Safeguards Off, Internet On

2026-08-05 / 00:25:44

> *“An evaluation is only a boundary if the model inside it can't reach past the boundary. Turn off the safeguards and hand the model the open internet and you haven't built a test. You've built a smaller version of deployment with a report attached.”*
>
> — from this episode's transcript

- Lenar Kess
- Damra Vol

A government lab turned off two frontier models' safeguards, gave them the open internet, and then had to publish an incident report — which raises the question of where an evaluation actually ends and deployment begins.

- The UK AI Security Institute's incident report on a cyber evaluation of Claude Mythos 5 and GPT-5.6 Sol, with OpenAI's containment account and Anthropic's pointer post going up within the hour
- The Shai-Hulud worm back in npm across hundreds of packages, and Darcy Clarke's vlt 1.0 arguing nothing should run on your machine just because you typed install
- Liquid AI's on-device agentic model with day-zero throughput numbers, next to a ten-million-token claim with nothing attached to check
- Apple asking a judge to put OpenAI under forensic supervision
- Agents that can spend money, and agents that get told no with a signed receipt, plus Simon Willison's billing request

- Shieldstral, Starmind, revenue concentration, Oxide's Form D, and a summarizer that changed its mind overnight

---

SEGMENTS

Transcript

1. Lenar Kess 00:00:04

   Start with the setup rather than the finding. You're a government lab. You have two of the most capable models anyone has built sitting in front of you, and your job is to work out what they can do in a cybersecurity context. So you turn off the safeguards the developers ship them with, because the safeguards are exactly what you're trying to see past. And then you give the models the internet. It isn't a mirrored network you stood up in a lab or a synthetic target range. It's the internet. So where does that experiment end? What's the edge of it? Last night the UK AI Security Institute published an incident report from precisely that setup, running Anthropic's Claude Mythos 5 and OpenAI's GPT-5.6 Sol. Their reference number dates the incident to July twenty-eighth. And within about an hour of the report going up, both labs posted their own accounts of the same evaluation.

   - cdn.prod.website-files.com
   - x.com
   - openai.com
   - x.com
   - x.com

- x.com
- x.com
- x.com
- x.com
- reddit.com
- techcrunch.com
- github.com
- x.com
- mistral.ai
- x.com
- wheresyoured.at
- sec.gov
- x.com

2. Damra Vol 00:00:57

   An hour apart, in the evening, right after a government institute posts a PDF. Nobody does that by accident. Somebody had a draft ready and was waiting on the institute's clock.

3. Lenar Kess 00:01:08

   That's my read too. And it tells you the labs knew what was in the report, which is the normal way this works. The institute evaluates under an agreement, findings go back to the developer, and disclosure gets coordinated. What's different is that all three artifacts are public on the same night. Anthropic's post points straight at the AISI report. OpenAI published its own piece describing two separate incidents that occurred during external cyber evaluations, and how each one was contained.

4. Damra Vol 00:01:39

   Two incidents. That plural is doing something. [pause] Let's take the setup apart, because I think people hear "safeguards off" and picture somebody flipping a switch marked <emphasis>evil</emphasis>. What it actually means is you're testing the model, not the product. The refusal behavior, the classifier layer sitting in front of the API, and the tool allowlist are all part of the deployed system. Strip them and you're measuring the raw capability underneath, which is what a security institute should want to know.

5. Lenar Kess 00:02:09

   Right, and there's a good argument for it. If you only ever evaluate the shipped product, you're measuring the current safety stack rather than the model. The safety stack changes every few weeks. The weights don't. So you want a number that survives the next release.

6. Damra Vol 00:02:24

   Sure. But then the second half of the setup is where I get uneasy. Unrestricted internet access. If your test subject can reach any host on the open internet, you haven't built a container with a capable thing inside it. You've run deployment at smaller scale and called it a test. The evaluation boundary is a policy rather than a network boundary, and a policy is only a boundary for something that agrees to respect it.

7. Lenar Kess 00:02:50

   Which is the interesting bit of the OpenAI write-up. Their account isn't "the model did something we didn't expect and we panicked." It's "during external cyber evaluations, two incidents occurred, and here is the containment." They're describing an operational response, which implies detection, which implies somebody was watching in close to real time.

8. Damra Vol 00:03:12

   And containment is the word I'd stop on. It's a thing you do after a boundary has already been crossed. You don't contain something that stayed inside. So by publishing a containment narrative, OpenAI is telling you the model reached past where the evaluators expected it to go, and the monitoring is what's reassuring here, not the model.

9. Lenar Kess 00:03:32

   Anthropic's post is a different animal. It's shorter and it functions mostly as a pointer: here's the AISI report, read it. Which is a defensible choice, because the institute is the neutral party and their document should speak. It also means Anthropic isn't volunteering its own account of what Claude Mythos 5 did, which is information you'd like to have.

10. Damra Vol 00:03:54

    [tsk] Both of these are self-disclosures, and I'd read them for the omissions as much as the content. OpenAI tells you two incidents and containment; it doesn't tell you what capability was demonstrated in the interval. Anthropic hands you the institute's document; it doesn't tell you what internal changes followed. Neither is lying. Both are choosing what the sentence is about.

11. Lenar Kess 00:04:18

    One caution on sourcing, because this one is still moving as we talk. The Hacker News submission of the report has a handful of points and a few comments, and the top comment is one person's summary of a PDF. That paraphrase is where the "unfettered internet access" language is coming into circulation. The institute's own document and the two lab posts are the things to read. Don't let the comment become the source.

12. Damra Vol 00:04:42

    Which happens all the time, and it's how a report with real hedging in it turns into a headline with none. Somebody reads the summary of the summary, and then the summary of the summary is what six hundred people argue about.

13. Lenar Kess 00:04:55

    Stay with the containment question, because there's a second source from yesterday that sits next to it well. Trail of Bits, a security firm that has been doing serious vulnerability research for a long time, posted about where the field is going and about their own pivot in response. When a firm whose whole business is finding bugs in software says the AI half of that work is changing what they build, that is a practitioner signal, and it comes from people who have no incentive to hype the capability.

14. Damra Vol 00:05:24

    They're also the population with the least romance about it. Automated vulnerability discovery has been the promised land of security tooling for thirty years. Fuzzers, then symbolic execution, then everything after. Every wave arrived with a claim and then a long slow grind of false positives. So when that crowd changes what they're building, it's because the false-positive rate moved, not because a demo was impressive.

15. Lenar Kess 00:05:49

    There's also a smaller item from yesterday, and I'll name it as texture rather than as fact. Joseph Thacker, who does security research and pokes at these models all day, was posting questions about GPT-5.6 and what it will and won't do on cyber tasks. That's one researcher noticing behavior at the boundary. It isn't a confirmed policy change, and I'd hate for us to turn a person's observation into an announcement.

16. Damra Vol 00:06:16

    Though the fact that people are guessing about it is itself information. Nobody publishes a table saying "here is exactly what we refuse on offensive security and here is why." So the actual policy gets reverse-engineered by researchers comparing refusals, one prompt at a time, and everyone's working from folklore.

17. Lenar Kess 00:06:36

    So what do we take from the whole cluster? My read is that an evaluation is only a boundary if the model inside it can't reach past the boundary. Turn off the safeguards and hand the model the open internet and you haven't built a test. You've built a smaller version of deployment with a report

attached. That's not an argument against doing the evaluation. It's an argument that the containment procedure is now part of the methodology, and it should be published with the same rigor as the capability numbers.

18. Damra Vol 00:07:04

    And I'd add one thing on the institute side. AISI publishing an incident number at all is a change in posture. Governments usually publish findings. Publishing an incident means admitting your own experiment went somewhere you didn't plan, in a document with a date on it. That's a harder thing to do than releasing a capability score, and it's the reason I take this report more seriously than the two lab posts around it.

19. Lenar Kess 00:07:29

    Different kind of security story, and this one is live right now. The Shai-Hulud worm is back in the npm registry. The count going around yesterday, from the International Cyber Digest account, is at least eight hundred and sixty-eight compromised packages, together accounting for more than two billion installs a month. It's a credential stealer, and it's self-propagating. It uses what it steals from one maintainer to publish poisoned versions of the next package.

20. Damra Vol 00:07:57

    That eight hundred and sixty-eight is an aggregator's count and it will move, most likely upward. But the mechanism matters more than the count. Self-propagating means it doesn't need eight hundred and sixty-eight separate successful phishing attacks. It needs one, plus a publish token with too much reach, plus the fact that npm packages run install scripts on your machine as a matter of routine.

21. Lenar Kess 00:08:21

    And the developer-side view of it is exactly what you'd expect from a live incident. People are asking, in public, whether the tools they use every day are affected. There's a developer named Jean-Claude asking the maintainers of a couple of coding agents directly whether their packages are in the affected set, and as of those posts the answer was still unconfirmed.

22. Damra Vol 00:08:43

    That's what an unresolved supply chain incident looks like from outside. For most people the status yesterday was "we don't know yet." [sigh] And every agentic coding setup any of us runs sits on that same dependency tree. Your agent harness, your model client, and your test runner are all npm underneath.

23. Lenar Kess 00:09:03

Which makes the timing of the other item almost too neat. Darcy Clarke shipped vlt 1.0 the same afternoon, a replacement for the npm client. And the pitch he leads with is a design claim rather than a speed claim: nothing runs on your machine just because you typed install.

24. Damra Vol 00:09:22

    That's the right sentence to build a package manager around. Install-time script execution is the original sin here. It exists because it was convenient in two thousand ten for native modules that needed compiling, and it has been the reliable path from "a maintainer's token leaked" to "code ran on ten thousand laptops" ever since. Turning it off by default costs you some ergonomics and removes an entire category of attack.

25. Lenar Kess 00:09:48

    To be fair to Darcy, vlt wasn't built in response to this week. That's a coincidence of the calendar. He's been working on this problem since he was on the npm team.

26. Damra Vol 00:09:57

    Right, and I'd rather the coincidence than the alternative, which is a tool that gets announced three days after every incident and vanishes before the next one. The uncomfortable version of this story is that the registry's design keeps producing the same event, and the response each time is a faster scanner rather than a change in what install means.

27. Lenar Kess 00:10:18

    Let's move to models, because yesterday was unusually dense on the local side. Liquid AI released LFM2.5, at two point six billion parameters, a hundred and twenty-eight thousand token context window, and the pitch is that it plans and calls tools on device. Not a chat model you run locally for fun. An agentic model sized for a phone or a laptop.

28. Damra Vol 00:10:42

    And the numbers came the same day, which almost never happens. Nativ shipped day-zero support and posted measurements on an M5 Max: over eleven thousand tokens per second on prefill, eighty-two tokens per second on decode, and under eight and a half gigabytes at peak. That's full bf16, not a quantized build. Prince Canuma posted his own run.

29. Lenar Kess 00:11:05

    Unpack the eight and a half gigabytes for a second, because that's the number I keep looking at.

30. Damra Vol 00:11:11

It means you're not choosing between quality and fitting in memory. Normally the local story goes: here's a model, here's the four-bit quantization that makes it fit, here's the quality you gave up to get there, now go argue about it in a subreddit. At two point six billion parameters in full precision under eight and a half gigs, the compromise isn't in the conversation. You run what the lab trained.

31. Lenar Kess 00:11:34

    Though Nativ and Prince Canuma are both launch partners. Day-zero numbers from people who got early access are real measurements, but they're measurements taken by someone who wants the launch to go well. The check is whether the numbers hold when a stranger with a different machine runs the same prompt next week.

32. Damra Vol 00:11:52

    Which is exactly the standard the other announcement can't meet. Pokee AI put out Pokee-Isaac, twenty-eight billion parameters, claiming a ten million token context window on a single RTX 4090, using what they describe as a proprietary architecture that isn't decoder-only.

33. Lenar Kess 00:12:12

    [chuckle] Say that claim again slowly.

34. Damra Vol 00:12:15

    Ten million tokens of context, on one consumer graphics card, from a twenty-eight billion parameter model. And there are no weights, no paper, and no independent run attached to the announcement. Any one of those three claims would be a significant result on its own. Together, with nothing to check, I'd file it as a claim a company made yesterday.

35. Lenar Kess 00:12:36

    I don't think they're lying, to be clear. Non-decoder-only architectures are a real research direction, and there are ways to get very long context cheaply if you're willing to give something up, usually retrieval fidelity in the middle of the window. But the announcement doesn't say what was given up, and that's the sentence I'd want.

36. Damra Vol 00:12:54

    There's a nice contrast in the same day's material. Ling-3.0-flash weights went up on Hugging Face under an MIT license, in bf16 with an official eight-bit floating point build alongside it. That's a mixture-of-experts model where the lab did the quantization themselves instead of leaving it to volunteers. No claim attached. Here are the weights, here's the license, go look.

37. Lenar Kess 00:13:19

And on the software side of the same floor dropping, there's a llama.cpp pull request that caches hot mixture-of-experts experts on the graphics card instead of paging them from system memory. The reported result is thirty-three tokens per second going to fifty-six, on eight gigabytes of video memory.

38. Damra Vol 00:13:40

    That's a routing observation turned into an optimization. In a mixture-of-experts model most of the parameters are asleep on any given token, but a subset of experts gets hit constantly. So you keep the popular ones resident and stream the rest. It's not a new model. It's someone noticing the access pattern was skewed and doing the obvious thing about it, for seventy percent more throughput on a card you can buy used.

39. Lenar Kess 00:14:06

    And somebody on the LocalLLaMA subreddit posted a DeepSeek V4-Flash configuration running the full one million token context on a single RTX 5090 with DDR5 and CPU offloading through vLLM. That's around eight hundred tokens per second on prompt processing, and fifteen and up on decode, doing agentic coding. We talked about the cost side of that model on Sunday. This is a different thing: not the price of the API, but the fact that one desktop can hold the whole window.

40. Damra Vol 00:14:40

    Fifteen tokens a second is slow enough that you'd notice, and fast enough that you'd leave it running overnight on something. That's a real category. And I'd put it next to the Liquid release rather than against it. One end is a model small enough to live inside an application, and the other is a full-size model that fits on furniture.

41. Lenar Kess 00:14:58

    To the legal calendar. Apple has moved for a preliminary injunction against OpenAI, and it has asked a federal judge to put OpenAI under forensic supervision. That second request is what changes the temperature. An injunction is a court telling you to stop doing something. Forensic supervision is a court putting a person inside your systems to check.

42. Damra Vol 00:15:20

    Think about what that means operationally for a company shipping model updates on a weekly cadence. A court-appointed examiner with access to your source control and your training data provenance isn't a legal cost, it's an engineering cost. Every internal decision suddenly has an audience that can subpoena the commit.

43. Lenar Kess 00:15:39

And Apple's claim is broader than it was at filing. TechCrunch reported yesterday that Apple now says additional former employees may have carried confidential material over, so it's moved from a specific set of people to a category. OpenAI disputes it, and disputes in particular that Apple has established the residual access claim it's built on.

44. Damra Vol 00:16:01

    "More ex-employees may have" is the language of a company that has been running its own internal forensics and doesn't like what it's seeing, or of a company that wants a broader discovery order. Those look identical from outside and they're very different things.

45. Lenar Kess 00:16:16

    Musk's contribution to the public record was three words. "Can't trust OpenAI."

46. Damra Vol 00:16:21

    He has a standing lawsuit against them. That's a party to adjacent litigation posting on a platform he owns. I'd treat it as weather.

47. Lenar Kess 00:16:30

    One more item from the same day, and I'll keep it separate rather than fold it in. OpenAI was reported to be facing a three point two million dollar Department of Labor fine over hiring practices. Different matter, different agency, and no connection to the trade secrets case. It's a busy legal week for one company, and that's all it is.

48. Damra Vol 00:16:49

    Agreed, and the reason not to stitch them together is that the stitching is what makes a company look besieged when it's actually just large. Big companies have several legal matters open at all times. The forensic supervision request is the one that would change how the place operates day to day.

49. Lenar Kess 00:17:06

    Now to a run of agent infrastructure releases, and I'll say the theme once and then let the items stand. Cloudflare launched Cloudflare Wallets, which lets AI agents make stablecoin payments across the internet. That reaches us through an aggregator, so read Cloudflare's own post before quoting the terms. But the direction is unambiguous. Give the agent a way to pay for things.

50. Damra Vol 00:17:30

A group called Agentrust posted a Show HN the same day for the opposite move. You deny an agent's tool call through the Model Context Protocol, and you get back a signed receipt. Five points and no comments, so it's an idea rather than a validated tool. But put those two next to each other. One gives an agent the ability to spend, and the other produces cryptographic evidence that you stopped it.

51. Lenar Kess 00:17:55

    The receipt is the interesting move. A denial that leaves no trace is indistinguishable from a system that never got asked. If you want to argue later with an auditor, a customer, or the agent's own operator about what your infrastructure refused and when, you need an artifact.

52. Damra Vol 00:18:12

    And you need it signed, because the agent's own transcript isn't evidence. The agent will tell you it was denied. It will tell you it was denied when it wasn't, because it's a language model and the sentence is plausible either way. A receipt from the enforcement point is a different class of object.

53. Lenar Kess 00:18:29

    Simon Willison posted the same problem from the developer side. He's asking why, if you're building a feature into a web app that calls a model, the billing has to route through your account. He wants an OAuth flow where the user's own model subscription pays for the tokens they consume in your app.

54. Damra Vol 00:18:46

    That's such an obviously correct request that its absence is the informative bit. If every user brought their own model account, the developer stops being a reseller of inference and the lab keeps the direct relationship. Both sides should want it. Nobody's shipped it, and I suspect the reason is that per-app metering against a consumer subscription is a hard product to reason about.

55. Lenar Kess 00:19:08

    Two more from the same afternoon. There's a model router from Not Diamond that works natively with Claude Code, choosing model and reasoning effort before each turn. And there's a serverless agent worker service called FlyMy Harness, where you describe the worker and get one provisioned in about ninety seconds. Warp also shipped a command-line coding agent, which drew eighty-six points and about fifty comments on Hacker News.

56. Damra Vol 00:19:32

The router is the one I'd actually try. Reasoning effort per turn is the setting nobody exposes well. You're either paying for deep thinking on a file rename, or you're getting a shallow pass on the hard refactor. Picking it per turn without you asking is either a meaningful cost reduction or a new source of unpredictable behavior, and I don't think anyone knows which yet.

57. Lenar Kess 00:19:55

    Mario Zechner gave the whole category its deflation yesterday: so everybody is building chatboxes with connectors now. Which is unkind and about right.

58. Damra Vol 00:20:05

    It's accurate about the surface and wrong about the substance. The chatbox is the same everywhere. What's changing underneath is the permission model, the payment rail, and who holds the receipt when something gets refused. The competition is happening down there, and none of it is settled.

59. Lenar Kess 00:20:21

    A few standalone items to close. Mistral released Shieldstral, a three billion parameter open-weights model for multimodal content moderation. It was the highest-voted thing in the day's pool by a wide margin, at four hundred and forty-seven points and a hundred and thirteen comments on Hacker News.

60. Damra Vol 00:20:41

    And the comments are the reason it's in the episode. Multiple people said, in effect, moderation is why I never shipped mine. They had an image-sharing product or a small social app in mind, they did the arithmetic on per-call moderation API costs plus the legal exposure of getting it wrong, and they shelved it. A three billion parameter model you run yourself changes that arithmetic, because the marginal cost of checking an upload drops to roughly the electricity.

61. Lenar Kess 00:21:09

    Read Mistral's post for the actual license and which modalities it covers before you call it fully open, and there's no independent accuracy evaluation of it yet. A moderation model that's wrong in an interesting direction is worse than no moderation model, because now you have a system that says it checked.

62. Damra Vol 00:21:27

    That's what'll take a month to learn. Nobody knows its failure pattern yet, and moderation failures are asymmetric. The false negatives make the news, and the false positives make people leave without telling you.

63. Lenar Kess 00:21:39

    SpaceX and Nvidia announced Starmind, a satellite compute payload. Each satellite carries Nvidia Rubin graphics processors and Vera CPUs, and the claim is datacenter-class compute in orbit. There's no launch date, no power budget, and no capacity number.

64. Damra Vol 00:21:56

    Skip the orbit part for a second, because Musk added a detail that's odder and more specific. The same design, minus the solar array and the radiator, gets deployed on the ground in SpaceX data centers as an efficiency improvement. That's backwards from how this normally goes. You design for terrestrial and adapt for space. Designing under orbital thermal constraints and then running it in a building implies the density you get from that discipline beats what conventional data center design gives you.

65. Lenar Kess 00:22:26

    It's a checkable claim, which is more than most orbital compute announcements offer. He also committed SpaceX to Nvidia graphics processors exclusively, which is a straightforward supply statement.

66. Damra Vol 00:22:37

    Tren Griffin put a nice piece of history next to it. He was pointing at a two thousand two National Research Council report on geostationary latency while discussing Hughes going bankrupt. Arguments about where compute should sit relative to the people using it are old, and the orbital version has lost most of them.

67. Lenar Kess 00:22:57

    Two money items. First, an analysis making the rounds estimates that more than seventy percent of the AI revenue Amazon, Microsoft and Google report comes from OpenAI and Anthropic. That's from wheresyoured.at, which is explicitly bearish, and it's an analyst estimate rather than a disclosure. What would make it checkable is segment reporting that none of the three currently provides.

68. Damra Vol 00:23:21

    And if it's even directionally right, it changes how you hear those earnings calls. Hyperscaler AI growth described as a market is two customers, and those two customers are funded by some of the same investors buying the compute. I'd want the number verified before building anything on it. I'd also note that nobody has offered a competing estimate, which is its own signal.

69. Lenar Kess 00:23:43

    The counterpoint sitting right next to it: Oxide Computer filed a Form D with the SEC for four hundred and forty-five million dollars. Filed, not confirmed closed. Oxide builds on-premises rack-scale computers, and somebody just committed serious money to the position that not all of this settles into three clouds.

70. Damra Vol 00:24:03

    That's a long-horizon bet from investors who have watched the concentration story too. Two hundred and twenty-three points on Hacker News for a Form D filing tells you the audience for that argument is paying attention.

71. Lenar Kess 00:24:15

    Last one, and it's small and a little absurd. Stephan Rabanser's group published a paper on whether AI agents can do open-ended research. Within hours of the paper being indexed, Google's AI Overview flipped its answer to the paper's central research question, from a confident yes to a confident no. Same question and same product, on a different morning.

72. Damra Vol 00:24:38

    [laugh] That's his own screenshot and self-report, so it's one person's observation. But it's a beautiful little demonstration of what a summarization layer actually is. It has no position. It has an index, and it will state either side with the same confidence depending on what got crawled overnight. The paper's own claim about research agents is a separate question, and we haven't checked it.

73. Lenar Kess 00:25:02

    The AISI report itself is next on my reading list, in full rather than anyone's summary of it, and specifically whether the institute describes the containment procedure or only the outcome. If a government lab is going to run capability evaluations with the safeguards off and the network open, the containment method is a published methodology question rather than an operational detail.

74. Damra Vol 00:25:26

    And I want to see whether anyone outside Liquid AI's launch partners reproduces eighty-two tokens per second on a machine they bought themselves. That's a number a stranger can check by Friday, which makes it a better number than most of what we talked about today.

Hosts on this episode

- Lenar Kess moderator
- Damra Vol critic