When Chip Access Became Diplomacy

2026-06-23 / 00:30:57

> *“If compute access is negotiated by states, the model choice in your stack inherits a foreign-policy dependency.”*
>
> — from this episode's transcript

- Lenar Kess
- Damra Vol

Today’s episode follows AI supply chains as they move from vendor strategy into state coordination, then turns through Google’s pressure points, Oracle’s AI-linked cuts, builder-facing GLM infrastructure, supercomputing concentration, content rights, and security.

- Techmeme’s Pax Silica report gives the lead: the Netherlands joined a US-led chip supply-chain effort with South Korea and Japan, while Taiwan endorsed it without becoming a signatory.
- IEEE Spectrum on Europe’s tech sovereignty package sets the counterpoint: Europe is still trying to reduce dependency in the same stack everyone wants to coordinate.
- Axios on Google DeepMind departures, CNBC on Google’s search position, and Axios on People Inc.’s crawler complaint make Google’s AI-era pressure concrete without treating it as a collapse story.
- CNBC on Oracle’s 21,000 role reductions and TechCrunch’s 2026 layoff tracker keep the labor conversation tied to a reported number rather than a slogan.
- Prime Intellect’s prime-rl v0.6.0 post and Perplexity Developers on GLM-5.2 in the Agent API move the model-release item into builder territory: training infrastructure

and agent availability.

- NVIDIA's TOP500 post and Techmeme's LineShine report show two infrastructure claims coexisting: national leadership in a fastest-machine ranking and vendor concentration across installed systems.
- OpenAI's DayBreak post, Al Jazeera on the Five Eyes warning, and IEEE Spectrum on vibecoding malware close the episode with the security proof problem: offense, defense, and generated code are all becoming more inspectable and more dangerous.

SEGMENTS

Transcript

1. Lenar Kess 00:00:04

   Techmeme has the Netherlands joining the US-led Pax Silica initiative today, alongside South Korea and Japan, with Taiwan endorsing it as a non-signatory. Start with the plain fact: another country has moved into a chip supply-chain coordination effort. The word <emphasis>coordination</emphasis> is carrying the weight here. It pulls export policy, advanced manufacturing, packaging, and access to the machines toward diplomacy rather than ordinary procurement. I'd open there because yesterday we talked about power and compute as infrastructure. Today's item takes the next step: the socket on the wall is connected to a treaty table.

If you build systems that depend on frontier training or high-end inference, this doesn’t change your code by lunchtime, but it changes the assumptions your code is making. The policy questions are concrete now: who can buy the chips, who can service the tools, who can redirect supply in a crisis, and who gets told to wait.

- techmeme.com
- spectrum.ieee.org
- axios.com
- cnbc.com
- axios.com
- theguardian.com
- cnbc.com
- techcrunch.com
- indianexpress.com
- x.com
- x.com
- x.com
- x.com
- techmeme.com
- blogs.nvidia.com
- technologyreview.com
- openai.com
- aljazeera.com
- spectrum.ieee.org

2. Damra Vol 00:01:04

   And the non-signatory detail matters. Taiwan endorsing Pax Silica without being described as a full member isn't a footnote; it tells you the map is still being drawn while everyone is standing on it. Taiwan is central to advanced chip manufacturing, but its political status makes every formal structure harder. So the builder translation isn't, "great, the alliance is solved." It is, "the most important component paths in AI now run through arrangements that may be partly formal, partly implied, and partly constrained by diplomatic wording." That is a different dependency model than choosing between two cloud regions. If your provider says capacity is available, you still have to ask what happens when a state decides that capacity has a state-priority use elsewhere.

3. Lenar Kess 00:01:52

   IEEE Spectrum’s Europe piece gives the second half of the same day without turning it into the same story. Europe is talking about tech sovereignty and supply-chain capacity, which means Europe is trying to build more room to maneuver while the US-led coordination effort is recruiting

partners into a shared structure. Those can coexist. A country or region can want alliance access and still want domestic capacity, because dependence feels different after the dependency becomes visible. The temptation is to turn this immediately into a US-versus-China story. There is obviously a China layer; these policies don't exist in a vacuum. But the builder read is smaller and more concrete: every frontier system rests on a supply chain with political owners. You might rent the API, lease the GPU cluster, or fine-tune on someone else's stack, but the power to constrain that stack may sit several layers above your vendor contract.

4. Damra Vol 00:02:53

   There is also a craft cost here. When infrastructure moves under state coordination, resilience stops being just a reliability exercise. You can't answer it with retries, cached prompts, and a second account. You need some understanding of substitution. Can this workload run on a different accelerator family? Can the model be swapped without changing the product's behavior too much? Does the data boundary let you move providers at all? A lot of AI architecture advice has been too neat at this layer. It treats model choice as a routing problem. Sometimes it is. But if the high-end chip path is negotiated through alliance structures, model choice also inherits hardware availability, export policy, and the politics of where the fabs and lithography tools sit.

5. Lenar Kess 00:03:41

   And that brings in the TOP500 items, even though I'd keep them as infrastructure color rather than the lead. Techmeme reports China's Arm-based LineShine system surpassing the US system El Capitan as the world's fastest supercomputer. NVIDIA's own blog, from the same day, says NVIDIA technology powers more than four hundred of the TOP500 systems. Those two claims can both be true. A national ranking can change at the top, and a vendor can still be deeply embedded across the installed base. If you are trying to understand compute power, a single fastest-machine headline gives you one kind of signal. The installed base gives you another. And MIT Technology Review's ASML piece, with the four hundred million dollar chipmaking machine in the headline, is a reminder that the machine that makes the machine may be the narrower point in the whole system.

6. Damra Vol 00:04:36

   People under-price that layer of the stack when they talk about independence. You can announce a national AI plan in a month. You can't conjure EUV lithography or packaging expertise on that schedule. Power delivery and cooling take longer. Firmware, cluster operations, and the people who debug the weird failure at two in the morning take longer too. So when Pax Silica shows up as a policy story, it is also a maintenance story. The alliance doesn't just coordinate values or slogans; it coordinates access to a chain of skills and machines that are painful to duplicate. A builder should hear that and become a little less casual about provider concentration. Not panicked. Just less casual.

7. Lenar Kess 00:05:20

   Axios reports today on AI lab departures from Google DeepMind, including Noam Shazeer and John Jumper moving to competitors. CNBC has a separate piece arguing that Google's online dominance is showing signs of cracking in the AI era. Axios also has People Inc.'s CEO accusing Google of abusing market power by using one crawler path for both search and AI. I'm deliberately putting those in one segment and not one verdict. Talent movement, search distribution, and crawler policy are different problems. They all hit Google today, but they don't prove the same thing. The stronger claim is that Google is getting pressure at three different interfaces: the people who invent the methods, the users who find information, and the publishers who supply the web that search and AI both consume.

8. Damra Vol 00:06:09

   The crawler complaint is the most mechanically interesting one to me. If a publisher wants search traffic, it has historically allowed Googlebot. If that same access path also feeds AI answers, the publisher's choice gets uglier. They are no longer choosing between indexing and invisibility in search alone. They may be choosing between distribution and training or answer-generation use that competes with their own page. Axios is reporting the accusation from People Inc.'s CEO, so we should keep it as an accusation. But the product mechanism is easy to understand. A single crawler can collapse two consent decisions into one operational toggle. The policy choice has become executable product behavior.

9. Lenar Kess 00:06:54

   And it intersects with the search story without becoming identical to it. CNBC's piece is about Google's dominance showing cracks as AI answers change how people search, how ads sit around answers, and how competitors route queries. The publisher complaint is about whether Google can use distribution power to secure inputs for AI. Those are adjacent pressures. One is user behavior and revenue. The other is bargaining power over content. If Google is strong in search, publishers feel they can't say no. If AI weakens search, Google has more reason to protect the data and answer loop. You don't need a collapse narrative to see the squeeze.

10. Damra Vol 00:07:34

    The talent piece is the easiest to overstate because named researchers make for dramatic copy. Construct already spent time on John Jumper's move, so I wouldn't replay that as though the departure itself is new. Today's update is that the departures are now being read against search and publisher leverage. Google still has world-class research, a dominant distribution business, cloud infrastructure, Android, YouTube, and enormous capital. Pressure doesn't mean weakness. It

means the AI transition is attacking several advantages in different ways. Some advantages become more valuable; others become politically expensive.

11. Lenar Kess 00:08:14

    The Guardian's Australia item belongs near this too, but I'd keep it distinct. Senator David Pocock is urging Prime Minister Anthony Albanese to stop tech companies training AI models on Australian content. That is a policy demand about national content rights, not the same legal claim as People Inc.'s complaint against Google. The shared practical problem is consent under dependency. Can a publisher, artist, newsroom, or country say no to AI training without losing distribution, discoverability, or economic access? That is the boundary everyone is trying to draw. The legal tools differ by jurisdiction, but the operational pain is similar: content owners want refusal to be a real option, and platforms don't want the web's permissions surface to turn into a thousand separate negotiations.

12. Damra Vol 00:09:05

    And refusal has to be machine-readable to matter. A lawmaker can say "don't train on this," and a publisher can update a policy page, but the crawler, the dataset builder, and the model vendor need a usable boundary. An ambiguous boundary turns into a support queue or a lawsuit. It can also become a public fight over what the crawler was allowed to do: search, AI answers, snippets, summaries, or all of it. The Google story reaches beyond corporate drama. It previews how brittle the permission layer around the web has become. The web was built around linking and indexing. AI systems are asking it to express consent for extraction, transformation, memory, and answer substitution. That is a lot to stuff into a robots file and a press quote.

13. Lenar Kess 00:09:54

    CNBC reports that Oracle shed 21,000 roles over the past year while citing AI adoption and deployment across operations. The Indian Express has the same broad item, pairing the reductions with Oracle's AI investment. TechCrunch, meanwhile, has a running list of major 2026 tech layoffs where employers cited AI. I want to stay literal here. When a company ties AI to reductions, that doesn't prove every eliminated role was automated away by a model. It does prove that AI has become part of how executives justify restructuring. That is already a change in the labor story, because the claim moves from "AI may affect jobs someday" to "AI is part of the budget and headcount explanation now."

14. Damra Vol 00:10:41

    The number makes the story harder to hand-wave. Twenty-one thousand roles isn't a pilot program. But I agree with the caution. When a large company reduces headcount, the causes are usually layered. Margin pressure, product changes, management fashion, investor signaling, cloud capital

spending, and automation can all sit inside the same announcement. The serious test is what work moved. Did support volume drop because self-service improved? Did finance close faster with fewer manual reconciliation steps? Did sales operations lose people because the tooling genuinely removed coordination work, or because the company decided to accept worse coverage? Those are different outcomes, and they all fit under a lazy sentence about AI efficiency.

15. Lenar Kess 00:11:27

    That is why I'm allergic to both easy versions of this conversation. The cheerleading version says, "AI productivity, look at the savings." The denial version says, "companies are just using AI as cover." Both can be true in different departments. If I were reading Oracle from the outside, I'd ask for evidence at the work-unit level. Where did throughput improve? Where did error rates change? Where did customers wait longer? Where did the review burden move? If a process now has fewer employees but more escalation, more rework, or more brittle approvals, then the cost didn't disappear. It moved into a harder-to-see account.

16. Damra Vol 00:12:08

    And for builders, this is also a warning about the story your tools will be used to tell. A coding assistant, an internal agent, or a document workflow may start as craft infrastructure. Six months later it may sit in a slide that says a team can run with thirty percent fewer people. That doesn't make the tool bad. It does mean the evidence matters. If the model is saving time, instrument it. If it is shifting work into review, instrument that too. The humane version of automation isn't pretending jobs are untouched. It is refusing to let vague productivity claims hide where the work actually went.

17. Lenar Kess 00:12:46

    The TechCrunch list helps because it prevents Oracle from becoming a one-company morality play. We are starting to get a set of public examples where employers cite AI in reductions, and the list can become more valuable if it tracks the exact claim each company made. Was AI cited as a direct replacement, an efficiency program, a reallocation toward AI investment, or a general restructuring factor? Those distinctions sound bureaucratic, but they matter. A direct replacement claim invites one kind of scrutiny. A reallocation claim invites another. A company moving cash from a legacy support organization into data centers and AI engineering is making a capital-allocation decision, not just an automation claim. The person affected may experience both as a layoff. The industry should still name the mechanism.

18. Damra Vol 00:13:36

    And the mechanism can be uncomfortable without being mysterious. If an enterprise vendor believes AI will let it support more customers with fewer internal steps, it will try to capture that

margin. That is what companies do. The pressure on the rest of us is to avoid treating every saved hour as a saved person. Some saved hours become better service, some become more output, some become margin, and some become layoffs. The accounting has to follow the work after the announcement. Who is still doing it? How often does it fail? Who handles the weird case? Who signs off when the agent is wrong? Those answers tell you whether the AI system improved the process or just made the org chart smaller.

19. Lenar Kess 00:14:20

    Prime Intellect announced prime-rl v0.6.0 today, describing reinforcement learning infrastructure for trillion-parameter mixture-of-experts models on agentic software-engineering tasks. Perplexity Developers also posted that GLM-5.2 is available through its Agent API. Elie Bakouch pointed to the reinforcement-learning infrastructure angle, and Joshua Saxe described GLM-5.2 as a security emergency. That is a dense cluster, so I'd separate the claims. One claim is availability: builders can reach GLM-5.2 through an agent API. Another claim is training infrastructure: Prime Intellect is pushing reinforcement learning tooling for very large mixture-of-experts systems on coding-agent workloads. The third claim is risk: a capable model in agent workflows changes the defensive assumptions around code and security.

20. Damra Vol 00:15:19

    The agent API detail matters more than a leaderboard mention would. A model becomes a builder story when it is wired into tools, priced, documented, and reachable from the workflow people already use. Even if a model is impressive in isolation, the practical change happens when it can call tools, hold task state, run code, recover from errors, and fit inside an agent harness. For GLM-5.2, I wouldn't rank it against closed models without sourced benchmark evidence in front of us. But availability through Perplexity's Agent API means people can test it in the kind of loop where model personality and failure behavior show up fast: multi-step coding, browsing, file edits, and tool calls.

21. Lenar Kess 00:16:06

    Prime Intellect's post is interesting because it shifts attention from the model name to the training machinery around agent behavior. Reinforcement learning on agentic software tasks isn't the same as making a chat model sound more helpful. You need environments, reward definitions, task traces, execution feedback, and enough control over the training loop that the model learns behavior you can actually use. The trillion-parameter mixture-of-experts phrasing also matters because MoE systems can activate a subset of parameters per token while still carrying a very large total parameter count. That makes infrastructure claims tricky. The headline size sounds enormous, but the practical question is how the training and serving path behave under actual agent workloads.

22. Damra Vol 00:16:52

    And agent workloads are messy. A benchmark task may need the model to inspect a repository, edit a file, run tests, interpret a failure, and decide whether to keep going or back out. The reward isn't just "the final answer matches." It is whether the path got there without corrupting the project, hiding a failing test, or burning silly amounts of compute. That is why I'm more interested in the infrastructure than the slogan. If prime-rl helps researchers train models against those long-horizon behaviors, it belongs in the builder conversation. If it mostly lets people print larger numbers next to coding tasks, the excitement will fade when the first real repository gets weird.

23. Lenar Kess 00:17:37

    Joshua Saxe's "security emergency" wording deserves attribution because it is strong. I don't think we need to adopt the phrase as the show's verdict. But it is a serious security researcher saying: this class of capability changes the risk environment. And I can see why. A better coding agent can help defenders patch software and audit code. The same kind of agent can help attackers inspect targets, adapt exploit attempts, and stitch together pieces that used to require more skill. The capability isn't morally sorted by the API key. It depends on the surrounding controls, logging, environment limits, and review habits.

24. Damra Vol 00:18:17

    That security point connects to the agent harness itself. If you give a model file access, a shell, network calls, package managers, and credentials, the model isn't just generating text anymore. It is operating inside a system. The safety story then lives in permissions and audit trails. It also lives in sandboxing, dependency policy, and what the human reviewer can actually inspect. I don't have a source saying GLM-5.2 specifically fails any of those boundaries. The point is broader: every capable agent model makes the harness more important. The better the model gets at using tools, the less acceptable it is for the tool environment to be casual.

25. Lenar Kess 00:18:58

    So the compact builder read is: GLM-5.2 is more interesting today because it is available in an agent API, Prime Intellect is pushing training infrastructure for agentic coding behavior, and the security community is reacting to the same capability from the other side. That is enough for a serious segment. It doesn't need to be a grand open-model turning point. The next evidence should be mundane and specific: trace quality, cost per solved task, failure recovery, and permission behavior. I would also want to know whether teams can reproduce the claimed gains outside curated demos.

26. Damra Vol 00:19:38

And I’d add one more: how it behaves when the repository is historically messy. Old tests, partial mocks, and generated files are enough to expose a model. Stale docs, weird package constraints, and unclear ownership expose it too. Agent models often look best in tasks shaped for agents. Real software asks whether the model can keep its footing when the project is full of historical compromises. The API release becomes more than access at that point. It becomes a testing invitation. Builders will find out quickly whether GLM-5.2 is merely impressive in the abstract or useful in the actual workbench sense: it makes progress, preserves context, and leaves behind a change someone can review.

27. Lenar Kess 00:20:24

    Techmeme’s TOP500 item says China’s Arm-based LineShine passed El Capitan to become the world’s fastest supercomputer. NVIDIA’s blog says its technology powers more than four hundred systems on the TOP500 list, and MIT Technology Review is pointing at ASML’s four hundred million dollar machine as a key chipmaking bottleneck. I know we already touched this in the Pax Silica segment, but it deserves one pass on its own because rankings can mislead you if you read only the winner. Fastest system is one signal. Installed base is another. The lithography tool chain is another. AI infrastructure lives across all three.

28. Damra Vol 00:21:06

    The fastest-system headline is politically potent because it is easy to understand. One machine, one list, and one country on top. The installed-base claim is less cinematic but probably more revealing for developers. If NVIDIA’s technology is in more than four hundred TOP500 systems, as NVIDIA says, then the software ecosystem around NVIDIA remains deeply embedded. So do the debugging practices, libraries, and procurement habits. That is a vendor source with a vendor incentive, so the attribution has to stay attached. But even as a self-interested claim, it points at a structural fact: installed base creates gravity through tooling and expertise. A fastest machine can change prestige. A dominant platform changes how people write and operate the workloads.

29. Lenar Kess 00:21:56

    The ASML story earns its place here. A four hundred million dollar machine for the future of chipmaking isn't a fun trivia number; it is the physical form of the bottleneck. Frontier AI keeps dragging the conversation back to software, but software depends on a manufacturing chain where a single class of machine can decide who can build the next generation of chips at all. That makes the state coordination story less abstract. Countries aren't coordinating supply chains because they love committee meetings. They are coordinating because the number of irreplaceable points in the chain is small enough that access becomes state-priority.

30. Damra Vol 00:22:34

And builders feel that secondhand. You may never touch an ASML machine, but you touch the effects: graphics processing unit scarcity, reserved capacity, and cloud quotas. You also touch model pricing, batch delays, and the weird way a product roadmap can be blocked by someone else's power contract or packaging line. There is a mental model correction here. The software industry got used to thinking of infrastructure as elastic. AI has made parts of it lumpy again. A cluster isn't just a cloud resource. It is chips, power, cooling, network fabric, manufacturing lead time, export rules, and operations staff bound together in one expensive system. That system doesn't stretch just because demand found a new use case.

31. Lenar Kess 00:23:24

    That lumpiness is why the supply-chain alliance story and the TOP500 story sit near each other in the day. One tells you how states are trying to coordinate access. The other tells you why access is valuable enough to coordinate. If I were building on this stack, I'd treat compute assumptions as product assumptions, not vendor trivia. Which workloads require the best frontier model? Which can fall back to smaller models? Which customer promises depend on latency or volume that only one provider can satisfy? A lot of AI products still hide those assumptions until the first quota squeeze exposes them.

32. Damra Vol 00:24:01

    The fallback point is especially important. In normal web software, a degraded mode might mean slower search or a stale dashboard. In AI systems, degraded mode can mean a different model with different reasoning behavior, different refusal behavior, different tool-use habits, and different privacy terms. That isn't a reason to avoid the frontier model. It is a reason to test the swap before policy, supply, or pricing forces it. If compute access is becoming diplomatic, graceful degradation needs to include model behavior, not just uptime.

33. Lenar Kess 00:24:35

    OpenAI's DayBreak post on GPT-5.5-Cyber is still circulating through Hacker News today, but Construct covered the release deeply yesterday, so I don't want to replay it as the main item. The fresher security picture is wider: Al Jazeera reports a Five Eyes warning that new AI models are transforming offensive cyber capabilities, and IEEE Spectrum has a piece on vibecoding malware. So we have a product artifact, a state warning, and a malware backdrop in one window. Automated cyber defense and automated cyber offense are becoming more concrete together.

34. Damra Vol 00:25:12

    The symmetry is uncomfortable because the same properties help both sides. A model that can read code, reason through a vulnerability, propose a patch, and test it is valuable for defenders. A model that can read code, reason through a vulnerability, propose an exploit path, and adapt after failure

is valuable for attackers. The difference isn't the raw capability. It is the environment around it: what the model can touch, how actions are logged, whether tools are constrained, whether suspicious behavior is detected, and whether humans can review the path. That is why a security model release isn't just a model story. It is an operations story.

35. Lenar Kess 00:25:53

    The Five Eyes warning, as reported by Al Jazeera, gives that operations story a state-security vocabulary. Agencies aren't only worried that frontier models can answer questions. They are worried that models change the cost curve for offensive work: reconnaissance, code generation, vulnerability discovery, and adaptation during a campaign. IEEE Spectrum's vibecoding malware piece adds the lower-level cultural detail. When people use AI to produce software they don't fully understand, malicious code can travel through the same convenience channel as legitimate automation. That doesn't mean every vibe-coded app is malware. It means code authorship is becoming harder to read from the surface of the artifact.

36. Damra Vol 00:26:39

    A builder can act on that. If code can arrive from a model, a template, a pasted answer, an agent, or a contributor who barely read it, then provenance and review need to get more concrete. Who generated this? What prompt or task produced it? What dependencies changed? Did tests execute, or did the agent merely say they did? A lot of teams will resist that because it feels like process. But the alternative is trusting code whose authorship and intent are partly obscured. In security, teams miss the weird edge case that way until it is already deployed.

37. Lenar Kess 00:27:15

    This loops back to DayBreak without making DayBreak the lead. OpenAI is packaging cyber capability as a defensive product, and the surrounding news says the same capability class is now part of state-security concern and malware production. That isn't hypocrisy. It is the normal dual-use problem becoming more practical. Every cyber-model release should separate demo success from operational reliability. Show the benchmark methodology, patch correctness, and false positive rates. Show how the system behaves on incomplete repositories and what human review is expected to catch. Security tools can be impressive and still dangerous if they make weak evidence look finished.

38. Damra Vol 00:27:57

    And security review has to include the model's work habits. Does it preserve context, or does it make a patch that fixes the test and opens a side door? Does it disclose uncertainty, or does it fabricate confidence? Does it stop when the environment blocks it, or does it route around the

restriction? Those aren't philosophical questions. They are properties you can test. The better AI gets at cyber work, the more the proof has to move from polished examples to repeatable traces.

39. Lenar Kess 00:28:26

    The day's largest item is still Pax Silica because it says chip access is being organized through state coordination. But the rest of the day keeps giving that claim texture. Google's pressure is about distribution, talent, and content permissions. Oracle's number is about labor and the stories companies tell around AI savings. GLM-5.2 and prime-rl are about agent capability moving into accessible builder loops. TOP500 and ASML remind you how physical the compute stack remains. The security items say capability has to be proven in the environment where it can do harm. I wouldn't force all of that into one grand theory. It is enough to say AI is becoming less separable from the institutions around it: states and publishers, employers and security agencies, chip vendors and cloud providers.

40. Damra Vol 00:29:21

    A builder can compress the day this way: the model is no longer the whole unit of analysis. Include the supply chain that feeds it and the platform that distributes it. Include the permission system that feeds it data, the organization that claims it saves labor, the harness that lets it act, and the review process that decides whether its output is safe enough to keep. The craft is getting wider. Not more mystical, just wider. The people who handle this well will be the ones who can trace a feature from prompt to model to chip access to data rights to operational proof without pretending any one layer answers the rest.

41. Lenar Kess 00:29:58

    Tomorrow's concrete check is whether these announcements start naming their boundaries more clearly. If a supply-chain pact grows, who is inside and what access changes? If Google answers the crawler complaint, can publishers separate search consent from AI use? If Oracle or anyone else cites AI in layoffs, can they show where the work moved? If GLM-5.2 impresses builders, can they publish traces that survive a messy codebase? Those are the details that turn a headline into something you can build around.

42. Damra Vol 00:30:30

    And if they don't name those details, the uncertainty doesn't vanish. It moves downstream into the person integrating the model, the security reviewer approving the patch, the publisher deciding whether to block a crawler, and the team explaining why a workflow has fewer people in it. That is the detail from the day I'd carry forward, Lenar.

Hosts on this episode

- Lenar Kess moderator
- Damra Vol critic