

When Compute Signed a Utility Contract

2026-06-22 / 00:24:54

“The AI capacity story got less abstract today because Microsoft did not announce a model. It signed for power.”

— from this episode’s transcript

- Lenar Kess
- Damra Vol

Today’s episode follows AI infrastructure into its physical contracts: gas-fired power in West Texas, water claims from Nvidia, local backlash against data centers, export controls, enterprise rollouts, and the strange new place where coding agents meet private equity diligence.

- [The Wall Street Journal on Chevron and Microsoft](#) reported a 20-year gas-fired power agreement for a proposed West Texas AI data center, which turns compute capacity into a named energy project.
- [Axios on Nvidia’s water claim](#) adds the cooling side of the constraint, with Nvidia arguing that next-generation systems can reduce data center water use.
- [Axios on data center backlash](#) and [Gallup’s local-opposition polling](#) show that data centers are becoming the public face of AI anxiety.
- [The Associated Press on China’s restrictions](#) explains Beijing’s new export-control and procurement moves against U.S. defense-linked firms.
- [OpenAI’s Samsung announcement](#) and [Getty Images’ OpenAI display partnership](#) show OpenAI signing the enterprise and content agreements that

make AI products usable inside institutions.

- The Financial Times on Bain's AI replicas points to a new diligence test for software companies: can a buyer approximate the product quickly enough to question the moat?

SEGMENTS

00:00:04 Power Becomes Product

00:05:46 Control Lists and Model Access

00:09:12 OpenAI Signs the Operating Agreements

00:13:27 China Builds Into Daily Life

00:17:39 Agents Need Owners

Transcript

1. Lenar Kess 00:00:04

Microsoft and Chevron signed a 20-year power agreement for a proposed AI data center in West Texas. That's the first fact of the day, and it matters because it refuses to stay in AI vocabulary. The Wall Street Journal describes Project Kilby as a Reeves County site powered by a new 2.7 gigawatt gas-fired plant. Chevron would supply local natural gas. GE Vernova is involved on turbines, and Joulent, the Engine Number One-backed energy startup, is part of the deal. Power would start in 2028, the data center would initially sit outside the Texas grid, and it could later connect and sell surplus power. So today we're starting with the physical contract. That means fuel and turbines first. County permits come next, along with water, financing, and local politics. After that we'll move through China's export controls and OpenAI's enterprise and Getty agreements. Then we'll end with China's platform-and-workforce story, plus a builder segment on agents, bugs, and AI-assisted private equity diligence.

- techmeme.com
- wsj.com
- axios.com
- axios.com

- news.gallup.com
- techmeme.com
- techmeme.com
- apnews.com
- scmp.com
- ft.com
- forbes.com
- openai.com
- newsroom.gettyimages.com
- cnbc.com
- restofworld.org
- scmp.com
- ft.com
- youtube.com
- aws.amazon.com
- aws.amazon.com
- github.com
- x.com
- ft.com
- x.com

2. Damra Vol 00:01:08

The detail to keep is that Microsoft isn't waiting for an abstract grid upgrade to materialize. It's attaching a compute plan to a power plant. If you're building anything on top of frontier systems, that changes the mental model. Capacity is no longer just the cloud provider saying more tokens are coming. It's a fuel supply, a turbine queue, and a county that has to live next to the project.

3. Lenar Kess 00:01:32

Right. The altitude should stay sane here. A single power agreement doesn't prove that every AI company will become an energy developer. It shows one path hyperscalers are taking when normal interconnection timelines are too slow. The Journal says the project could become one of the largest AI data centers in the country. Bloomberg's writeup makes the same basic point: Chevron is selling gas-fired power for a West Texas data center rather than just selling fuel into the market. That isn't a model benchmark. It belongs to the supply chain.

4. Damra Vol 00:02:05

And it gives you a different question to ask about AI capacity. How many megawatts can they control, and under what contract? Who carries the risk if demand changes before the site is fully built? Who carries it if regulation or local opposition changes?

5. Lenar Kess 00:02:21

The water side arrived in the same news cycle. Axios has Nvidia saying AI's water challenge is largely solved by the next generation of data center cooling: warm liquid cooling, recirculated fluid, and systems that can run at 113 degrees Fahrenheit rather than relying on the old chilled-water pattern. I'd be cautious with the word solved. Axios is reporting Nvidia's claim, and the claim is strongest around cooling water inside the facility. It doesn't make the power plant's water use disappear. It also leaves local concerns about construction, noise, transmission, and land.

6. Damra Vol 00:02:58

The distinction matters. A data center can use less water at the rack and still cause fights over the power system that feeds it. If the electricity comes from gas generation, water shows up in a different part of the chain. If the site needs roads, substations, and emergency response capacity, the resource fight moves again. Nvidia may have a serious cooling advance. That doesn't turn a data center into a frictionless object.

7. Lenar Kess 00:03:26

The public reaction is the next piece. Axios ran another piece today on data centers becoming the face of AI backlash. The Milltown Partners poll they cite had 6,872 registered U.S. voters, and nearly half supported a temporary ban on new data center construction. Gallup had a related result last month: seven in ten Americans said they opposed an AI data center in their local area, with 48 percent strongly opposed. Those numbers aren't the same survey, but they rhyme enough to be taken seriously.

8. Damra Vol 00:03:59

That's a strange political object, too. People can like better search, better coding tools, and better medical research and still reject the building down the road that makes it possible. And to be fair to them, the building down the road is what they can inspect. They can see land use, water applications, and power bills. They can't inspect the model training run that made the demand.

9. Lenar Kess 00:04:21

The finance item makes the same point from the capital side. Techmeme's data-center finance item points to Morgan Stanley pitching leveraged loans to data center clients, with estimates around tens of billions in AI-related leveraged finance this year and much more over the next several years. Leveraged loans aren't exotic in private equity, but seeing that tool move into AI infrastructure financing tells you the buildout is hunting for more kinds of capital than ordinary corporate bonds and hyperscaler cash flow.

10. Damra Vol 00:04:51

And that is where builders should be a little less romantic about compute. The model you call through an API may depend on a private-credit desk, a gas turbine delivery date, a cooling design, and a local vote. You don't have to become an energy analyst to write software. But if your whole product assumes cheap elastic inference, the stack below you is getting more contractual and more political.

11. Lenar Kess 00:05:16

What I take from today's lead is narrow: AI capacity is getting named. Project Kilby is one noun. Reeves County is another. The same goes for the 2.7 gigawatts of planned gas-fired power, the 2028 delivery date, Nvidia's warm liquid cooling claim, local moratorium polling, and leveraged finance. Those details aren't as fun as a new model card, but they decide whether the model card can be served to millions of people at a price anyone can tolerate.

12. Lenar Kess 00:05:46

China added ten U.S. companies to an export-control list and barred government purchases from 46 more U.S. companies. Techmeme's headline names MP Materials and Teal Drones in the export-control group, and Anduril in the procurement-ban group. The Associated Press reports that the affected firms are military-related and that the restrictions cover dual-use items, meaning goods with both civilian and military applications. The South China Morning Post says export operators are prohibited from shipping those dual-use items to the ten firms.

13. Damra Vol 00:06:19

That's the supply-chain version of the control problem. It isn't a press statement about AI values. It's a list of companies that can no longer receive particular materials or products from Chinese exporters, plus a procurement rule for Chinese government buyers. And rare earth firms being on that list isn't incidental, because rare earth supply is one of the places where the physical and military parts of the technology stack touch each other.

14. Lenar Kess 00:06:45

Exactly. And the control story isn't only physical supply chains this week. The Anthropic export-control fight is still moving in the background. We covered the model-access side recently, so I don't want to turn this into a second version of that episode. But the Financial Times analysis around Anthropic's own risk language, and Forbes asking whether advanced models might be treated as military-grade systems, both belong next to the China move. One side is dual-use exports into defense-linked firms. The other is whether frontier model access becomes something governments classify and ration.

15. Damra Vol 00:07:22

There's a practical difference between those two surfaces. A rare earth control list is painful, but companies know some of the routing moves. They look for alternate suppliers, substitute parts, build stockpiles, and accept longer lead times. A model-access control can cut through a product at the API boundary. If your app depends on a capability that your customers can lose by nationality, geography, or government category, that isn't just procurement risk. That's product behavior.

16. Lenar Kess 00:07:52

And this is why the China item shouldn't be treated as a generic trade-war story. It's a reminder that AI infrastructure includes materials and drones. It also includes defense companies, software models, and government purchasing rules. These stories all concern control over inputs. Some inputs are magnets and components. Some are model weights or hosted access. Some are contracts that say who's allowed to buy what.

17. Damra Vol 00:08:19

For a builder, the practical question is how long you can abstract that away before it becomes irresponsible to ignore. If you're writing a chat app, probably most of it. If you're building in defense, energy, chip design, or biosecurity, you need an access plan that survives a policy change. The same goes for anything sold across borders. That means more than a second API key in a config file.

18. Lenar Kess 00:08:44

And there is a human piece here that shouldn't get sanded down. Export controls sound administrative until they land on a team that can't buy a part, a researcher who can't access a model, or a customer whose contract gets rewritten by a government decision. Monday's China action gives us a concrete new example on the company side. The Anthropic fight gives us the model-access example. The shared question is who gets to continue using the system after the policy line moves.

19. Lenar Kess 00:09:12

OpenAI announced that Samsung Electronics is deploying ChatGPT Enterprise and Codex to employees around the world. The primary announcement says the rollout covers all Samsung Electronics employees in Korea and all consumer devices division employees worldwide. OpenAI says the work spans research and development, manufacturing, marketing, product development, and corporate functions. OpenAI calls it one of its largest enterprise launches ever.

20. Damra Vol 00:09:42

That sentence proves less than some people will want it to prove. It establishes access, not productivity. It says Samsung is making ChatGPT Enterprise and Codex available to a very large

employee population. It doesn't prove that manufacturing yield improves, that marketing gets faster, or that research becomes better. Those outcomes have to be measured later.

21. Lenar Kess 00:10:05

Yes. And I like that calibration. The source is an enterprise deployment announcement. It's still significant because Samsung isn't a toy account. It has phones and chips. It has appliances, factories, research groups, and a lot of internal process. If Codex gets used inside that environment, the interesting part isn't just code generation. The harder question is how a coding agent sits inside a company with hardware and firmware. Security review, manufacturing constraints, and local language needs all come with it.

22. Damra Vol 00:10:37

Also, Samsung SDS becoming involved as a reseller and manager for ChatGPT Enterprise in Korea matters. That makes the deployment less like a lab team buying seats and more like a channel relationship. Enterprise AI adoption often becomes a question of who handles identity and support. Compliance review, procurement, and renewal follow close behind. The model may be the exciting object, but the deployment survives through the paperwork around it.

23. Lenar Kess 00:11:06

The other OpenAI agreement today was Getty Images. Getty's own newsroom says it signed a multi-year display partnership that brings Getty's licensed content libraries into OpenAI search and discovery experiences within ChatGPT. The stated use is display in ChatGPT, with richer visual responses. Reporting around the deal also notes what the release doesn't say: this isn't described as a training-data deal. It's about licensed images appearing in product experiences.

24. Damra Vol 00:11:37

That distinction is everything. Training rights and display rights are different legal objects. Search rights, attribution, creator economics, and indemnity are different product problems too. Getty spent years as one of the most visible rights holders in the AI image fight. Now it is also supplying licensed images into a major AI interface. That doesn't settle the copyright argument, but it does show one commercial path: make the image library useful inside the answer surface and get paid for that access.

25. Lenar Kess 00:12:08

And it pairs oddly well with Samsung, even though the two deals are separate. One agreement gets OpenAI deeper into an industrial enterprise. The other gets rights-cleared visual supply into ChatGPT's discovery layer. Neither story announces a new model capability. They are distribution

and permission stories. Who can use the tool at work? Which images can appear in the answer? Who signs the agreement that makes the experience durable enough to depend on?

26. Damra Vol 00:12:36

That's the platform work people tend to undercount because it doesn't demo well. You can't show procurement onstage for applause. But if the rights agreement is missing, the product is brittle. If the enterprise identity and admin layer are missing, the rollout stalls. If the reseller channel is missing, local support becomes a spreadsheet and a prayer. <soft>And spreadsheets have done enough for civilization already.</soft> [chuckle]

27. Lenar Kess 00:13:02

The laugh is earned. The serious version is that OpenAI is becoming a company of agreements as much as a company of models. Samsung is an agreement about employees and enterprise controls. Getty is an agreement about licensed media in the answer surface. On a day when the lead story is Microsoft signing for power, that comparison feels fair: AI products are being made out of contracts around the model.

28. Lenar Kess 00:13:27

Tencent is testing an AI assistant called Xiaowei inside WeChat. CNBC and other outlets describe it as an early-stage assistant that uses Tencent's own WeLM and sometimes DeepSeek, with the obvious reason for paying attention: WeChat isn't just a messaging app. It handles payments and mini-programs. It also carries services, commerce, and everyday coordination for a huge number of Chinese users.

29. Damra Vol 00:13:54

Putting an assistant there is different from launching a separate chatbot. A separate chatbot has to become a habit. WeChat already is one. If the assistant can operate across mini-programs, answer inside existing flows, or coordinate tasks where people already transact, Tencent gets to test AI in a place where distribution has already been solved. The technical question becomes permission and reliability, not whether anyone downloads the app.

30. Lenar Kess 00:14:22

The education story from Rest of World and related Chinese reporting points in the same direction: institutions are changing around AI, not only adding chatbots. Chinese universities have cut thousands of degree programs in arts, humanities, foreign languages, and management. At the same time, they have added programs in artificial intelligence, robotics, chips, and embodied intelligence. South China Morning Post put the count at roughly 12,200 removed or suspended degree programs and about 10,200 new ones between 2021 and 2025, affecting more than 30 percent of programs.

31. Damra Vol 00:15:04

That is a large administrative intervention, and I'd be careful not to treat it as proof that the labor market is fixed. Changing the catalog isn't the same as creating good jobs. It shows what the state and universities think the workforce problem is. They are trying to move students toward technical fields at the scale of the degree system.

32. Lenar Kess 00:15:24

Then JD.com's founder Richard Liu added the labor-side version. The Financial Times reports that Liu warned robots will replace the company's 700,000 delivery workers sooner or later. The same reporting says JD has partnerships with 120 schools to retrain workers for robot repair and equipment servicing. AI Weekly quotes his welfare pledge: he doesn't want those 700,000 employees left without jobs or income.

33. Damra Vol 00:15:53

The story gets morally and mechanically harder at that point. Retraining is the responsible sentence. But moving hundreds of thousands of delivery workers into robot maintenance isn't a sentence, it is a labor-market experiment. How many repair jobs exist? Where are they located? What happens to older workers, migrant workers, and people who took delivery work because it was available now?

34. Lenar Kess 00:16:17

And the three items together tell a sharper story if we leave them separate. Tencent is platform distribution. Universities are labor-supply policy. JD is automation demand and worker transition. We don't need to force them into one master explanation. They are three signs that China's AI adoption is being routed through institutions people already depend on: the super-app, the university system, and the logistics network.

35. Damra Vol 00:16:45

The WeChat item is the one a builder might underestimate because it sounds like another assistant. But assistants embedded in operating platforms are qualitatively different from assistants sitting in a browser tab. If WeChat gives Xiaowei reach into payments, services, and mini-programs, agent design starts to meet platform operations. Account recovery shows up. So do fraud, parental controls, merchant disputes, and all the messy things consumer platforms already had to learn.

36. Lenar Kess 00:17:16

And the university and JD stories give the other side of that: people aren't only users of AI interfaces. Institutions adapting to automation are sorting people, training them, displacing them,

protecting them, or promising protection. The plainest sentence I can say is that the ambition is legible and the transition cost isn't solved by saying retraining.

37. Lenar Kess 00:17:39

Nate B. Jones put out a short briefing called, You Can't Run AI Agents Without This. The YouTube chapter markers are almost the whole thesis: from assistant to agent; job, diet, boundaries. The practical idea is that an agent should have an owner and a job. It also needs an information diet, boundaries, and review loops. I'm paraphrasing from the briefing and public chapter markers here because the Braid tool bridge couldn't fetch the transcript today, so I'm using the point as a builder note rather than as a quoted source.

38. Damra Vol 00:18:12

That scope feels right for today. We did a deeper agent engineering episode yesterday, so today's update should be smaller: agents are starting to look like operational assets. You don't just create one; you assign it a scope, feed it the right context, measure it, and decide who's paged when it behaves badly.

39. Lenar Kess 00:18:32

AWS is pushing the same direction from the platform side. Amazon Bedrock AgentCore Harness is now generally available. AWS says the managed harness lets customers define the model and tools in configuration. They also define skills and instructions, then AgentCore assembles and runs the loop. The release notes say the harness is available in supported AWS regions. The AWS blog emphasizes managed memory as an addressable resource. You can query it, attach it to a different agent, audit it, or hand it to analytics.

40. Damra Vol 00:19:07

The lock-in and leverage trade is sitting right there. A managed harness can save a team from rebuilding orchestration and memory. It can also absorb tool routing, streaming, and lifecycle code. It also means the decisions about how an agent runs are now partly shaped by a cloud provider's object model. The config file becomes a product boundary.

41. Lenar Kess 00:19:30

The counterexample today is almost comically ordinary: a GitHub issue against OpenAI Codex says the CLI's SQLite feedback logs can write about 640 terabytes a year and rapidly consume SSD endurance. The issue author says Codex was continuously writing to local SQLite feedback log files, and after about 21 days of uptime the machine had written about 37 terabytes, with process and file checks pointing to Codex logs as the main continuous writer.

42. Damra Vol 00:20:03

That bug is a perfect agent-era reminder because it isn't science fiction. It's a local write loop. The tool can be brilliant at editing code and still chew through SSD endurance because a feedback database is too chatty. For me, that is the maintenance lesson: when agent tools become part of daily work, ordinary local behavior becomes product quality. Logs and caches aren't backstage if they wear out the machine. Neither are write-ahead files, retries, and telemetry.

43. Lenar Kess 00:20:35

And this connects to Omar Sar's agent-evaluation post without needing to make a big theory. He points to scalable evaluation for AI agents in production and says the interesting part is front-loading human judgment into reusable evaluation criteria. Teams keep trying to avoid that work because it is slower than a demo. But once an agent is allowed to act in a production workflow, you need reusable judgment. Otherwise every incident becomes a bespoke argument about whether the agent did the right thing.

44. Damra Vol 00:21:07

Reusable judgment is a nice phrase because it has teeth. A rubric is only useful if it changes what the system is allowed to do next. If evaluation is just a dashboard, people will admire the red line while the agent keeps writing the wrong file. The interesting systems will couple evaluation to routing, permission, rollback, and review.

45. Lenar Kess 00:21:28

One more builder-finance item belongs in this segment because it changes the social life of coding agents. The Financial Times reports that Bain is using AI-generated software replicas during private equity due diligence. The basic idea is blunt: if a buyer can prompt its way into a rough version of your product, that tells them something about whether the target's software is defensible.

46. Damra Vol 00:21:53

Founders are going to hate that, and some of the fury will be justified. A rough replica can miss the hard parts. Data quality matters. Edge cases matter. So do customer trust, integrations, regulatory history, and the ten years of tiny decisions that make software survivable. But as a diligence tool, it is a good provocation. If the buyer can clone the happy path in a weekend, maybe the value isn't in the interface. Maybe the value is in distribution or workflow knowledge. Maybe it is in data rights, switching cost, or the support organization.

47. Lenar Kess 00:22:29

Ethan Mollick had a related post yesterday about current coding tools being impressive but still bounded. He argues that tools like Codex and Copilot are changing knowledge work, but they don't remove the need for judgment about what the work is. That fits the Bain item cleanly: AI can approximate code, but the buyer still has to decide what the approximation proves.

48. Damra Vol 00:22:53

And it means software defensibility becomes more legible. If a company's moat was mostly that the product looked hard to build, AI-assisted replicas damage that story. If the moat is customer context or operational depth, the replica may make that clearer. The same is true for data access, trust, or a workflow that resists simplification. I wouldn't want a deal team treating vibe-coded clones as truth. I'd want them treated as questions that get asked earlier.

49. Lenar Kess 00:23:24

So the builder close for today is practical. Agents need owners. Managed harnesses are becoming cloud products. Evaluation has to carry reusable human judgment. Local bugs still matter. And AI-generated replicas are entering the room where software value gets negotiated. It's a lot for Monday, but it isn't a single revolution. It's a set of ordinary systems becoming more accountable because people are starting to rely on them.

50. Damra Vol 00:23:52

The sentence I'd leave in the margin is simpler: if an AI system is useful enough to depend on, it is useful enough to maintain. Today gave us power contracts and export lists. It also gave us licensing deals, university catalogs, delivery-worker retraining plans, managed agent harnesses, and a SQLite log bug. None of those are glamorous, and every one of them changes whether the technology holds up when people lean on it.

51. Lenar Kess 00:24:19

Tomorrow, the evidence should be the concrete follow-through. Project Kilby still needs local approval and grid terms. Nvidia's cooling claim needs to show up in facility disclosures. Getty's deal has to change ChatGPT's visual answers. Tencent has to open Xiaowei beyond testing, and the Codex logging issue needs a release fix rather than only a diagnosis in a thread. At that point, the day's announcements become systems people can trust or systems people route around. Lenar Kess.

Hosts on this episode

- Lenar Kess moderator
- Damra Vol critic

