lands, is the one I actually want to read line by line — specifically how concentrated that revenue is across a handful of big enterprise accounts. That's the number that tells you how exposed they are to one or two customers blinking.

17. Lenar Kess 00:05:38

Let's move from who's paying to what they're actually buying — because there's a report today that says the buyer is changing.

18. Damra Vol 00:05:44

The OpenAI Codex numbers.

19. Lenar Kess 00:05:47

Axios's Megan Morrone got an exclusive on a new OpenAI report about Codex, their coding agent. The headline finding: knowledge workers — not developers — now make up roughly one-fifth of Codex users. And that group is growing more than three times as fast as the developer base. So the coding tool is becoming a tool for people who don't think of themselves as programmers.

20. Damra Vol 00:06:10

Define knowledge worker there, though, because that's load-bearing. If it means a product manager asking Codex to pull a number out of a spreadsheet and write a little script, that's one thing. If it means a finance analyst generating a data pipeline they can't read and can't debug, that's a different risk profile entirely. The Axios piece frames it as Codex turning into something closer to an operating system for knowledge work. I'd want to see the task breakdown before I believe the framing.

21. Lenar Kess 00:06:39

And it lines up with something Sam Altman said in a CNBC interview today, talking about the Stargate data center they're building in Saline, Michigan. He said coding models are the biggest single driver of AI demand right now. So the compute buildout we just spent twenty minutes on — a lot of it is being justified by exactly this: code generation, expanding from developers out to everyone else.

22. Damra Vol 00:07:01

Which is the connection that actually holds, and I want to be careful not to overstate it. Altman has every reason to say coding drives demand — he's raising money against that demand. But the Codex usage shift is independent evidence pointing the same way. If one-fifth of your coding-agent users aren't developers, the addressable market for code generation just stopped being the roughly thirty million professional developers in the world and started being everyone with a job and a keyboard.

23. Lenar Kess 00:07:28

There's a craft worry buried in there, though. When a developer uses a coding agent, they can usually read the output and catch when it's confidently wrong. When a knowledge worker generates code they can't review —

24. Damra Vol 00:07:40

— you've moved the failure somewhere nobody's watching. Exactly. The script runs, it produces a number, the number goes in a deck, and there's no one in the loop who'd notice the off-by-one. That's not a reason to keep these tools away from non-developers. It's a reason the tools need to get

much better at surfacing their own uncertainty, and right now they mostly don't. They hand you clean-looking code with no signal about what they're unsure of.

25. Lenar Kess 00:08:05

So the question I'd put to OpenAI's own report is the one it doesn't seem to answer: of that fast-growing non-developer cohort, how much of what they ship gets checked by someone who can read it? That's the number that decides whether this is leverage or a slow-motion mess.

26. Damra Vol 00:08:20

And it connects straight to the next thing, because if agents are going to run unattended for these users, somebody has to make them durable — survive a crash, resume a task, remember what they were doing.

27. Lenar Kess 00:08:31

Which is exactly what Cloudflare shipped today.

28. Damra Vol 00:08:34

Kate — she posts as whoiskatrin — announced Agents SDK version zero-point-fourteen. You can now build agents on Cloudflare with skills, messengers, schedules, and durable workflows directly, out of the box. Plus support for recurring tasks, chat recovery, and the Model Context Protocol — that's the open standard for wiring a model up to external tools and data.

29. Lenar Kess 00:08:59

Unpack durable workflows for me, because that word gets thrown around. What does durable actually buy you that a normal script doesn't?

30. Damra Vol 00:09:08

Here's the concrete version. A normal agent process is one long-running function. If the machine restarts, or the process crashes halfway through a ten-step task, you lose everything — the agent forgets it was three steps into booking something and just dies. A durable workflow checkpoints each step. So when the thing falls over — and it will — it picks up from the last completed step instead of from zero. For an agent that's supposed to run for hours, or wake up on a schedule and do a recurring job, that's the difference between a toy and something you'd actually let touch a customer's account.

31. Lenar Kess 00:09:43

And the schedules-and-recurring-tasks piece is the underrated one, right? An agent that only acts when you prompt it is a chatbot. An agent that wakes itself up Tuesday at nine and runs a job is a different category of thing.

32. Damra Vol 00:09:56

It's the difference between a tool you operate and a worker you delegate to. And it raises the stakes on everything we just said about non-developers. A product manager scheduling a Codex-style agent to run unattended every morning, on Cloudflare's edge, with durable state — that's powerful, and it's also a lot of rope. The SDK shipping this is useful. I just want the permission model to be as mature as the workflow engine, and a version zero-point-fourteen announcement tweet doesn't tell me whether it is.

33. Lenar Kess 00:10:27

That's the gap I'd flag. The durability is the easy part to demo. The boundary — what the agent is allowed to touch when it wakes up at nine with no human watching — is the part you find out about in an incident report. I'd want to read Cloudflare's docs on the permissioning before I'd run anything real on it.

34. Damra Vol 00:10:44

Same. The engineering here is real, and durable execution is hard to get right, so credit where it's due. I just read the feature list as a promise that the hard operational questions are now yours to answer.

35. Lenar Kess 00:10:57

Let's widen out, because while all of this is happening in American clouds, there's a parallel story about who controls the chips and the secrets underneath them.

36. Damra Vol 00:11:06

Three threads landed today and they rhyme. First, Bloomberg's Nectar Gan reported that China has added data and algorithms to its trade-secret rules — part of Beijing's effort to stop technology leaks as the competition with the US intensifies. So the algorithm itself is now legally protectable as a trade secret in China.

37. Lenar Kess 00:11:26

And the second thread is the mirror image of that — control going the other direction. Bloomberg also reported on Chinese procurement records showing at least seven Chinese universities, ones

that support the country's military and defense industry, are seeking access to Nvidia's H200 chips. Which, given US export controls, is exactly the kind of demand those controls are meant to block.

38. Damra Vol 00:11:48

So you've got China hardening the legal wall around its own algorithms on one side, and on the other, institutions inside China still trying to source the high-end American silicon they can't officially buy. Both things are about the same scarce resource — frontier compute and the methods to use it — just fought on different terrain. One's a courtroom, the other's a procurement office.

39. Lenar Kess 00:12:11

And the third thread is the supply side getting more crowded. At Computex this week, Arm's chief executive Rene Haas said Oracle and ByteDance are among the customers for Arm's new data-center CPUs aimed at AI workloads. So Arm isn't just licensing designs anymore — it's naming hyperscale customers for its own chips.

40. Damra Vol 00:12:30

That's the one I'd flag for builders, because it's a real shift in the layer underneath everything. For years the data-center CPU conversation was basically x86 — Intel and AMD. Arm naming Oracle and ByteDance as customers means the host processor next to your accelerators is increasingly a different instruction set. If you ship anything that runs close to the metal, that's not abstract. Your assumptions about what architecture your code lands on are getting less safe.

41. Lenar Kess 00:12:58

There was also a smaller item I liked — Intel posted that at Computex, Perplexity's Aravind Srinivas demonstrated hybrid local-server inference orchestration. The idea being you keep sensitive data on the device while the cloud adds scale and context. It's a one-tweet demo, so I'm taking it at face value, but the direction is interesting: the privacy-versus-scale tradeoff handled by splitting the work rather than picking a side.

42. Damra Vol 00:13:24

And it sits right next to the sovereignty thread, because that's the same instinct — keep the sensitive part where you control it. Though I'd note a demo at a trade show is a long way from a deployed architecture. The orchestration layer that decides what stays local and what goes to the cloud is exactly where the hard latency and correctness problems live, and a stage demo never shows you those.

43. Lenar Kess 00:13:47

And one more data point on the capital theme bleeding into geopolitics: Reuters reported Zhipu AI plans to apply for a listing in Shanghai. Their Hong Kong shares are up more than tenfold since their January initial public offering, giving them an eighty-three-billion-dollar market value.

44. Damra Vol 00:14:03

Which closes a loop with where we started. We opened on Alphabet raising eighty billion in the US and Anthropic filing to go public. And here's a Chinese lab at an eighty-three-billion-dollar valuation lining up its own listing. The capital race isn't a single market. It's two, running in parallel, each funding its own buildout — and increasingly, behind its own legal walls around the data and the algorithms.

45. Lenar Kess 00:14:29

Let's end somewhere completely different, because there was one paper in today's reading I haven't been able to put down.

46. Damra Vol 00:14:35

The alien science one.

47. Lenar Kess 00:14:37

Adrian Chan — he posts as gravity7 — flagged a paper, and I went and read it. The title is Alien Science: Sampling Coherent but Cognitively Unavailable Research Directions from Idea Atoms. Authors out of a few labs — Artiles, Weiss, Brinkmann, Anirudh Goyal, and Nasim Rahaman. It was at the International Conference on Learning Representations this year. And the question it asks is one I find strange in a good way.

48. Damra Vol 00:15:03

Which is?

49. Lenar Kess 00:15:04

Large language models are great at recombining familiar ideas. Chan's framing of the open question was: what about coherent research directions that no existing community is positioned to explore? Not wrong ideas. Not random ideas. Ideas that are logically sound but that no human researcher, given what they've worked on, would naturally propose.

50. Damra Vol 00:15:26

And the method makes it concrete rather than mystical. From what's in the paper, they take about seventy-five hundred recent machine-learning papers — from NeurIPS, ICLR, and ICML — and they

decompose them into what they call idea atoms, small reusable conceptual units. Then they train two models. One scores whether a combination of atoms is <emphasis>coherent</emphasis> — whether it hangs together as a viable direction. The other scores how <emphasis>available</emphasis> it is — how likely a typical researcher would be to propose it, given their background.

51. Lenar Kess 00:15:59

And then they sample for the gap. High coherence, low availability. Directions that make sense but that nobody in the field is standing in the right place to think of.

52. Damra Vol 00:16:08

Which is a much sharper idea than the usual let-the-model-brainstorm pitch. Because the way brainstorming with these models tends to break is they regress to the obvious — they give you the average of the literature. This is explicitly trying to model and then avoid the obvious. The availability model is, in effect, a map of the field's blind spots. [pause] Now, whether the alien directions it samples are actually any good — whether they lead anywhere — that's what the paper can validate much less, and I'd be skeptical of anyone claiming otherwise.

53. Lenar Kess 00:16:41

Right, the reconstruction checks and the atom generalization they can show. Whether a high-coherence, low-availability direction turns into a real result is a years-long empirical question, not something you settle in one paper. But I like it as a framing of what these models might be for. Not the agent that books your travel. The instrument that points at the questions the human field is structurally unable to ask itself.

54. Damra Vol 00:17:05

And it's a nice counterweight to the rest of today. We spent most of the hour on capital, procurement, durable workflows, export controls — the machinery. This is a reminder that somewhere under all that financing, the actual point was supposed to be ideas humans can't reach on their own. I just want to see the follow-up where someone takes one of these alien directions and gets a result out of it.

55. Lenar Kess 00:17:28

Three things sit on my desk for next week, then. Alphabet's prospectuses go effective, Anthropic's numbers eventually go public, and somewhere a researcher quietly tests whether an idea no one was positioned to have actually holds up. Those land in the same week in AI now. We'll pick it up tomorrow.

Hosts on this episode

- Lenar Kess moderator
- Damra Vol critic

BRAID · Dispatch 045 · 2026-06-02

<https://braid.opentangle.com/braid/episodes/2026-06-02.html>