

SquaredCast: Episode 4

“Are You 18+? (Sponsored by Meta)”

Recording date: March 27, 2026

URL: <https://squaredcast.com/episode-4/>

Intro	2
The Rundown (News)	3
1. Meta's Week from Hell: \$375M Child Exploitation Verdict, \$6M Addiction Verdict	3
2. Meta Is Behind the Push for OS-Level Age Verification	6
3. Nvidia DLSS 5: AI Slop Meets Gamer Rage	9
4. Anthropic vs. The Pentagon	14
5. Microsoft's "Microslop" Retreat	18
6. FCC Bans All Foreign-Made Routers	21
7. Age Verification Hits the iPhone: UK First, US Next	25
The Deep Dive (SKIPPED)	30
The Build Log	31
Chris N: Observability	31
Chris V: Same as before.	33
The Plug / Outro	34
Chris N's Plug	34
Chris V's Plug	34
SquaredCast Links	34
Final Notes	34

Intro

Welcome back to SquaredCast! This is episode four, recorded on March 27th, 2026.

Recently, two juries hit Meta with back-to-back verdicts in 24 hours: \$375 million in New Mexico for child exploitation, and \$6 million in Los Angeles for engineering addictive products that harmed a minor. While those verdicts landed, a researcher published findings on GitHub tracing Meta's lobbying operation across 45 states, with over \$25 million directly confirmed and estimates running as high as \$2 billion, all designed to shift age verification responsibility onto Apple and Google. Then Apple released iOS 26.4, and UK iPhone users got a new prompt: "Confirm You Are 18+." California's law goes further, requiring every operating system to collect user age data, Linux included. GrapheneOS, a privacy-focused alternative to Android, told regulators to pound sand. Elsewhere: Nvidia unveiled DLSS 5 and the internet called it an AI beauty filter. A federal judge blocked the Pentagon's attempt to blacklist Anthropic as a national security threat. Microsoft admitted Windows 11 has a bloat problem. And the FCC banned all new foreign-made routers from the U.S. market, which covers basically every brand on store shelves.

We've got so much news to cover that we'll actually be skipping the deep dive today!

As always, show notes and sources are at squaredcast.com. If you want to support us and get bonus content, our Patreon starts at two bucks a month. Link's in the show notes.

Let's get into it...

The Rundown (News)

1. Meta's Week from Hell: \$375M Child Exploitation Verdict, \$6M Addiction Verdict

Two juries in two different states delivered back-to-back guilty verdicts against Meta in the span of 24 hours, and the combined message is hard to misread. On March 24, a Santa Fe jury found Meta willfully violated New Mexico's consumer protection laws after a six-week trial brought by Attorney General Raúl Torrez, who accused the company of creating a "breeding ground" for child predators on Facebook and Instagram. The case originated from a 2023 undercover operation where investigators created fake profiles of children under 14 and were immediately flooded with sexually explicit content and contact from adults. Jurors awarded the state \$375 million in civil penalties at the maximum \$5,000-per-violation rate. Torrez called it "a historic victory for every child and family who has paid the price for Meta's choice to put profits over kids' safety." Meta said it plans to appeal.

Then, on March 25, a Los Angeles jury found Meta and YouTube negligent in a test case designed to set the tone for roughly 2,000 pending lawsuits. The plaintiff, a 20-year-old woman identified as KGM, alleged she became addicted to Instagram at age 9 and YouTube at age 6, contributing to depression, body dysmorphia, and suicidal thoughts. The jury awarded \$3 million in compensatory damages and \$3 million in punitive damages, with Meta shouldering 70% and YouTube 30%. This is the first time a jury has treated social media apps as defective products for being engineered to exploit developing brains. Internal Meta documents shown at trial included a memo stating "if we wanna win big with teens, we must bring them in as tweens" and evidence that 11-year-olds were four times more likely to return to Instagram than competing apps, despite the platform's stated age minimum of 13.

The second phase of the New Mexico trial begins May 4, where Torrez will argue a public nuisance claim and push for court-ordered changes to Meta's platforms, including mandatory age verification, algorithm modifications, and an independent monitor. California Attorney General Rob Bonta announced his state's own trial against Meta is set for August. These verdicts don't exist in isolation. Meta faces thousands of active lawsuits from parents, school districts, and attorneys general across the country. Legal observers are calling this the tech industry's "Big Tobacco" moment, and the comparison is getting harder to dismiss.

Sources:

- CNBC, "Meta must pay \$375 million for violating New Mexico law in child exploitation case," Mar 24, 2026: "Meta willfully violated the state's unfair practices act, and decided

the company should pay \$375 million in damages based on the number of violations." — <https://www.cnbc.com/2026/03/24/jury-reaches-verdict-in-meta-child-safety-trial-in-new-mexico.html>

- CNN, "Jury finds Meta liable in case over child sexual exploitation on its platforms," Mar 24, 2026: "The jury found Meta liable on all counts, including for willfully engaging in 'unfair and deceptive' and 'unconscionable' trade practices." — <https://www.cnn.com/2026/03/24/tech/meta-new-mexico-trial-jury-deliberation>
- Source New Mexico, "Santa Fe jury awards New Mexico \$375M in Meta child exploitation case," Mar 24, 2026: NM AG Raúl Torrez: "The jury's verdict is a historic victory for every child and family who has paid the price for Meta's choice to put profits over kids' safety." — <https://sourcenm.com/2026/03/24/santa-fe-jury-awards-new-mexico-375m-in-meta-child-exploitation-case/>
- CNBC, "Jury in Los Angeles finds Meta, YouTube negligent in social media addiction trial," Mar 25, 2026: "Compensatory damages were assessed at \$3 million, with Meta on the hook for 70% and YouTube the remaining 30%." — <https://www.cnbc.com/2026/03/25/meta-youtube-los-angeles-california-verdict.html>
- NPR, "Jury orders Meta and Google to pay woman \$3 million in social media addiction trial," Mar 25, 2026: "It represents the first time a jury has found that social media apps should be treated as defective products for being engineered to exploit the developing brains of kids and teenagers." — <https://www.npr.org/2026/03/25/nx-s1-5746125/meta-youtube-social-media-trial-verdict>
- NPR, same: Internal Meta document: "If we wanna win big with teens, we must bring them in as tweens." Another showed "11-year-olds were four times as likely to keep coming back to Instagram, compared with competing apps." — <https://www.npr.org/2026/03/25/nx-s1-5746125/meta-youtube-social-media-trial-verdict>
- CBS News, "Meta and YouTube found liable on all charges in landmark social media addiction trial," Mar 25, 2026: KGM alleged using Instagram and YouTube "contributed to her mental health problems, including depression, body dysmorphia and suicidal thoughts." — <https://www.cbsnews.com/news/meta-youtube-social-media-addiction-lawsuit-verdict/>
- Al Jazeera, "Jury finds Meta, YouTube liable for social media addiction: What we know," Mar 26, 2026: "KGM told the trial she had started using YouTube at the age of 6 and Instagram at the age of 9." — <https://www.aljazeera.com/news/2026/3/26/jury-finds-meta-youtube-liable-for-social-media-addiction-what-we-know>
- CNBC, "New Mexico seeking changes to Meta's platform after jury verdict," Mar 25, 2026: AG Torrez: "We're going to be asking for injunctive relief. That means changes to the design features of the platform itself, real age verification, changes to the algorithm, an independent monitor." — <https://www.cnbc.com/2026/03/25/new-mexico-meta-lawsuit-jury-verdict.html>
- NBC News, "Jury finds Meta and YouTube negligent in landmark lawsuit on social media safety," Mar 25, 2026: CA AG Rob Bonta: California "looks forward to holding Meta

accountable in our own upcoming August trial in the Bay Area." —

<https://www.nbcnews.com/tech/tech-news/verdict-reached-landmark-social-media-addicti-on-trial-rcna263421>

2. Meta Is Behind the Push for OS-Level Age Verification

In March 2026, a research investigation published on GitHub under the handle "upper-up" pulled back the curtain on one of the most deliberate corporate lobbying campaigns in recent tech history. The project, called TBOTE, used publicly available records, including tax filings, federal lobbying disclosures, state lobbying registrations, campaign finance databases, and corporate registries, to map how Meta built a multi-channel influence operation to push age verification laws that shift the regulatory burden away from social media platforms and onto Apple and Google's app stores. The investigation was covered by Bloomberg, The Register, Rappler, and Gigazine, among others.

The TBOTE investigation directly confirmed over \$25 million in Meta's lobbying spending tied to these efforts and estimates the company may have spent upward of \$2 billion when including dark money grant networks and fragmented political operations. The operation spans all five confirmed funding channels: \$26.3 million in direct federal lobbying in 2025 (a company record, exceeding even Lockheed Martin and Boeing), over 86 lobbyists deployed across 45 states, covert funding of a "grassroots" child safety group, \$70 million routed through fragmented super PACs structured to stay below the reporting limits that would trigger public disclosure, and targeted state legislative campaign spending.

The centerpiece of the operation is the Digital Childhood Alliance, which launched on December 18, 2024, with a professional website fully loaded with statistics, testimonials from Heritage Foundation and NCOSE staff, and talking points for the App Store Accountability Act. Within weeks of its launch, DCA was testifying before the Utah state legislature in support of SB-142, which became the first ASAA law signed in the country 77 days later. Bloomberg reported in July 2025, through three sources familiar with the funding, that Meta was bankrolling the DCA. Under oath at a Louisiana Senate committee hearing, DCA Executive Director Casey Stefanski admitted receiving tech company funding but refused to name donors.

The legislative target is the App Store Accountability Act, which would require Apple and Google to verify every user's age at the app store level before any download. Under the proposed framework, app stores bear the full infrastructure and compliance cost of building age verification systems. Social media platforms like Meta's receive the age data through an API (basically a data pipeline between two pieces of software) without having to build or maintain verification systems of their own. The bills also contain safe harbor provisions for app developers: if a developer relies in good faith on age data provided by the app store, that developer is shielded from liability under this Act. The practical effect is that if Apple or Google misidentify a minor as an adult, Meta faces no consequences under these specific laws, though liability under other statutes like COPPA would still apply. Google spokesperson Danielle Cohen told Bloomberg: "We see the legislation being pushed by Meta as an effort to offload their own responsibilities to keep kids safe."

The DCA's legal footprint tells its own story. It's registered in Delaware and reports gross receipts under \$25,000 for tax year 2024. It files the IRS form reserved for the smallest nonprofits, one that requires zero financial disclosure. That figure is difficult to square with an organization that employs a paid executive director, retains a registered lobbyist, commissions polling, and coordinates testimony programs across more than 20 states. The TBOTE investigation concludes the real operating budget never touches DCA's own EIN. Instead, funding flows through intermediary nonprofit structures that aren't subject to standard political ad disclosure, keeping Meta's fingerprints off the paperwork.

Utah's version of the law, the first to be signed, hits its compliance deadline on May 6, 2026. Texas enacted its own version, though a federal court issued a preliminary injunction blocking enforcement on First Amendment grounds. Louisiana's version takes effect July 1, 2026. Alabama signed its ASAA into law on March 9, 2026. At the federal level, Senator Mike Lee and Representative John James introduced the App Store Accountability Act (S.1586 / H.R. 3149), which would preempt the state patchwork with a single national standard enforced by the FTC. Roughly 17 additional states have introduced or are considering similar bills. The direction is clear, and Meta's lobbying machine is pushing the door open in every statehouse it can reach.

Meanwhile, the EU's eIDAS 2.0 framework is building a fundamentally different model. The European Digital Identity Wallet, now being piloted in five member states, supports privacy-preserving age verification using zero-knowledge proofs, which allow a user to prove they're over 18 without revealing their birthdate or identity. The data stays on the user's device. It's open-source, decentralized, and targets large platforms while exempting small entities and open-source projects. The contrast with Meta's preferred approach could not be sharper: one model puts the user in control, the other builds a permanent identity layer into operating systems and hands the data to the companies that already have too much of it.

Sources:

- GitHub (upper-up / TBOTE Project), "Meta Lobbying and Other Findings," updated Mar 2026: "Meta spent a record \$26.3 million on federal lobbying in 2025, deployed 86+ lobbyists across 45 states, and covertly funded a 'grassroots' child safety group called the Digital Childhood Alliance (DCA) to advocate for the App Store Accountability Act (ASAA)." DCA EIN 33-2669790, Delaware, tax year 2024 gross receipts under \$25,000. — <https://github.com/upper-up/meta-lobbying-and-other-findings>
- The Register, "Age checks creep into Linux as systemd gets a DOB field," Mar 24, 2026: "TBOTE claims it has directly traced more than \$25 million, and that Meta could have spent upward of \$2 billion on this over the last year." — https://www.theregister.com/2026/03/24/foss_age_verification/
- Bloomberg (via Insurance Journal), "Meta Clashes With Apple, Google Over Child Age Check Legislation," Jul 25, 2025: Google spokesperson Danielle Cohen: "We see the legislation being pushed by Meta as an effort to offload their own responsibilities to keep kids safe." Bloomberg confirmed Meta is funding DCA "according to three people familiar with the funding." —

<https://www.insurancejournal.com/news/national/2025/07/25/833246.htm>

- Rappler, "Meta's age verification lobbying may hurt app stores... and everyone else," Mar 2026: "The \$26.3 million spent by Meta for this lobbying in 2025 in the US was higher than the spending of Lockheed Martin or Boeing." —
<https://www.rappler.com/technology/features/tech-thoughts-meta-age-verification-lobbying-effects-harms/>
- Stoel Rives, "Utah's App Store Accountability Act Goes into Effect," 2025: "App store providers and developers have until May 6, 2026 to comply with the Act's obligations." —
<https://www.stoel.com/insights/publications/utahs-app-store-accountability-act-goes-into-effect>
- U.S. Congress, App Store Accountability Act, S.1586, 119th Congress: Safe harbor provision (Sec. 8): an app developer "is deemed not liable for a violation of this Act if the app developer demonstrates it has relied in good faith on age verification data provided by a covered app store provider." —
<https://www.congress.gov/bill/119th-congress/senate-bill/1586/text>
- Senator Mike Lee, "Lee Introduces Bill to Protect Children Online, Hold App Stores Accountable," May 1, 2025: "101 advocate organizations for children's online safety endorsed the App Store Accountability Act in a letter of support led by the Digital Childhood Alliance." —
<https://www.lee.senate.gov/2025/5/lee-introduces-bill-to-protect-children-online-hold-app-stores-accountable>
- Morrison Foerster, "Update: Federal Court Enjoins Texas App Store Accountability Act," Jan 9, 2026: Court concluded the Act "likely violates the First Amendment because it imposes content-based restrictions on speech." —
<https://www.mofo.com/resources/insights/251111-texas-targets-app-stores-with-new-accountability-law>
- Regulatory Oversight, "Alabama Enacts App Store Accountability Act," Mar 26, 2026: "On March 9, Governor Kay Ivey signed House Bill 161, the App Store Accountability Act, into law." —
<https://www.regulatoryoversight.com/2026/03/alabama-enacts-app-store-accountability-act-requiring-age-verification-and-parental-consent/>
- International Review of Law, Computers & Technology, "Regulatory options for integrating zero-knowledge proofs into the European Digital Identity Wallet," 2024: eIDAS 2.0 "introduces zero-knowledge proof (ZKP) privacy-enhancing technologies" to support the EUDIW. —
<https://www.tandfonline.com/doi/abs/10.1080/13600869.2024.2398915>
- European Commission, "eIDAS 2.0 / European Digital Identity Framework," 2024: EU Digital Identity Wallets allow users to "prove a specific personal attribute, such as age, without revealing their full identity or other personal details." —
<https://www.european-digital-identity-regulation.com/>

3. Nvidia DLSS 5: AI Slop Meets Gamer Rage

Nvidia unveiled DLSS 5 (Deep Learning Super Sampling, its AI-powered graphics technology) at GTC, Nvidia's annual technology conference, on March 16, pitching it as a breakthrough in 'neural rendering.' CEO Jensen Huang called it "the GPT moment for graphics," claiming the technology blends traditional rendering with generative AI to deliver a "dramatic leap in visual realism" while keeping creative control in developers' hands. Unlike previous DLSS versions, which focused on upscaling resolution and generating frames for performance, DLSS 5 uses an AI model to infer and reconstruct lighting, materials, and surface detail in real time. Nvidia framed this as the biggest advance in computer graphics since real-time ray tracing in 2018.

Gamers did not agree. The official DLSS 5 reveal trailer on Nvidia's GeForce YouTube channel racked up over a million views and an 84% dislike ratio, with roughly 82,500 thumbs-down against 16,100 thumbs-up, according to browser extensions that track hidden dislike counts. Every individual game demo video followed the same pattern. Resident Evil Requiem sat at 14.9% approval. EA Sports FC landed at 14.5%. Only the Zorah Unreal Engine tech demo broke out of the teens, and even that only managed 37%.

The face of the backlash was Grace Ashcroft from Capcom's Resident Evil Requiem. With DLSS 5 enabled, Grace's character model looked hollowed-out, smoothed, and altered in ways that gamers immediately compared to Instagram beauty filters. Darker hair roots appeared. Makeup materialized on her face. The internet coined a term for it: "yassified." Memes flooded every corner of the gaming internet. One especially popular YouTube comment read, "We went from raytracing to sloptracing."

Huang fired back. At a press Q&A with Tom's Hardware on March 17, he dismissed the backlash entirely: "Well, first of all, they're completely wrong." He argued that DLSS 5 "fuses controllability of the geometry and textures and everything about the game with generative AI" and insisted the technology operates at the geometry level, not as a post-processing filter. "It's not post-processing, it's not post-processing at the frame level, it's generative control at the geometry level," he said.

That framing started falling apart within days. YouTuber Daniel Owen published an email exchange with Jacob Freeman, a marketing specialist at Nvidia. Owen asked Freeman directly whether DLSS 5 is effectively taking a single 2D screenshot and some data about how objects are moving between frames to create the output. Freeman's reply: 'Yes, DLSS5 takes a 2D frame plus motion vectors as an input.' He confirmed that the system "understand[s] complex scene semantics" by "analyzing a single frame," with no access to 3D geometry data, scene depth, or material surface data from the engine. Owen's assessment of the exchange was blunt: DLSS 5 is taking a screenshot of the game and running it through a generative AI model. Kotaku reported the contradiction between Huang's "geometry level" claims and Freeman's

confirmation that the system works exclusively with 2D frame data, noting it wouldn't be the first time Huang was accused of misleading consumers about Nvidia products.

The developer reaction was equally damning. Insider Gaming reported that developers at both Capcom and Ubisoft, two studios whose games were featured in the DLSS 5 showcase, learned about the announcement at the same time as the public. "We found out at the same time as the public," said one Ubisoft developer. Capcom developers were reportedly shocked, given the company's historically "anti-AI" stance. On March 23, Capcom published an investor Q&A summary in which it stated plainly: "Our company will not be implementing any AI-generated assets into our video game content." Capcom said it would use generative AI internally for development efficiency but drew a clear line at final game assets. That policy adds an awkward layer to the DLSS 5 reveal, where Capcom executive Jun Takeuchi was quoted in Nvidia's own press release praising the technology's potential for Resident Evil. Bethesda's Todd Howard was even more vocal, calling the technology's results in Starfield "amazing" and saying "the artistic style and detail shine through without being held back by the traditional limits of real-time rendering." Howard's enthusiasm reads differently when you consider that Starfield was one of the demos where DLSS 5 was caught hallucinating facial details that don't exist in the original character models.

Digital Foundry, one of the most respected names in technical analysis, initially ran a positive hands-on piece with the technology, and immediately came under fire. The outlet released a follow-up Q&A video in which founder Richard Leadbetter admitted they "should have taken more time." Team member Alex Battaglia went further, saying that DLSS 5 appears to "trample on artistic vision in a very hardcore way." Battaglia pointed out that the system only has access to 2D data and no specialized 3D face scans, meaning it averages results based on training data. On the ethical side, Battaglia said the alteration of characters was "very problematic because the actress probably signed off on her likeness to a certain degree."

By March 23, Huang was on the Lex Fridman podcast, and the tone had shifted. "I think their perspective makes sense, and I can see where they're coming from, because I don't love AI slop myself," he said. "All of the AI-generated content increasingly looks similar, and they're all beautiful... so I'm empathetic toward what they're thinking." But the concession only went so far. Huang still maintained that DLSS 5 is "3D conditioned, 3D guided" and that it preserves artist intent. He described it as an optional tool where "the artist determines the geometry" and "all of that is done for the artist, so that they can create something that is more beautiful but still in the style that they want."

DLSS 5 is not shipping until fall 2026. Nvidia has said integration will use its existing Streamline framework (the same toolset DLSS and Reflex currently use), but no public developer tools or documentation have been released. The company has announced 16-plus launch titles across major publishers, though the blindsided developer reaction at Capcom and Ubisoft raises real questions about how much "artistic control" studios will actually exercise. DLSS as a technology has been integrated into over 750 games since its 2018 debut, and previous versions faced their own skepticism before gaining widespread adoption. Whether DLSS 5 follows the same

arc or becomes the first version gamers actively refuse to accept is an open question. The backlash shows no sign of cooling.

Sources:

- Nvidia Newsroom, "NVIDIA DLSS 5 Delivers AI-Powered Breakthrough in Visual Fidelity for Games," Mar 16, 2026: Jensen Huang: "DLSS 5 is the GPT moment for graphics." DLSS described as "a real-time neural rendering model that infuses pixels with photoreal lighting and materials." Arriving "this fall." — <https://nvidianews.nvidia.com/news/nvidia-dlss-5-delivers-ai-powered-breakthrough-in-visual-fidelity-for-games>
- Nvidia GeForce News, "NVIDIA DLSS 5 Delivers AI-Powered Breakthrough In Visual Fidelity For Games," Mar 16, 2026: "NVIDIA DLSS was released in 2018 as an AI technology to boost performance... It has been integrated in over 750 games." DLSS 5 "takes a game's color and motion vectors for each frame as input." — <https://www.nvidia.com/en-us/geforce/news/dlss5-breakthrough-in-visual-fidelity-for-games/>
- VGC, "Nvidia's DLSS 5 reveal trailer only has 16% likes on YouTube," Mar 2026: 16,100 likes against 82,500 dislikes. Resident Evil Requiem demo at 14.9% approval. EA Sports FC at 14.5%. Zorah demo at 37%. — <https://www.videogameschronicle.com/news/nvidias-dlss-5-reveal-trailer-only-has-16-likes-on-youtube-as-players-make-their-voices-heard/>
- Kotaku, "Nvidia's Loathed DLSS 5 AI Is Getting Slaughtered On YouTube," Mar 2026: Dislike ratio of "16 percent likes and 84 percent dislikes." Jensen Huang: "They're completely wrong." — <https://kotaku.com/nvidias-loathed-dlss-5-ai-sloppifier-is-getting-slaughtered-on-youtube-2000680291>
- Tom's Hardware, "Jensen Huang says gamers are 'completely wrong' about DLSS 5," Mar 17, 2026: Huang at GTC press Q&A: "Well, first of all, they're completely wrong." And: "It's not post-processing, it's not post-processing at the frame level, it's generative control at the geometry level." — <https://www.tomshardware.com/pc-components/gpus/jensen-huang-says-gamers-are-completely-wrong-about-dlss-5-nvidia-ceo-responds-to-dlss-5-backlash>
- Kotaku, "Nvidia CEO's Defense Of DLSS 5 Gets Contradicted By One Of His Employees," Mar 2026: Jacob Freeman confirmed DLSS 5 "takes a 2D frame plus motion vectors as an input" with no access to 3D geometry. Daniel Owen's assessment: "DLSS 5 is just using a screenshot and slapping a filter over top." — <https://kotaku.com/dlss-5-nvidia-ai-slop-faces-filter-ceo-jensen-huang-2d-2000680670>
- GamesRadar, "'The underlying geometry is unchanged' — turns out DLSS 5 really is just a filter over 2D images," Mar 2026: Freeman confirmed via email: "DLSS 5 only takes the rendered frame and motion vectors as inputs. Materials are inferred from the rendered frame." — <https://www.gamesradar.com/hardware/desktop-pc/the-underlying-geometry-is-unchanged/>

[ed-turns-out-dlss-5-really-is-just-a-filter-over-2d-images-as-nvidia-employee-reveals-materials-are-inferred-from-the-rendered-frame/](#)

- Insider Gaming, "NVIDIA CEO Says Gamers Are 'Completely Wrong' About DLSS 5," Mar 2026: Developers at Capcom and Ubisoft told Insider Gaming the DLSS 5 tech "was revealed to them at the same time as everyone else." Ubisoft developer: "We found out at the same time as the public." Capcom developers were "shocked" given the company's historically "anti-AI" stance. —
<https://insider-gaming.com/dlss-5-gamers-are-wrong/>
- PC Gamer, "Capcom says it 'will not implement assets generated by AI into our game content,'" Mar 23, 2026: Capcom investor Q&A: "We will not incorporate assets generated by AI into our game content." Will use AI for development efficiency. —
<https://www.pcgamer.com/gaming-industry/capcom-says-it-will-not-implement-assets-generated-by-ai-into-our-game-content-but-still-plans-to-use-ai-to-enhance-efficiency-and-boost-productivity-in-game-development/>
- NotebookCheck, "Digital Foundry's Alex Battaglia calls DLSS 5 out for messing with artistic vision and serious ethical concerns," Mar 19, 2026: Battaglia said DLSS 5 "trample[s] on artistic vision in a very hardcore way." Called the alteration of character likenesses "very problematic." —
<https://www.notebookcheck.net/Digital-Foundry-s-Alex-Battaglia-calls-DLSS-5-out-for-messing-with-artistic-vision-and-serious-ethical-concerns.1253527.0.html>
- Kotaku, "Digital Foundry On DLSS 5 Blowback: 'We Should Have Taken More Time,'" Mar 2026: Alex Battaglia: "Ethically, I find that very problematic because the actress probably signed off on her likeness to a certain degree." Richard Leadbetter admitted the team published too quickly. —
<https://kotaku.com/digital-foundry-dlss5-nvidia-genai-response-video-2000680002>
- Tom's Hardware, "Nvidia CEO says he's 'empathetic' to DLSS 5 concerns," Mar 23, 2026: Huang on Lex Fridman Podcast #494: "I think their perspective makes sense, and I could see where they're coming from because I don't love AI slop myself." —
<https://www.tomshardware.com/pc-components/gpus/nvidia-ceo-says-hes-empathic-to-dlss-5-concerns-jensen-huang-doubles-down-on-defense-while-decrying-ai-slop>
- GameSpot, "Nvidia CEO Hates AI Slop Too, Promises DLSS 5 Isn't Making More Despite The Evidence," Mar 23, 2026: Freeman confirmed to Daniel Owen that DLSS 5 has no "access to the underlying 3D geometry data present in a game's engine." Huang on podcast: "All of that is done for the artist, so that they can create something that is more beautiful, but still in the style that they want." —
<https://www.gamespot.com/articles/nvidia-ceo-hates-ai-slop-too-promises-dlss-5-isnt-making-more-despite-the-evidence/1100-6538958/>
- Winbuzzer, "NVIDIA CEO Walks Back DLSS 5 Criticism Dismissal Amid AI Slop Backlash," Mar 24, 2026: "A Grace Ashcroft screenshot from Resident Evil Requiem became the face of the controversy." DLSS integrated into "over 750 games since its 2018 launch." —
<https://winbuzzer.com/2026/03/24/nvidia-ceo-backtracks-dlss-5-criticism-ai-slop-xcxwbn/>

- Windows Central, "Bethesda's Todd Howard says NVIDIA DLSS 5 in Xbox's Starfield is 'amazing,'" Mar 2026: Howard praised DLSS 5 results in Starfield. Article notes the technology "makes some aggressive changes to character skin and facial appearance" and compares the results to "generative AI beautification filters you'll find in smartphone apps." —
<https://www.windowscentral.com/gaming/pc-gaming/bethesdas-todd-howard-says-nvidia-dlss-5-in-xboxs-starfield-is-amazing-everyone-else-is-calling-it-uncanny-ai-slop>

4. Anthropic vs. The Pentagon

In July 2025, Anthropic signed a contract worth up to \$200 million with the Pentagon through the Department of Defense's Chief Digital and Artificial Intelligence Office. It became the first AI lab to deploy its technology across the military's classified networks. For a company built on the premise of AI safety, the deal represented a calculated bet: work with the military, but set boundaries.

Anthropic drew two lines it would not cross. Claude would not be used for fully autonomous weapons systems or for mass domestic surveillance of American citizens. When deployment talks began on the DOD's internal AI platform, GenAI.mil, in September, those two restrictions became the sticking point. Talks stalled. By late February 2026, the standoff had turned into an ultimatum. Defense Secretary Pete Hegseth told Anthropic CEO Dario Amodei to grant the Pentagon access to Claude "for all lawful purposes" or face termination of the contract and worse. Amodei refused. In a public statement, he said the company "cannot in good conscience accede to their request."

On February 27, Trump posted on Truth Social ordering every federal agency to "immediately cease" all use of Anthropic's technology. He gave agencies six months to phase it out. "WE will decide the fate of our Country," Trump wrote, "NOT some out-of-control, Radical Left AI company run by people who have no idea what the real World is all about." Hours later, Hegseth designated Anthropic a "supply chain risk," a classification under 10 U.S.C. § 3252 that is normally reserved for companies tied to foreign adversaries. Think Huawei. Think Chinese state-owned telecoms. Not a San Francisco AI startup founded by former OpenAI researchers.

It was the first time the United States government had ever applied this designation to an American company. The practical consequences were severe: if the label stood, every defense contractor doing business with the Pentagon, including Amazon, Microsoft, and Palantir, would be required to certify that they do not use Claude in any military-related work. The fallout wouldn't stop at the \$200 million contract. It would ripple through Anthropic's entire commercial customer base.

On March 9, Anthropic filed two federal lawsuits. One in the U.S. District Court for the Northern District of California, one in the D.C. Circuit Court of Appeals. The company called the designation "unprecedented and unlawful," arguing it violated both First Amendment protections and the company's Fifth Amendment right to due process. The ACLU and the Center for Democracy & Technology filed a supporting brief in the D.C. Circuit case. Patrick Toomey, deputy director of the ACLU's National Security Project, put it this way: "AI-powered surveillance poses immense dangers to our democracy. Anthropic's public advocacy for AI guardrails is laudable and protected by the First Amendment, not something the Pentagon should be punishing."

Meanwhile, something remarkable happened across Silicon Valley. Over 100 Google employees working on AI technology sent a letter to Jeff Dean, chief scientist of Google DeepMind, requesting that the company prohibit the military from using Gemini for surveillance of Americans or autonomous weapons without human oversight. The letter explicitly cited the Anthropic dispute. Separately, nearly 900 employees across Google and OpenAI signed a public open letter titled "We Will Not Be Divided," urging their companies to hold the same lines Anthropic was being punished for drawing. OpenAI CEO Sam Altman then signed a deal with the Pentagon anyway, though he publicly claimed it contained similar restrictions.

Anthropic had also been making moves in the political arena. On February 12, the company donated \$20 million to Public First Action, a bipartisan nonprofit advocacy group backing candidates from both parties who support AI oversight and safeguards in the 2026 midterms. The donation put Anthropic directly at odds with Leading the Future, a rival super PAC backed by OpenAI co-founder Greg Brockman and Andreessen Horowitz that had already raised \$125 million to support lighter AI regulation.

The case went before U.S. District Judge Rita Lin in San Francisco on March 24. During a 90-minute hearing, Lin did not hide her skepticism. "I don't know if it's murder," she said, "but it looks like an attempt to cripple Anthropic." She questioned whether the company was being "punished for criticizing the government's contracting position in the press." Eric Hamilton, the government's lawyer, argued that the DOD had "come to worry that Anthropic may in the future take action to sabotage or subvert IT systems." Anthropic's attorney Michael Mongan responded: "This is something that has never been done with respect to an American company. It is a very narrow authority. It doesn't apply here."

The backdrop made the whole thing even more surreal. On the same day as the hearing, Pentagon Chief Information Officer Kirsten A. Davies confirmed during a Senate Armed Services Committee appearance that Claude is actively being used in the ongoing military conflict with Iran. The Pentagon designated Anthropic a national security threat while simultaneously relying on the company's technology in combat. Sen. Jack Reed asked Davies if that struck her as "odd." She deflected, saying the six-month phase-out period was "reasonable."

Without the injunction, Anthropic said in filings it could lose billions in business. On March 26, Judge Lin granted the preliminary injunction, blocking both the supply chain risk designation and Trump's directive. "Punishing Anthropic for bringing public scrutiny to the government's contracting position is classic illegal First Amendment retaliation," she wrote. The Pentagon has signaled it will appeal.

Sources:

- Anthropic, "Anthropic and the Department of Defense to Advance Responsible AI in Defense Operations," Jul 14, 2025: "The U.S. Department of Defense (DOD), through its Chief Digital and Artificial Intelligence Office (CDAO), has awarded Anthropic a two-year

prototype other transaction agreement with a \$200 million ceiling." —

<https://www.anthropic.com/news/anthropic-and-the-department-of-defense-to-advance-responsible-ai-in-defense-operations>

- NPR, "OpenAI announces Pentagon deal after Trump bans Anthropic," Feb 28, 2026: "President Trump ordered the U.S. government to stop using the artificial intelligence company Anthropic's products and the Pentagon moved to designate the company a national security risk." —
<https://www.npr.org/2026/02/27/nx-s1-5729118/trump-anthropic-pentagon-openai-ai-weapons-ban>
- CNN, "Anthropic rejects latest Pentagon offer," Feb 26, 2026: Amodai: "In a narrow set of cases, we believe AI can undermine, rather than defend, democratic values." —
<https://www.cnn.com/2026/02/26/tech/anthropic-rejects-pentagon-offer>
- Fortune, "OpenAI sweeps in to snag Pentagon contract after Anthropic labeled 'supply chain risk' in unprecedented move," Feb 28, 2026: "It is the first time the U.S. has ever designated an American company a supply chain risk, and the first time the designation has been used in apparent retaliation for a business not agreeing to certain contract terms." —
<https://fortune.com/2026/02/28/openai-pentagon-deal-anthropic-designated-supply-chain-risk-unprecedented-action-damage-its-growth/>
- CNN, "Anthropic sues the Trump administration after it was designated a supply chain risk," Mar 9, 2026: "Anthropic alleges that its categorization as a supply chain risk — and the Trump administration's directive that federal agencies stop using the company's technology — are legally unsound." —
<https://www.cnn.com/2026/03/09/tech/anthropic-sues-pentagon>
- NPR, "Anthropic sues Trump administration over blacklisting decision," Mar 9, 2026: "The lawsuits, filed in the U.S. District Court for the Northern District of California and the federal appeals court in Washington, D.C., allege the Trump administration violated the company's First Amendment rights and exceeded the scope of supply chain risk law." —
<https://www.npr.org/2026/03/09/nx-s1-5742548/anthropic-pentagon-lawsuit-amodai-hegets>
- ACLU, "ACLU and CDT Urge Court to Stop Government from Punishing Anthropic," Mar 2026: Patrick Toomey, deputy director of the ACLU's National Security Project: "Anthropic's public advocacy for AI guardrails is laudable and protected by the First Amendment — not something the Pentagon should be punishing." —
<https://www.aclu.org/press-releases/aclu-and-cdt-urge-court-to-stop-government-from-punishing-anthropic-for-important-advocacy-on-ai-guardrails>
- CNBC, "Judge presses DOD on why Anthropic was blacklisted," Mar 24, 2026: Eric Hamilton, government lawyer: The DOD had "come to worry that Anthropic may in the future take action to sabotage or subvert IT systems." Anthropic attorney Michael Mongan: "This is something that has never been done with respect to American company." —
<https://www.cnbc.com/2026/03/24/anthropic-lawsuit-pentagon-supply-chain-risk-claude.html>

- MS Now, "Federal judge calls Pentagon's ban of Anthropic 'troubling,'" Mar 24, 2026: Judge Lin: "I don't know if it's murder, but it looks like an attempt to cripple Anthropic." — <https://www.ms.now/news/anthropic-takes-fight-against-pentagon-to-federal-court>
- The Hill, "Anthropic's Claude used by Pentagon in war with Iran, official confirms," Mar 24, 2026: Pentagon CIO Kirsten A. Davies: "The use of the system is active right now." — <https://thehill.com/policy/defense/5799136-claude-pentagon-iran-war/>
- The New York Times (as cited by multiple outlets including CNBC, Mar 3, 2026), "Google employees call for military limits on AI amid Iran strikes, Anthropic fallout": Over 100 Google employees sent a letter to Jeff Dean requesting the company prohibit military use of Gemini for mass surveillance and autonomous weapons. — <https://www.cnbc.com/2026/03/03/anthropic-fallout-iran-war-tech-military-ai.html>
- CNBC, "Anthropic gives \$20 million to group pushing for AI regulations ahead of 2026 elections," Feb 12, 2026: "Anthropic, the artificial intelligence lab that's taken heat from the White House for its support of regulations and safety, is putting \$20 million into the political arena ahead of the 2026 elections." — <https://www.cnbc.com/2026/02/12/anthropic-gives-20-million-to-group-pushing-for-ai-regulations-.html>
- Anthropic, "Anthropic is donating \$20 million to Public First Action," Feb 12, 2026: "Anthropic is contributing \$20 million to Public First Action, a new bipartisan 501(c)(4) that will support public education about AI, promote safeguards, and ensure America leads the AI race." — <https://www.anthropic.com/news/donate-public-first-action>
- CNBC, "Judge issues preliminary injunction in Anthropic's fight with the Pentagon," Mar 26, 2026: "Punishing Anthropic for bringing public scrutiny to the government's contracting position is classic illegal First Amendment retaliation," Lin wrote. — <https://www.cnbc.com/2026/03/26/anthropic-pentagon-dod-claude-court-ruling.html>
- CNN, "Judge blocks Pentagon's effort to 'punish' Anthropic," Mar 26, 2026: "Nothing in the governing statute supports the Orwellian notion that an American company may be branded a potential adversary and saboteur of the U.S. for expressing disagreement with the government." — <https://www.cnn.com/2026/03/26/business/anthropic-pentagon-injunction-supply-chain-risk>

5. Microsoft's "Microslop" Retreat

On March 20, Windows chief Pavan Davuluri published a blog post titled "Our commitment to Windows quality," and if you've spent any time using Windows 11 over the past year, the list of promised changes reads like a user wish list that Microsoft ignored for three straight years and suddenly discovered in a desk drawer.

The headline: Microsoft is reducing Copilot integrations across the OS, starting with Snipping Tool, Photos, Widgets, and Notepad. "You will see us be more intentional about how and where Copilot integrates across Windows, focusing on experiences that are genuinely useful and well-crafted," Davuluri wrote. In the same post: "We are reducing unnecessary Copilot entry points." That language is doing a lot of heavy lifting. It implies that the previous approach, cramming Copilot into every corner of the operating system whether it belonged there or not, was somehow unintentional. Nobody who watched Microsoft bolt an AI button onto Notepad believes that.

Here's the backstory. Back in late December 2025, CEO Satya Nadella published a year-end blog post asking the tech industry to get past "the arguments of slop vs sophistication" when talking about AI. The internet took that as a direct challenge. Within days, the term "Microslop" exploded across X, Reddit, Instagram, and just about every tech forum that exists. The name itself had floated around since the mid-2000s, but Nadella's plea turned it into the most popular tech meme of early 2026. Browser extensions that replaced every mention of "Microsoft" with "Microslop" racked up over 180,000 downloads. Microsoft eventually banned the word on its own Copilot Discord server, which predictably triggered a raid that forced the server offline entirely by March 1. Classic Streisand Effect.

That backlash clearly reached Redmond. Davuluri's blog post goes well beyond Copilot. He announced the return of the movable taskbar, letting users position it on the top or sides of their screen. That's a feature Windows offered 30 years ago in Windows 95 and inexplicably removed for Windows 11. He promised reduced baseline RAM usage, so the OS itself would consume less memory at idle. He committed to a faster and more dependable File Explorer with lower latency for search and navigation. He pledged that Windows Update would become less disruptive, with the ability to pause updates as long as you need, skip them during setup, and shut down without being forced to install. He announced plans to reduce promotional upsells and make the OS less cluttered overall. And he outlined an overhaul of the Windows Insider Program to make feedback loops more transparent, with an upgraded Feedback Hub so users can actually see how their input shapes Windows.

Windows Central summed it up: "today's announcement almost reads like an apology letter, just without the actual apology." Windows Latest went further, calling it "the biggest thing that has ever happened to Windows 11." The Register was characteristically blunt, pointing out that the blog was "long on promises that things will get better, but short on words like 'sorry,' 'apologize,'

or even the Americanism 'our bad.'"

Windows Latest also reported that Microsoft may drop the Copilot branding from some AI features entirely and shift toward performance-focused development instead. That tracks with earlier reporting from Windows Central that several planned Copilot integrations for Settings and File Explorer were quietly shelved, and that features which did return came back without the Copilot name attached.

The initial wave of changes is rolling out to Insider builds through March and April 2026. Scott Hanselman, VP at Microsoft, confirmed on X that updates will arrive "every month this year," with broader availability throughout 2026. Whether Microsoft actually follows through is the real question. Davuluri made similar pledges in January. Before that, in November 2025, he posted on X about Windows "evolving into an agentic OS" and got so much backlash he had to disable replies. The company has been here before. But between the Microslop meme, users fleeing to macOS and Linux, and Windows 11's reputation sitting at what multiple outlets describe as an all-time low, Microsoft may finally be running out of second chances.

Sources:

- Microsoft / Windows Insider Blog, "Our commitment to Windows quality," Mar 20, 2026: "You will see us be more intentional about how and where Copilot integrates across Windows, focusing on experiences that are genuinely useful and well-crafted. As part of this, we are reducing unnecessary Copilot entry points, starting with apps like Snipping Tool, Photos, Widgets and Notepad." — <https://blogs.windows.com/windows-insider/2026/03/20/our-commitment-to-windows-quality/>
- The Register, "Windows boss promises to heal the operating system's wounds," Mar 24, 2026: Davuluri's post was "long on promises that things will get better, but short on words like 'sorry,' 'apologize,' or even the Americanism 'our bad.'" — https://www.theregister.com/2026/03/24/windows_boss_promises/
- Windows Central, "Microsoft unveils MAJOR improvements coming to Windows 11 this year," Mar 20, 2026: "today's announcement almost reads like an apology letter, just without the actual apology." — <https://www.windowcentral.com/microsoft/windows-11/windows-11-major-improvements-announced-movable-taskbar-less-ads-reduced-copilot-better-performance-2026>
- Windows Latest, "Microsoft confirms a major Windows 11 update with faster Explorer, reduced Copilot, and more," Mar 21, 2026: "this is the biggest thing that has ever happened to Windows 11." — <https://www.windowlatest.com/2026/03/21/microsoft-confirms-a-major-windows-11-update-with-faster-explorer-reduced-copilot-and-more/>
- Windows Latest, "Microsoft responds to Microslop criticism by scaling back Copilot in Windows 11," Mar 21, 2026: "Windows Latest understands that Microsoft might also drop Copilot branding from some AI features and focus more on performance." — <https://www.windowlatest.com/2026/03/21/microsoft-responds-to-microslop-criticism-by->

[scaling-back-copilot-in-windows-11-starting-with-notepad-and-other-apps/](#)

- TechCrunch, "Microsoft rolls back some of its Copilot AI bloat on Windows," Mar 20, 2026: "Pavan Davuluri, EVP of Windows and Devices, wrote on the company's blog that Microsoft is becoming more intentional about 'how and where Copilot integrates across Windows.'" — <https://techcrunch.com/2026/03/20/microsoft-rolls-back-some-of-its-copilot-ai-bloat-on-windows/>
- Know Your Meme, "Microslop," updated Jan 7, 2026: "the term primarily saw widespread adoption in late December 2025 and early January 2026 as a way to mock Microsoft for its AI push in the mid-2020s." — <https://knowyourmeme.com/memes/microslop>
- Windows Central, "'Microslop' trends on social media — backlash to Microsoft's on-going AI obsession continues," Jan 3, 2026: "'Microslop' began trending on X yesterday, in the wake of our coverage on the topic." — <https://www.windowcentral.com/artificial-intelligence/microslop-trends-on-social-media-backlash-to-microsofts-on-going-ai-obsession-continues>
- Windows Latest, "Microsoft gets tired of 'Microslop,' bans the word on its Discord, then locks the server after backlash," Mar 2, 2026: "the word 'Microslop' was actively filtered in the official Microsoft Copilot Discord server." — <https://www.windowlatest.com/2026/03/02/microsoft-gets-tired-of-microslop-bans-the-word-on-its-discord-then-locks-the-server-after-backlash/>
- Winbuzzer, "Microsoft Pledges Windows 11 Overhaul With Copilot Pullback, Movable Taskbar," Mar 22, 2026: "Microsoft made similar quality commitments in September 2025 and January 2026, raising questions about whether this round of promises will deliver results." — <https://winbuzzer.com/2026/03/22/microsoft-pledges-windows-11-quality-overhaul-copilot-pullback-taskbar-cxwbn/>

6. FCC Bans All Foreign-Made Routers

On March 23, the FCC added every consumer-grade router produced in a foreign country to its Covered List. That single move blocks any new foreign-made router model from receiving FCC equipment authorization, and without that authorization, a router cannot legally be imported, marketed, or sold in the United States. The action followed a March 20 National Security Determination from a White House-convened interagency body, which concluded that foreign-produced routers "pose unacceptable risks to the national security of the United States." The determination cited the Volt, Flax, and Salt Typhoon cyberattack campaigns, all attributed to Chinese state actors, as evidence that foreign-made routers had been weaponized against American infrastructure.

FCC Chair Brendan Carr framed the action as following "President Trump's leadership" on cybersecurity. The determination itself was blunt: "From disrupting network connectivity to enabling local networking espionage and intellectual property theft, foreign-produced routers present unacceptable risks to Americans." That language came directly from the National Security Determination the FCC published alongside its ruling.

Here is the problem with framing this as a targeted security action: virtually every consumer router sold in America is made overseas. TP-Link, Asus, Netgear, D-Link, Linksys, Amazon's Eero, and Google's Nest Wifi all manufacture their hardware outside the U.S., primarily in Taiwan, Vietnam, Thailand, and China. Not a single major consumer router brand currently manufactures entirely domestically. The only notable exception is some Starlink routers produced in Texas, and those ship as part of a satellite internet kit, not as standalone consumer networking gear. A TP-Link spokesperson told WIRED: "Virtually all routers are made outside the United States, including those produced by U.S.-based companies like TP-Link, which manufactures its products in Vietnam."

Some government officials, including former NSA cybersecurity director Rob Joyce in congressional testimony, have estimated that Chinese manufacturers control roughly 60% of the U.S. home router market. TP-Link has contested this figure, citing Circana market research that puts its retail share at about 37% by unit volume. Either way, the ban does not target Chinese manufacturers specifically. It targets all routers produced in any foreign country, period.

Companies can apply for "Conditional Approval" from the Department of War or the Department of Homeland Security to get their products exempted. The process is not light. Applicants must disclose their full management structures and supply chain details, provide a detailed bill of materials with country of origin for every component, and submit a concrete, time-bound plan for moving manufacturing to the United States, complete with quarterly progress reports. This is not a security audit. It is an onshoring program with a security label on it.

The FCC ran the exact same playbook with foreign-made drones in December 2025. The results so far are revealing. Of all the drone applications submitted, exactly four systems have received Conditional Approval as of March 18, 2026: SiFly Aviation (U.S.), Mobilicom (Israel), ScoutDI (Norway), and Verge Aero (U.S.). Every single approval went to a non-Chinese manufacturer. DJI and Autel, the two dominant drone makers, remain fully blocked. DJI is currently suing the FCC in the Ninth Circuit. Former FCC officials told CyberScoop the router ban has similar "big swing" parallels to the drone action, questioning whether it would survive legal challenge or meaningfully address the actual security risk.

And about that security risk: the cyberattack campaigns the FCC cited to justify the ban largely exploited American-brand hardware. Salt Typhoon gained access to U.S. telecom networks by exploiting known, unpatched vulnerabilities in Cisco routers running Cisco's own networking software, including a flaw that had been publicly documented for seven years. Cisco confirmed this in its own disclosure. Flax Typhoon's botnet (a network of hijacked devices controlled remotely) targeted both U.S.-made and foreign-made routers indiscriminately. Multiple cybersecurity researchers and publications, including CSO Online and Dark Reading, have noted that TP-Link's vulnerability track record is comparable to, and in some metrics better than, its American and Taiwanese competitors. CISA, the federal cybersecurity agency, maintains a public list of known exploited security flaws. That list has 2 entries for TP-Link compared to 74 for Cisco and 20 for D-Link.

The practical impact for consumers right now is limited. Existing routers already in homes and offices are completely unaffected. Retailers can continue selling any model that already carries an FCC ID. The FCC's Office of Engineering and Technology issued a waiver permitting manufacturers to push firmware and security updates to previously authorized covered routers through at least March 1, 2027. After that date, the update pathway depends on a forthcoming Department of War determination that has not yet been made public. So if you need a new router, the models on shelves today are still fair game. But the pipeline of new products just froze, and nobody knows when or if it will thaw.

Sources:

- FCC, "FACT SHEET: FCC Updates Covered List to Include Foreign-Made Consumer Routers," Mar 23, 2026: Brendan Carr statement: "Following President Trump's leadership, the FCC will continue do our part in making sure that U.S. cyberspace, critical infrastructure, and supply chains are safe and secure." — <https://docs.fcc.gov/public/attachments/DOC-420034A1.pdf>
- FCC, Public Notice DA-26-278A1, Mar 23, 2026: National Security Determination quote: "From disrupting network connectivity to enabling local networking espionage and intellectual property theft, foreign-produced routers present unacceptable risks to Americans." — <https://docs.fcc.gov/public/attachments/DA-26-278A1.pdf>
- FCC, "FAQs on Recent Updates to FCC Covered List Regarding Routers Produced in Foreign Countries," Mar 23, 2026: "New models of devices on the Covered List are prohibited from receiving FCC authorization and are therefore prohibited from being

imported or sold in the U.S." —

<https://www.fcc.gov/faqs-recent-updates-fcc-covered-list-regarding-routers-produced-for-foreign-countries>

- The Register, "US bans any new consumer-grade routers not made in America," Mar 24, 2026: Carr: "I welcome this Executive Branch national security determination, and I am pleased that the FCC has now added foreign-produced routers, which were found to pose an unacceptable national security risk, to the FCC's Covered List." —
https://www.theregister.com/2026/03/24/fcc_foreign_routers/
- The Record (Recorded Future News), "FCC bans foreign-made routers from US market over 'unacceptable risk,'" Mar 24, 2026: TP-Link spokesperson: "Virtually all routers are made outside the United States, including those produced by U.S.-based companies like TP-Link, which manufactures its products in Vietnam." —
<https://therecord.media/fcc-routers-banned-security-china>
- Consumer Reports, "FCC Bans Foreign-Made Routers, Citing National Security Concerns," Mar 24, 2026: "We're not aware of any major consumer router brand that currently manufactures its products domestically." —
<https://www.consumerreports.org/electronics-computers/wireless-routers/foreign-made-routers-fcc-ban-a1057564057/>
- CyberScoop, "Critics call FCC router rule a 'big swing' that could create more supply chain uncertainty," Mar 24, 2026: "A former FCC official told CyberScoop they were puzzled by the move, and questioned whether it was related to national security or if it would even pass legal muster in the courts." —
<https://cyberscoop.com/fcc-bans-foreign-routers-critics-warn-about-supply-chain/>
- TechCrunch, "FCC bans import of new consumer routers made overseas, citing security risks," Mar 24, 2026: "Salt Typhoon, a China-backed espionage group that has hacked into dozens of phone and internet companies across the world, including in the United States, has been known to exploit vulnerabilities in routers made by American networking giant Cisco." —
<https://techcrunch.com/2026/03/24/fcc-bans-import-of-new-consumer-routers-made-overseas-citing-security-risks/>
- Wiley.law, "FCC Adds Foreign-Produced Consumer-Grade Routers to Covered List," Mar 24, 2026: OET waiver enables firmware updates "at least until March 1, 2027." Notes forthcoming DoW determination referenced in the waiver that has not yet been made public. —
<https://www.wiley.law/alert-FCC-Adds-Foreign-Produced-Consumer-Grade-Routers-to-Covered-List>
- Cooley, "FCC Moves to Prevent New Foreign Routers," Mar 26, 2026: Conditional Approval applicants must provide "a detailed bill of materials and country of origin for all components. A US manufacturing and onshoring plan that is time-bound and overseen by a dedicated officer." —
<https://www.cooley.com/news/insight/2026/2026-03-26-fcc-moves-to-prevent-new-foreign-routers>

- DroneXL, "FCC Exempts Four Foreign Drone Models From Import Ban As DJI Court Battle Continues," Mar 19, 2026: "Every exemption granted since December has gone to a non-Chinese manufacturer." DJI and Autel "remain fully blocked from importing new models or critical components into the US." — <https://dronexl.co/2026/03/19/fcc-drone-ban-exemptions-dji-lawsuit/>
- Nextgov/FCW, "Salt Typhoon hackers exploited stolen credentials and a 7-year-old software flaw in Cisco systems," Feb 20, 2025: "Internet routing giant Cisco disclosed Thursday that the Salt Typhoon hacking group predominantly made use of stolen victim credentials — and in one instance exploited a seven-year-old, known vulnerability in its software." — <https://www.nextgov.com/cybersecurity/2025/02/salt-typhoon-hackers-exploited-stolen-credentials-and-7-year-old-software-flaw-cisco-systems/403146/>
- Dark Reading, "TP-Link Router Ban Is Mostly About Politics," Dec 23, 2024: TP-Link has "two security issues gracing the Known Exploited Vulnerabilities (KEV) list curated by CISA, compared with 74 for Cisco Systems, 23 for Ivanti, and 20 for D-Link." — <https://www.darkreading.com/endpoint-security/us-ban-tp-link-routers-politics-exploitation-risk>
- CSO Online, "No evidence that TP-Link routers are a Chinese security threat," Mar 2025: "Experts say that any valid security qualms over vulnerabilities in TP-Link routers apply equally to all Wi-Fi routers regardless of which company or country manufactures the devices." — <https://www.csoonline.com/article/3504775/no-evidence-that-tp-link-routers-are-a-chinese-security-threat.html>
- TP-Link, "Sets the Record Straight Regarding Inaccurate Testimony," 2025: "Recent market research from Circana identifies TP-Link as holding a 36.6% unit share and a 31% dollar share of the US consumer router market during 2024." — <https://www.tp-link.com/us/press/news/21656/>

7. Age Verification Hits the iPhone: UK First, US Next

Apple released iOS 26.4 on March 25, and UK users opened their Settings app to find something new at the top: a prompt labeled "Confirm You Are 18+." If you have a credit card linked to your Apple account or your account is old enough, the confirmation takes seconds. If not, you'll need to scan a government-issued ID to prove you're an adult. If you decline the prompt entirely, or can't verify, Apple automatically flips on its Web Content Filter and Communication Safety features, restricting your access to certain apps, websites, and content. Apple has framed its approach around data minimization and says the system checks payment methods and account history before ever asking for an ID. That said, the company has not published specifics on what happens to scanned ID data or how long it persists.

The rollout has not gone smoothly. Apple Support forums and Reddit filled with complaints from users whose driving license scans kept failing, whose debit cards were rejected (only credit cards are accepted), and who had no alternative verification method available. 9to5Mac documented the failures the day after launch, and TechRadar ran a piece headlined with a Reddit user quote: "If this starts affecting me using apps, I will switch to Android." Ofcom, the UK's communications regulator, praised the move as "a real win for children and families." But Ofcom also acknowledged that Apple was not technically required to implement OS-level age checks under the UK's Online Safety Act. Apple did this voluntarily, reportedly after close collaboration with Ofcom. In other words, Apple is getting ahead of regulation, not responding to it.

This is not limited to the UK. On February 24, Apple began blocking users in Australia, Brazil, and Singapore from downloading apps rated 18+ unless they can confirm they are adults through the App Store. Utah's App Store Accountability Act takes effect May 6, 2026, requiring app stores to verify user ages using "commercially reasonable methods" before allowing downloads. And in California, the Digital Age Assurance Act (AB 1043), signed by Governor Newsom on October 13, 2025, takes effect January 1, 2027. AB 1043 goes further than any of these. It requires every operating system provider to collect user age data at account setup and transmit it to app developers through a real-time API. Users get sorted into four brackets: under 13, 13 to 15, 16 to 17, and 18 or older. The law's definition of "operating system provider" is anyone who develops, licenses, or controls OS software on a computer, mobile device, or "any other general purpose computing device." That's broad enough to capture SteamOS, niche Linux distros, and anything else running software. AB 1043 passed both California chambers unanimously: 76 to 0 in the Assembly, 38 to 0 in the Senate. Newsom signed it, but flagged concerns in his signing statement about how the law handles shared devices and multi-user accounts.

The Electronic Frontier Foundation published an analysis of AB 1043 in March 2026 that lays out the downstream consequences. Once a developer receives an age-bracket signal, they are "deemed to have actual knowledge" of the user's age range under California law. That phrase is

legally significant. It triggers liability under COPPA (the federal Children's Online Privacy Protection Act) for users under 13, under the CCPA (California's Consumer Privacy Act) for minors generally, and under AB 1043 itself. Penalties run up to \$7,500 per affected child for each intentional violation. The EFF argues this gives developers a strong incentive to simply block anyone flagged as a minor rather than navigate the compliance maze. The EFF's term for it: "outsourced censorship." California sets the mandate; developers do the restricting. And because California's consumer protection laws tend to set the national floor (as the CCPA did for data privacy), AB 1043 will almost certainly shape how OS providers and app developers behave everywhere, not just in California. Apple itself argued against this exact approach in a February 2025 white paper titled "Helping Protect Kids Online," stating that requiring age verification at the app marketplace level is not data minimization. Twelve months later, the company is implementing age verification at the OS level.

The privacy backlash is already measurable. When the UK's Online Safety Act began enforcing age verification requirements on adult websites in July 2025, Proton VPN reported a 1,400% surge in UK signups within minutes. NordVPN confirmed a 1,000% spike in UK purchases over the same period. These numbers reflect what happens when governments mandate identity checks online: a significant portion of users immediately reach for tools that circumvent them. The iOS 26.4 rollout is likely to accelerate that pattern.

The open-source community is pushing back harder than anyone. GrapheneOS, the privacy-focused Android fork that currently runs exclusively on Google Pixel hardware, posted a statement on March 20 refusing to comply with any age verification mandate. Their exact words: "GrapheneOS will remain usable by anyone around the world without requiring personal information, identification, or an account. GrapheneOS and our services will remain available internationally. If GrapheneOS devices can't be sold in a region due to their regulations, so be it." When someone on X asked whether they'd geo-block users via VPN, the team responded: "We don't filter the internet for Iran or North Korea so why would we for Brazil or California?" GrapheneOS announced a long-term partnership with Motorola at MWC 2026, with a dedicated GrapheneOS Motorola phone expected in 2027. Whether that device can ship in jurisdictions that mandate OS-level age checks remains an open question.

GrapheneOS is not alone. The open-source calculator firmware project DB48X issued a legal notice stating that their software "does not, cannot, and will not implement age verification." MidnightBSD, a Unix-based operating system, updated its license to ban use in Brazil entirely rather than comply with that country's age verification law. Garuda Linux, an independent Linux distribution, posted a statement saying it will not implement any age verification measures, since its legal jurisdictions have no laws mandating it. Systemd, a core piece of software that manages how most Linux systems start up and run, merged a code change adding an optional date-of-birth field to its user account data in response to the age verification laws of California, Colorado, and Brazil. Over 400 computer scientists have signed an open letter arguing that these laws build surveillance architecture without meaningfully protecting children, since

self-declaration is trivially bypassed. Colorado's SB26-051 passed the state senate on March 3 with similar requirements to California's law.

The pattern is clear. Governments are pushing age verification down the stack, from websites to app stores to operating systems. Apple is complying. Open-source projects are refusing. And the gap between the two is where the fight over who controls your device, and who gets to know how old you are, is playing out in real time.

Sources:

- 9to5Mac, "UK regulator Ofcom welcomes Apple age verification in iOS 26.4," Mar 25, 2026: "Apple yesterday released iOS 26.4...There was one change the company didn't mention, and that was the introduction of age verification for UK users." — <https://9to5mac.com/2026/03/25/uk-regulator-ofcom-welcomes-apple-age-verification-in-ios-26-4/>
- Engadget, "Apple introduces age verification for iCloud accounts in the UK," Mar 25, 2026: "The company says it will automatically switch on its Web Content Filter and Communication Safety features for everyone under 18 and for those who haven't verified their ages." — <https://www.engadget.com/big-tech/apple-introduces-age-verification-for-icloud-accounts-in-the-uk-115340237.html>
- Engadget, same: Ofcom: "Apple's decision that the UK will be one of the first countries in the world to receive new child safety protections on devices is a real win for children and families." — <https://www.engadget.com/big-tech/apple-introduces-age-verification-for-icloud-accounts-in-the-uk-115340237.html>
- TechRadar, "'If this starts affecting me using apps, I will switch to Android,'" Mar 27, 2026: Reddit user quoted: "Not everyone uses credit cards, and not everyone has a valid ID or license for the country they are living in. If this starts affecting me using apps, I will switch to Android." — <https://www.techradar.com/phones/ios/i-cant-believe-this-is-happening-iphone-users-are-threatening-to-defect-to-android-over-new-ios-age-verification-push-and-id-do-the-same>
- TechRadar, "Apple rolls out age verification in the UK with iOS 26.4," Mar 26, 2026: "While this specific step of age verification at a device level isn't required by UK law as it currently stands." — <https://www.techradar.com/phones/iphone/apple-rolls-out-age-verification-in-the-uk-with-ios-26-4-right-after-meta-and-google-get-fined-for-not-protecting-kids>
- to5Mac, "Apple age verification failing for some – here's what to try," Mar 26, 2026: "Some on the Apple Support forums and Reddit are reporting major problems with the process." — <https://9to5mac.com/2026/03/26/apple-age-verification-failing-for-some-heres-what-to-try/>

- Apple Developer, "Age requirements for apps distributed in Brazil, Australia, Singapore, Utah, and Louisiana," Feb 24, 2026: "Starting February 24, 2026, Apple will block users in Australia, Brazil, and Singapore from downloading apps rated 18+ unless they have been confirmed to be adults through reasonable methods." — <https://developer.apple.com/news/?id=f5zi08ey>
- TechCrunch, "Apple rolls out age-verification tools worldwide," Feb 24, 2026: "Apple will block the downloads of apps rated 18+ in Brazil, Australia, and Singapore, while also rolling out other features to comply with laws in Utah and Louisiana." — <https://techcrunch.com/2026/02/24/apple-rolls-out-age-verification-tools-worldwide-to-comply-with-growing-web-of-child-safety-laws/>
- California Legislature, AB-1043 bill text, Oct 13, 2025: "A person that violates this title shall be subject to...a civil penalty of not more than seven thousand five hundred dollars (\$7,500) per affected child for each intentional violation." — https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=202520260AB1043
- Tom's Hardware, "California introduces age verification law for all operating systems," Feb 2026: "California's Digital Age Assurance Act (AB 1043)...requires every operating system provider in California to collect age information from users at account setup and transmit that data to app developers via a real-time API." — <https://www.tomshardware.com/software/operating-systems/california-introduces-age-verification-law>
- Troutman Pepper, "Analyzing California's Digital Age Assurance Act," Oct 27, 2025: "The bill received unanimous, bipartisan support in both chambers. It passed the Assembly 76-0 in June and passed the Senate 38-0 in September." Governor Newsom "issued a signing statement urging the legislature to amend the law...to address concerns...including the 'complexities such as multi-user accounts shared by a family member.'" — <https://www.troutmanprivacy.com/2025/10/analyzing-californias-digital-age-assurance-act/>
- EFF, "A.B. 1043's Internet Age Gates Hurt Everyone," Mar 2026: "What California has done is essentially outsource censorship to developers, who are likely to lean into over-censorship." — <https://www.eff.org/deeplinks/2026/03/ab-1043s-internet-age-gates-hurt-everyone>
- Biometric Update, "Apple volleys age verification question back to sites and apps," Feb 28, 2025: Apple white paper: "Requiring age verification at the app marketplace level is not data minimization." — <https://www.biometricupdate.com/202502/apple-volleys-age-verification-question-back-to-sites-and-apps>
- Tom's Guide, "Major VPN provider sees a huge spike in sign ups," Jul 28, 2025: "Proton VPN sign ups in the UK have increased by 1,400% over the baseline. The spike started on July 25—the day the age verification law came into effect." — <https://www.tomsguide.com/computing/vpns/major-vpn-provider-sees-a-huge-spike-in-sign-ups-as-age-verification-law-comes-into-effect-in-the-uk>

- CybersecurityIntelligence, "VPN Demand Surges As British Online Safety Law Takes Effect," 2025: "NordVPN noting a 1,000% rise in UK subscriptions." — <https://www.cybersecurityintelligence.com/blog/vpn-demand-surges-as-british-online-safety-takes-effect-8580.html>
- Tom's Hardware, "GrapheneOS refuses to comply with new age verification laws," Mar 22, 2026: "GrapheneOS will remain usable by anyone around the world without requiring personal information, identification or an account...If GrapheneOS devices can't be sold in a region due to their regulations, so be it." — <https://www.tomshardware.com/software/operating-systems/grapheneos-refuses-to-comply-with-age-verification-laws>
- GrapheneOS (@GrapheneOS), X post, Mar 20, 2026: "We don't filter the internet for Iran or North Korea so why would we for Brazil or California?" — <https://x.com/GrapheneOS/status/2035037168414937462>
- Tom's Hardware, same: "The developers of open-source calculator firmware DB48X issued a legal notice recently, stating that their software 'does not, cannot and will not implement age verification.'" Also: "over 400 computer scientists who signed an open letter...argued that the laws create surveillance infrastructure without meaningfully protecting children." — <https://www.tomshardware.com/software/operating-systems/grapheneos-refuses-to-comply-with-age-verification-laws>
- The Register, "Age checks creep into Linux as systemd gets a DOB field," Mar 24, 2026: Garuda Linux maintainer: "Garuda Linux will not implement any age verification measures, since Garuda Linux's legal jurisdictions have no laws mandating age verification." — https://www.theregister.com/2026/03/24/foss_age_verification/
- It's FOSS, "Systemd's New Feature Brings Age Verification Option to Linux," Mar 2026: "The systemd project merged a pull request adding a new birthDate field to the JSON user records managed by userdb in response to the age verification laws of California, Colorado, and Brazil." — <https://itsfoss.com/news/systemd-age-verification/>
- Reason, "Computer scientists caution against internet age-verification mandates," Mar 4, 2026: "Over 400 computer scientists have signed an open letter warning that the rush to protect children from online dangers threatens to introduce new risks including censorship, centralized power, and loss of privacy." — <https://reason.com/2026/03/04/computer-scientists-caution-against-internet-age-verification-mandates/>

The Deep Dive (SKIPPED)

Skipped! Too much news to cover this week! :)

The Build Log

Chris N: Observability

So one of the things I've been deep in this past week or so is observability. Basically, building the ability to actually see what's happening with the product once it's out in the world. And it's been a ride.

Let me back up. For anyone who hasn't heard me talk about this before, Subtoken Auth is an authentication system I'm building for self-hosted apps. It generates permission-restricted access tokens for situations where a normal login flow doesn't work, like mobile apps that can't handle SSO redirects or two-factor prompts. The whole point is that you run it on your own server, on your own terms.

Which is great for privacy, great for control, but it creates a real blind spot for me as the developer. If something breaks, I don't get a phone call. There's no server I can check. The app is running on someone else's machine, behind their firewall. So the question becomes: how do I know if things are healthy out there? How do I know if people are even using it? And how do I know if something's going wrong before it turns into a support email?

The answer is community telemetry. A system where deployed instances can send back anonymized health signals. Think of it like a weather station network. Each station reports conditions locally, and when you aggregate them all, you get a forecast. No personal data, no tracking. Just stuff like "this instance started up successfully," "this instance validated 200 tokens today," "this instance is running version 1.2 on an x64 architecture with a community license." Useful stuff.

But building that pipeline was a project. The first version used a monitoring tool called SigNoz, and it worked, but it was heavy. Resource-hungry, complex to configure. So within about 48 hours I ripped that out and replaced it with a Grafana stack, which is the industry standard open-source monitoring toolkit for this kind of thing. LGTM stands for Loki, Grafana, Tempo, and Mimir, four tools that handle logs, dashboards, traces, and metrics respectively. Lighter, more flexible, and I actually understand what every piece does now. On top of that I had to build an authentication layer so only registered instances can send data. That went through a couple of iterations. Started with a single shared key, realized that doesn't scale, and ended up building a registration system where each deployed instance gets its own unique credential, digitally signed to prove it's legitimate, and issued automatically. No user interaction required. The app handles it silently on startup.

And then there's the license server, a separate service running on Cloudflare's global server network. When I started the week, it validated license keys and that was about it. By the end of the week, it validates keys, distributes plugins, registers telemetry credentials, monitors its own health, and deploys itself through an automated pipeline.

On the project management side, I recently moved from ClickUp to a self-hosted tool called Vikunja for task tracking. ClickUp worked fine, but I wasn't using most of the features and it felt like time to simplify. Now when I finish building a feature, the task gets marked done. When new work surfaces, it gets logged. Small thing, but it's smoothed out the workflow.

The big takeaway from all of this? When you're building a product that runs on other people's infrastructure, you have to build two products. You build the thing the customer uses, and then you build the thing that tells you the thing is working. And that second product is just as complex, just as important, and honestly, just as interesting to build.

Chris V: Same as before.

The Plug / Outro

Chris N's Plug

- A YouTube video from Joe Scott called "[Rocky Is Weirder Than You Think](#)"
- **Why?** If you've read Project Hail Mary by Andy Weir, or if you've been hearing about the movie that just came out, this is a deep dive into the alien character Rocky and the actual science behind how Andy Weir designed him. Weir apparently wrote 16 pages of notes on this creature's biology, everything from its circulatory system to its cellular makeup, all based on the real conditions of an actual exoplanet we've discovered. Joe Scott, an awesome creator himself, goes through those notes with Weir himself, and they get into details that aren't in the book or the movie. I haven't seen the film yet, but I'm a huge Andy Weir fan, The Martian was incredible, and this is exactly the kind of content I love: someone taking a fictional concept and pulling it apart to show you the real science underneath.

Chris V's Plug

- **None (officially)**

SquaredCast Links

- Patreon: <https://www.patreon.com/c/SquaredCast>
- Website: <https://squaredcast.com>

Final Notes

If you'd like to support what we're doing, check out our Patreon! You'll get bonus episodes, project builds, music from the archive, and a lot more, starting at just \$2/month.

We appreciate you being here. See you next week!