

Security Now! #1073 - 04-07-26

The FCC Bans New Consumer Routers

This week on Security Now!

- Apple's 26.4 age queries catches many by surprise.
- LinkedIn's 2.7 MB of privacy-invading javascript.
- Microsoft starts forcing Win11 24H2 to 25H2.
- Cisco loses source code to the Trivy supply-chain mess.
- Proton introduces privacy-first voice and video "Meet."
- GitHub to fix lagging security of its Actions feature.
- Cloudflare reaffirms the privacy of its 1.1.1.1 DNS.
- Cloudflare uses AI to re-code better secure Wordpress.
- The FCC drops a ban on all new consumer-grade routers.

In electronics, this symbol is a resistor which resists the flow of electrons. When used as it is here, it also resists the flow of people:



Our listener, Seth Smith, wrote that he was walking in his neighborhood, saw that nutty zig-zag concrete paved path and thought it would make a good picture of the week. I agree completely. For me, it's utterly inexplicable. It looks like the city's parks and sidewalks contractor may have mixed too much concrete and then needed to then use it all up. THIS ONE is a true mystery.

Security News

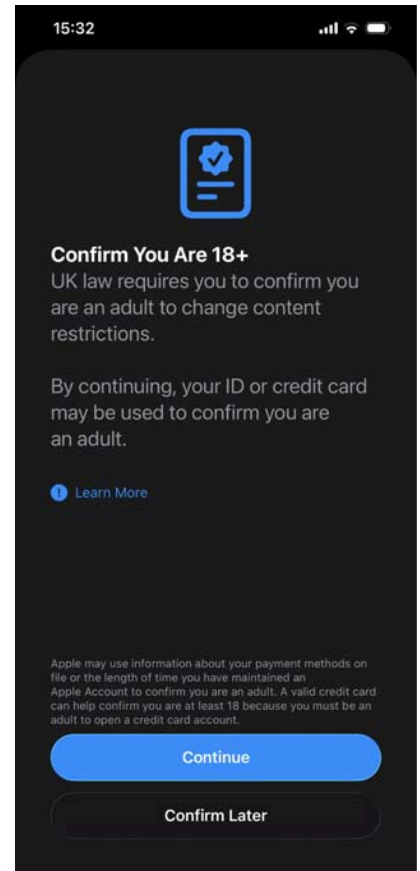
26.4

Last week's Apple upgrade to release 26.4 has sent age-confirming shockwaves through the UK. A listener, Dan Bright in Scotland sent an email: *"Hi Steve, Just FYI, although you likely know already: I'm in the UK and updated my iPhone to 26.4, to be immediately presented with an age verification process, which I'm instructed needs to be completed to enable age restricted content settings to be changed. Please find screenshot attached."* —>

Someone using the handle "Red" in GRC's SecurityNow! newsgroup posted: *"I am in the UK, have had an Apple Account for more than 18 years as I had an original iPod Touch (electronics similar to original iPhone) and after installing iPadOS 26.4 the system said "your Apple Account is older than 18 years, you are good." I wonder if that is a good method of doing age verification. Lots in the UK report problems, lots of people don't have credit cards (as you need to be 18 to have credit, so it's a common check) and Apple's system doesn't accept a UK Passport as proof of age."*

I poked around the Internet, reading feedback on The Guardian, 9to5Mac and elsewhere. Nothing stood out as worth sharing in greater detail. If I were to sum it all up with a generalization I'd call the nature of the reactions *"get off my lawn!"*. Normies, who don't listen to this podcast and who have not had any reason to track the rapidly changing landscape of online age verification, will be understandably surprised and annoyed by this apparently sudden need for their iDevices to need proof of their ages. For those who have been paying attention or listening to this podcast, this will be no surprise; one way or another, it will be coming sooner or later to every device we own.

As I noted last week, even reliably re-identifying an anonymous user remotely across a network has proven to be a challenge. Now we're needing to reliably and anonymously assert anyone's age. In my opinion, what Apple has done is exactly the right thing. Yes it's true that this will annoy some. 9to5Mac quoted a reader of theirs who commented on their coverage. He wrote:



This is quite a big failure by Apple. I use a debit card rather than a credit card. I've had one from the same bank for almost 40 years. I don't have a photo driving licence. My Apple account is about 14.5 years old. I can't verify my age despite being just over 60 years old. Even if they add a passport, which should have been usable from the start, I don't have one. As far as I understand, age verification is not required at device level, at least not yet, so Apple could either remove it, or make it opt-in. Whilst I can see how it's easier to have it on your device, so not having to verify age for all restricted websites and age related purchases, it needs to work for all, or not be forced on us. Besides, kids will find ways around it, and for now, from what I've seen, you can still get separate age verification for websites outside of Apple, unless they try to block people doing that.

Right... and, also, get off my lawn! This person wrote: *"it needs to work for all, or not be forced on us."* 100%! That would be great. But there's no magic solution. Partly, people are freaked out over **any** perceived loss of their (largely fictitious) anonymity online. Partly, people are upset over the imposition of any restrictions of any kind over what they can do online. They've never had any before, so why now all of a sudden? Are some freedoms being taken away from us and some restrictions imposed? Yes indeed. But everyone should blame their democratically elected politicians. The old adage of "don't shoot the messenger" applies here. The technologies are just doing the best they can to implement what the emerging regional laws require. As societies, we want to protect our children from all of the nastiness the world harbors. The anonymity that the Internet offers to criminals means that the Internet is likely to always contain more than its fair share of bad actors, just as it does today. That's unlikely to change.

So, might some of us, like this guy who grumbled to 9to5Mac, be inconvenienced by our collective desire to manage what kids can access online? Yes. That's going to happen. But that's the relatively small price we will need to pay. I don't see any way around it. And I love the idea that Apple is finally stepping up to this challenge. Having our platforms able to make these assertions for us – globally and anonymously – is the way to go.

All of this discussion of the age of our Apple accounts made me wonder whether there was any way for us to determine how long we've had our accounts. Since I have a credit card registered with Apple, they know I'm over 18. But since there are others who might be using debit cards and not have photo IDs I was curious. It appears that it's possible to bring up a web page at <https://privacy.apple.com>. You'll be asked to login with your Apple credentials when respond to a multi-factor prompt on one of your Apple devices. One you do that you'll be taken to a "Choose the data you want to download" page:





Choose the data you want to download

Select the data you'd like to download, and we'll prepare a copy for you. This process may take up to seven days. To ensure the security of your data, we use this time to verify that the request was made by you.

Your download will include:

- App usage and activity information as spreadsheets or files in .json, .csv, or .pdf format.
- Documents, photos and videos in their original format.
- Contacts, calendars, bookmarks and mail in .vcf, .ics, .html, and .eml format.

Your download will not include App, book, movie, TV show, or music purchases.

Back	Select all
 App install and push notification activity	<input type="checkbox"/>
 Apple Media Services Information Includes App Store, iTunes Store, Apple Books, Apple Music, Apple Podcasts, and Game Center activity	<input type="checkbox"/>
 Apple Account and device information	<input checked="" type="checkbox"/>
 Apple.com and Apple Store activity	Show more <input type="checkbox"/>

This is an amazingly comprehensive information request portal where you're able to download all sorts of information and data that Apple may have gathered and accumulated about you through your years of account ownership with them. The one item you need to select is "Apple Account

and device information.” I selected that one and pressed “Continue” at the bottom of the very long and comprehensive page of things I could request to receive from Apple. In fact, the list was so long and comprehensive that I was next asked how large a file I would be comfortable downloading. It defaulted to 1 gigabyte, but I chose the maximum offering of 25 gigs because the file is not going to be emailed. Once Apple has assembled the information I’ll receive another email, need to login again to prove my identity, then receive a link to download whatever Apple has to share. Since I initiated this last Saturday afternoon, and it’s expected to take as much as a week, I’m unsure when I’ll have any results to share. But in case anyone listening might also be curious to know how long they’re had their account, that appears to be the only way.

LinkedIn’s extensive use of visitor probing

Two weeks ago, Micah and I took a look at what we might term the “super-pixels” being used by Meta and TikTok to cause their own JavaScript code to be quietly run in the browsers of anyone visiting any website that hosted those “pixels” on behalf of Meta or TikTok. I noted at the time that the use of the term “pixel” was catching in my throat because what has evolved over time has rendered that term laughable.

So that we’re all starting off on the same page, the original idea was that a so-called “tracking pixel” could be hosted by a website a user was visiting. That “pixel” was simply a HTML URL for a single pixel-sized dot – a 1x1 jpeg, gif or png file. It might even be white or transparent, since it didn’t want to call any attention to itself.

Its entire purpose was to cause the user’s browser to fetch that tiny little innocuous 1x1 image dot from some other 3rd-party’s remote server. Just to be clear, it would be the visible website the user was visiting that would be delivering its pages to its visitors which contained the reference to that pixel. And since that pixel was referencing an image resource from another hosting domain, the user’s browser would quietly be making that request. Now we might wonder why the site being visited by its users might wish to add someone else’s invisible pixels to its own pages, and the somewhat distressing answer is that the site would be receiving payment by that 3rd-party site in return for the addition of those simple tiny pixels. So the obvious next question is why would some 3rd-party site be willing to pay 1st-party sites all across the Internet in return for hosting their little all-but-invisible pixels? And the answer to that is, of course, tracking. This was the emergence of the early Internet tracking economy.

When the user’s browser requested that tiny little invisible pixel from the remote 3rd-party server, its request contained a bunch of metadata information. The request’s “Referer” header would identify the entire URL of the page that site’s visitor was now viewing. And the request’s “Cookie” header would dutifully return the unique 3rd-party cookie that the 3rd-party site had previously given the user’s browser to hold. And, of course, the request would come from the user’s IP address. With these little tracking beacons scattered far and wide across the Internet, that 3rd-party site could just sit back and aggregate all the data that was available to it. The final bit of horror, which we covered here at the time, was that these tracking companies would create their own rewards and prize and sweepstakes websites, which they would advertise across the Internet in order to draw people in. When signing up for their chance to win nonexistent prizes, unwitting users would provide a ton of personal information, at least their names and email addresses and probably more. And since these bogus reward sites were being

hosted by the same companies who were littering the Internet with their tracking pixels, all of that anonymous tracking data aggregated over time – every website visited and the IP address used – would then be deanonymized when user provided their names and email addresses in return for nothing.

Sadly, those early days now look quaint in retrospect. As users became aware that the sites they were visiting were secretly betraying them behind their backs, compromising their privacy by embedding a pixel in return for payment, browser extensions such as our favorite uBlock Origin, but also Privacy Badger, Ghostery, AdGuard, Disconnect and NoScript were created to give users who cared some control over this egregious behavior.

The next thing to happen was the evolution of the embedded tracking object from a relatively benign – in retrospect – jpg, gif or png **image** pixel, into a reference to a remote host's HTML or JavaScript. Arranging to run a 3rd-party's remotely supplied JavaScript is the ultimate goal, and that can be done simply by directly referencing a 3rd-party JavaScript resource in the hosting page's HTML. Just like a site's own provided JavaScript, the 3rd-party JavaScript will be loaded into and run by every page the user displays. The problem we have now is that we've invited foreign code to run inside our web browser and the behavior of that code – the very code itself – is subject to unilateral change by that 3rd-party at any time. It is from such changes that the practice of web browser fingerprinting has evolved.

It should now be clear why the continued use of the term "pixel" for anything Meta or TikTok are doing today is laughable. A cute little "pixel" was 20 years ago. Today, what we have is "hostile, uncontrolled explicitly privacy compromising code execution by unseen 3rd parties." That's the threat environment that users and their browsers face today.

One of the points I wanted to make before we turn to last week's news about LinkedIn, is just how much of this behavior is completely hidden from anyone who is clicking links and wandering around the web. The expression "out of sight, out of mind" has never applied more than it does here. This unseen behavior has been a problem since the first use of a 3rd-party cookie for surreptitiously tracking user movement across the web. Through the intervening decades such behavior has exploded and the only thing that has any chance of reigning it in is government legislation that criminalizes its use when users have clearly stated their objections.

The examination of things that are going on behind people's back without their knowledge brings us to last week's LinkedIn revelation which has been dubbed "BrowserGate" by the apparently disgruntled developer who has a beef with LinkedIn's owner, Microsoft. The BrowserGate website is clearly passion-driven and its thesis raised some questions about its creator's motivations. But I'm getting ahead of the story. Let's first look at what the website at <https://browsergate.eu/> has to say.

Going to <https://browsergate.eu/> we're first confronted with the bold black headline "LinkedIn Is Illegally Searching Your Computer". The site then elaborates, writing:

Microsoft is running one of the largest corporate espionage operations in modern history. Every time any of LinkedIn's one billion users visits linkedin.com, hidden code searches their computer for installed software, collects the results, and transmits them to LinkedIn's servers

and to third-party companies including an American-Israeli cybersecurity firm. The user is never asked. Never told. LinkedIn's privacy policy does not mention it. Because LinkedIn knows each user's real name, employer, and job title, it is not searching anonymous visitors. It is searching identified people at identified companies. Millions of companies. Every day. All over the world. This is illegal and potentially a criminal offense in every jurisdiction we have examined.

I want to share what this author claims is the behavior of LinkedIn's downloaded code. And for the record, none of this behavior appears to be in dispute. It has been subsequently verified by independent researchers, including by BleepingComputer who has the advantage of objectivity. Under the heading of "What we found", this author writes:

Mass breach of personal data: LinkedIn's scan reveals the religious beliefs, political opinions, disabilities, and job search activity of identified individuals. LinkedIn scans for extensions that identify practicing Muslims, extensions that reveal political orientation, extensions built for neurodivergent users, and 509 job search tools that expose who is secretly looking for work on the very platform where their current employer can see their profile.

Under EU law, this category of data is not regulated. It is prohibited. LinkedIn has no consent, no disclosure, and no legal basis. Its privacy policy mentions none of this.

LinkedIn scans for over 200 products that directly compete with its own sales tools, including Apollo, Lusha, and ZoomInfo. Because LinkedIn knows each user's employer, it can map which companies use which competitor products. It is extracting the customer lists of thousands of software companies from their users' browsers without anyone's knowledge. Then it uses what it finds. LinkedIn has already sent enforcement threats to users of third-party tools, using data obtained through this covert scanning to identify its targets.

In 2023, the EU designated LinkedIn a regulated gatekeeper under the Digital Markets Act and ordered it to open its platform to third-party tools. LinkedIn's response: It published two restricted APIs and presented them to the European Commission as compliance. Together, these APIs handle approximately 0.07 calls per second. Meanwhile, LinkedIn already operates a [private] internal API called Voyager that powers every LinkedIn web and mobile product at 163,000 calls per second. In Microsoft's 249-page compliance report to the EU, the word "API" appears 533 times. "Voyager" appears zero times.

At the same time, LinkedIn expanded its surveillance of the exact tools the regulation was designed to protect. The scan list grew from roughly 461 products in 2024 to over 6,000 by February 2026. The EU told LinkedIn to let third-party tools in. LinkedIn built a surveillance system to find and punish every user of those tools.

LinkedIn ships your data to third parties. It loads an invisible tracking element from HUMAN Security (formerly PerimeterX), an American-Israeli cybersecurity firm, zero pixels wide, hidden off-screen, that sets cookies on your browser without your knowledge. A separate fingerprinting script runs from LinkedIn's own servers. A third script from Google executes silently on every page load. All of it encrypted. None of it disclosed.

Microsoft has 33,000 employees and a \$15 billion legal budget. We have the evidence. What we need is people and funding to hold them accountable.

Is this probably happening? As we'll see in a moment, apparently so. And thanks to the GDPR, much of what's being done behind the backs and without the explicit knowledge and permission of European Union citizens might well be illegal, as the creator of this website clearly believes. But knowing Microsoft, I would expect it to be covered by some vague consent to "business purposes" language which anyone can take to mean anything. The good news is, European regulators are generally unimpressed by such implied consent.

To help us examine this through a far less biased lens, two days ago, on Sunday, "The Next Web" did some great reporting. Even lacking the original author's bias, The Next Web's headline was "LinkedIn is secretly scanning your browser for 6,000 extensions, and you weren't told". Here's what they explained:

Every time you visit LinkedIn in a Chromium-based browser, a hidden JavaScript routine silently probes your browser for more than 6,000 installed extensions, collects 48 hardware and software characteristics about your device, encrypts the resulting fingerprint, and attaches it to every API request you make during your session. The practice, labelled "BrowserGate" by researchers, is not disclosed in LinkedIn's privacy policy. LinkedIn says it is a security measure; critics say it is covert surveillance of a billion users' browsing behaviour at industrial scale.

There is a routine that runs on your computer every time you open LinkedIn. You cannot see it, you were not told about it, and it is not described in the company's privacy policy. According to an investigation published in early April 2026 by Fairlinked e.V., a European association of commercial LinkedIn users, the platform injects a 2.7-megabyte JavaScript bundle into its website that silently scans visitors' browsers for the presence of more than 6,000 specific Chrome extensions, assembles a detailed fingerprint of their hardware, encrypts it, and transmits the result to LinkedIn's servers, where it is attached to every subsequent action taken during the session.

The investigation, independently confirmed by BleepingComputer, which verified the scanning behaviour through its own testing, has been dubbed "BrowserGate." LinkedIn disputes many of the report's characterisations. The technical facts are not in dispute.

LinkedIn calls its scanning system "Spectroscopy." When a user loads the LinkedIn website, the script fires off up to 6,222 simultaneous requests, each one probing for a specific browser extension by attempting to access files associated with that extension's ID. The presence or absence of a file in the response indicates whether the extension is installed. The entire operation runs silently in the background, without a visible prompt or notification of any kind.

Beyond extensions, the script collects 48 distinct characteristics of the user's device: CPU core count, available memory, screen resolution, timezone, language settings, battery status, audio hardware information, and storage capacity, among others. Individually, these attributes are unremarkable. Combined, they form a device fingerprint specific enough to identify a user even after cookies are cleared.

Once compiled, the data is serialised to JSON and encrypted using an RSA public key, LinkedIn's internal identifier for the key is "apfcDfPK", before being transmitted to telemetry endpoints including li/track and /platform-telemetry/li/apfcDf. The fingerprint is then permanently injected as an HTTP header into every API request made during the session, meaning LinkedIn receives it with every search, every profile view, every message sent.

Let me pause here for a moment. I haven't looked at the code, but what The Next Web described makes sense. Kinda. They wrote that the data was compiled, serialized to JSON and encrypted using an RSA public key. But my spidey sense tripped when I didn't see any mention of hashing and I did see that mention of reversible encryption thanks to the use of an RSA public key.

As we all know, the widely accepted way of "fingerprinting" a browser is to collect all of that random but very specific data then hash it down into an information-lossy irreversible hash. This creates a token that can be used to represent the user's browser as it moves about the web.

But my first question is why Microsoft would need to have that at all? This sort of fingerprinting is only used by 3rd parties who wish to track browsers as they move to other sites containing the same 3rd-party fingerprinting code. But Microsoft's LinkedIn users are already logged in with a 1st-party relationship with Microsoft. Why would Microsoft need to track them anywhere?

It seems to me that this is not a fingerprint at all in the traditional sense. I think that it must be a form of super-fingerprint. Microsoft is assembling those 48 data points into a JSON object which is then serialized. A random symmetric key will be derived and used to reversibly encrypt that serialized blob of data. That symmetric key will then be encrypted with the RSA public key contained within that massive 2.7 megabytes of JavaScript. That means that at any later date, Microsoft, or anyone else who might have the matching RSA private key, can decrypt the original symmetric key, then use that to decrypt and deserialize the JSON object to obtain the original 48 pieces of information.

Why would that be useful? The problem with using a hash to fingerprint is that thanks to the magic of cryptographic hashing, if even one single bit of the hash's input data are changed, on average half of the resulting hash's bits will be inverted. The point is that if just a single characteristic bit changes an entirely new and un-trackable hash results.

But Microsoft's super-fingerprint avoided the information-lossy hash. So they have presumably retained ALL of the information contained within those 48 pieces of information. That means that Microsoft's super-fingerprint can retain tracking – or more likely a tight association – even when some of the browser's captured data changes.

Since this is all sent back to the Microsoft LinkedIn mothership, what Microsoft probably does is fully decrypt all of that browser parameter data and keep it on file for every LinkedIn user. Over time, this would allow Microsoft to identify exactly how many and which web browsers each of their one billion LinkedIn users logs into. Perhaps that information might be useful for some security purpose.

The other thing that would also be interesting to check out would be what, exactly, those 48 pieces of information are. There might be a wolf hiding among the sheep. If the presumption was that everything was being hashed into a fingerprint that none of the specific information that Microsoft was collecting could be a big deal since it would be lost. But if we assume that Microsoft is collecting, reversibly encrypting and forwarding that to their mothership, it would be interesting so see exactly what they are collecting and retaining.

The Next Web has more to say, writing:

The question of which extensions LinkedIn is scanning for makes the surveillance more sensitive than simple fraud detection would require. According to the BrowserGate report, LinkedIn's list includes more than 200 products that compete directly with its own sales tools, among them Apollo, Lusha, and ZoomInfo. Because LinkedIn knows the employer of each registered user, systematically scanning for the presence of a competitor's tool gives the platform visibility into which companies are evaluating or deploying rival products.

The list also reportedly includes tools associated with neurodivergent conditions, religious practice, political interests, and job-hunting activity, categories that, in the European Union, qualify as sensitive personal data subject to heightened protection under the General Data Protection Regulation. Knowing that a user is running a job-search extension, for instance, is a meaningful inference about their employment intentions, drawn without consent.

The scale of the operation has grown substantially over time. LinkedIn began scanning for 38 specific extensions in 2017. By 2024, that number had grown to 461. By February 2026, the list had reached 6,167, a 1,252% increase in two years. BleepingComputer's testing confirmed the scanning was active as of early April 2026.

LinkedIn's response to BleepingComputer was pointed. A spokesperson said: "The claims made on the website linked here are plain wrong. The person behind them is subject to an account restriction for scraping and other violations of LinkedIn's Terms of Service. To protect the privacy of our members, their data, and to ensure site stability, we do look for extensions that scrape data without members' consent or otherwise violate LinkedIn's Terms of Service." The company added that it does not use the data to "infer sensitive information about members."

LinkedIn's characterisation of the source matters. Fairlinked e.V. is connected to Teamfluence Signal Systems OÜ, an Estonian company whose managing directors include Steven Morell and Jan Liebling. Teamfluence makes a Chrome extension, also called Teamfluence, that LinkedIn restricted for alleged terms of service violations. The company subsequently filed a preliminary injunction against LinkedIn Ireland Unlimited Company and LinkedIn Germany GmbH at the Regional Court of Munich, alleging violations of the Digital Markets Act, EU competition law, and German data protection rules. In January 2026, the Munich court denied the injunction, finding that LinkedIn's actions did not constitute unlawful obstruction or discrimination.

*The financial dispute between the parties does not change the technical findings, which were verified independently. It does mean the **framing** of those findings is contested, and readers should weigh both the substance of the claim and its provenance.*

This is not LinkedIn's first serious encounter with European data protection enforcement. In October 2024, the Irish Data Protection Commission, which regulates LinkedIn in the EU through its Irish subsidiary, fined the company €310 million, approximately \$334 million, for processing users' personal data for targeted advertising without a valid legal basis. The decision found that LinkedIn's consent mechanisms did not meet GDPR's requirement that consent be "freely given." LinkedIn was ordered to bring its data processing into compliance.

The BrowserGate investigation drops into that context. The legal question of whether scanning for 6,000 browser extensions constitutes processing of special-category personal data, and whether users' lack of awareness of the practice renders any implied consent invalid, is exactly the kind of question the Irish Data Protection Commission has already shown it is willing to adjoin in court. Europe's evolving digital regulation framework has been moving steadily toward requiring explicit disclosure of all significant data collection, and a scanning

operation of this scale, conducted without any mention in a privacy policy, appears difficult to square with that direction of travel.

LinkedIn is a Microsoft subsidiary, acquired in 2016 for \$26.2 billion. Microsoft has been aggressively expanding its AI capabilities in 2026, with LinkedIn's vast dataset of professional identity and employment history forming a significant part of the data infrastructure on which those capabilities rest. The relationship between LinkedIn's data collection practices and Microsoft's broader AI ambitions is not addressed in LinkedIn's privacy policy either.

LinkedIn has more than one billion registered users. The majority access the platform through Chromium-based browsers, meaning the Spectroscopy scan runs routinely on the devices of a significant fraction of the global professional workforce, collecting a fingerprint that is precise enough to persist across cookie resets and potentially across devices.

Short of using a non-Chromium browser such as Firefox, which would limit but not necessarily eliminate LinkedIn's fingerprinting capabilities, there is no user-facing setting that prevents the scanning. The platform does not offer an opt-out, because it does not disclose the practice in the first place. The 2026 push for governed and transparent AI and data practices is built on precisely the premise that invisible data collection of this kind should not be the default.

Whether regulators move quickly enough to change that default at LinkedIn's scale remains to be seen. Security firms increasingly built to detect exactly this kind of covert data harvesting are becoming a growth sector in their own right, a market indicator that the gap between what platforms collect and what users understand is still very wide. The year 2025 normalised AI-powered data collection at a pace that regulation has yet to match.

BrowserGate is a case study in what that lag looks like from the inside of a browser.

I thought that the summary statements in that article were spot on, and it is exactly what I was preparing to say if the article hadn't.

We began this topic by reminiscing over the quaint web browser pixel, which was literally a pixel image dot supplied by some other domain's web server. That has evolved – or perhaps devolved – into an astonishingly monstrous and invasive 2.7 megabyte, unsolicited blob of code that does actually, as observed, confirmed and reported by BleepingComputer scan the mass storage file system of its website visitors looking for the files belonging to 6,236 web browser extensions which are, arguably, absolutely none of its business.

As The Next Web stated, this may be illegal in the EU where, thankfully, privacy regulations are very strong and are only getting stronger. But whether or not this is legal, it seems pretty clear that things have gotten **way** out of hand, apparently due to a complete lack of adult supervision. Something really bizarre is going on at Microsoft's LinkedIn property. So regardless of the motivations of these begrudging developers, I'm very glad that the world has just received an absurdly clear example of the need to perhaps give our web browsers somewhat more control over, and more say, about what the JavaScript they host is allowed to do. If LinkedIn users were to receive a pop-up permission request saying: "The LinkedIn website you're visiting would like to rummage around inside your computer for a while, searching for the files belonging to 6,236 web browser extensions you may have installed. Do you consent to this?" I doubt much rummaging around would ensue.

Six thousand, two hundred and thirty-six individual web browser extensions individually checked for by searching the user's system's mass storage. Really Microsoft? What the hell is going on up there?

Win11 24H2 → 25H2

While I was over at BleepingComputer confirming that Lawrence Abrams had independently verified Microsoft's egregious JavaScript behavior I encountered the news that Microsoft has also now begun forcing upgrades of unmanaged Windows 11 PCs from 24H2 to 25H2. Last Friday, BleepingComputer reported:

Starting this week, Microsoft has begun force-upgrading unmanaged devices running Windows 11 24H2 Home and Pro editions to Windows 11 25H2. According to the company's Lifecycle Policy site, Windows 11 24H2 will reach end of support in roughly six months, on October 13, 2026. Also known as the Windows 11 2025 Update, Windows 11 25H2 began rolling out in September to eligible Windows 10 or 11 devices as a minor update installed through enablement packages less than 200 KB in size.

Microsoft said in a Monday update to the Windows release health dashboard: "The machine learning-based intelligent rollout has expanded to all devices running Home and Pro editions of Windows 11, version 24H2 that are not managed by IT departments. Devices running these editions will no longer receive fixes for known issues, time zone updates, technical support, or monthly security and preview updates containing protections from the latest security threats. These devices will automatically receive the update to Windows 11, version 25H2 when they're ready. No action is required, and you can choose when to restart your device or postpone the update."

Those who don't want to wait for the automatic upgrade can manually check whether the update is available in Settings > Windows Update and click the link to download and install Windows 11 25H2. If you're not ready to upgrade, you can also pause updates from Settings > Windows Update by selecting the amount of time you'd like to pause them. However, you must install the latest updates after the time limit has passed.

Microsoft also provides a support document and a step-by-step guide to help users resolve problems encountered during the Windows 11 25H2 upgrade process. Since the March 2026 Patch Tuesday updates were released, Microsoft has issued several emergency updates, including one that addresses a known issue breaking sign-ins with Microsoft accounts across multiple Microsoft apps, such as Teams and OneDrive. It also pushed out-of-band updates for hotpatch-enabled Windows 11 Enterprise devices that fixed a Bluetooth device visibility issue and security vulnerabilities in the Routing and Remote Access Service (RRAS) management tool.

I wanted to mention this because GRC's InControl freeware can also be used to give users control over this process. It configures Windows to appear as if it's under management, thus it is not "unmanged" and Microsoft will officially leave it alone. If you have used InControl to lock down your current Windows version and may wish to make the move to Win11 25H2, control can just as easily be released.

Cisco's source code escapes thanks to the Trivy vulnerability scanner

Last week's deep dive into the LiteLLM mess revealed that the proximate cause of LiteLLM's troubles was actually the use of a compromised free and open source vulnerability scanner called Trivy. It was widely expected that LiteLLM would not be alone in this and we have since learned that none other than Cisco Systems became another victim. BleepingComputer also reported on this, explaining:

Cisco has suffered a cyberattack after threat actors used stolen credentials from the recent Trivy supply chain attack to breach its internal development environment and steal source code belonging to the company and its customers. A source, who asked to remain anonymous, told BleepingComputer that Cisco's Unified Intelligence Center, CSIRT, and EOC teams contained the breach involving a malicious "GitHub Action plugin" from the recent Trivy compromise.

The attackers used the malicious GitHub Action to steal credentials and data from the company's build and development environment, impacting dozens of devices, including some developer and lab workstations. While the initial breach has been contained, BleepingComputer was told that the company expects continued fallout from the follow-on LiteLLM and Checkmarx supply chain attacks. As part of the breach, multiple AWS keys were reportedly stolen and later used to perform unauthorized activities across a small number of Cisco AWS accounts. Cisco has isolated affected systems, begun reimaging them, and is performing wide-scale credential rotation.

BleepingComputer has learned that more than 300 Cisco GitHub repositories were also cloned during the incident, including source code for its AI-powered products, such as AI Assistants, AI Defense, and unreleased products. A portion of the stolen repositories allegedly belongs to corporate customers, including banks, BPOs, and US government agencies. Multiple sources told BleepingComputer that more than one threat actor was involved in the Cisco CI/CD and AWS account breaches, with varying degrees of activity. BleepingComputer contacted Cisco with questions regarding the breach, but has not received a reply to our emails.

Cisco's breach was caused by this month's Trivy vulnerability scanner supply chain attack, in which threat actors compromised the project's GitHub pipeline to distribute credential-stealing malware through official releases and GitHub Actions. That attack enabled the theft of CI/CD credentials from organizations using the tool, giving attackers access to thousands of internal build environments.

Security researchers linked these supply chain attacks to the TeamPCP threat group based on the use of their self-titled "TeamPCP Cloud Stealer" infostealer. TeamPCP has been conducting a series of supply chain attacks targeting developer code platforms, such as GitHub, PyPi, NPM, and Docker. The group also compromised the LiteLLM PyPI package, which impacted tens of thousands of devices, and the Checkmarx KICS project to deploy the same information-stealing malware.

One snarky but understandable comment I saw from someone commenting upon the fact that some of Cisco's source code had escaped, wrote: "Maybe they can fix some bugs while they're in there." Yeah. We can only hope.

Introducing: Proton Meet

I know from seeing the domains of our Security Now! Listener email subscriptions that the Proton family of products are very popular among our listeners. So I wanted to note that last Tuesday Proton announced "Proton Meet", which Proton describes as a privacy-first end-to-end encrypted audio and video conferencing solution. They explained:

When meeting in person isn't an option, we turn to video calls for conversations too important for email or chat. Whether you're talking to a doctor, hosting an executive meeting, or checking in with your kids, you expect these interactions to be private and safe — but mainstream video conferencing services such as Zoom, Google, and Microsoft can eavesdrop on your conversations. Proton Meet gives you back your privacy and peace of mind by protecting your calls with end-to-end encryption, so nobody can listen in or use your conversations to sell ads, conduct surveillance, or train AI.

I was somewhat surprised by Proton's claim of eavesdropping and it appears that they're mostly referring to the leakage of event metadata. Also, their information may be a bit dated. Their claims included links for Zoom, Google and Microsoft. The Microsoft link talked about Outlook and the Zoom link was a posting written six years ago in 2020. This is not to suggest that I would not be far more inclined to trust Proton than Microsoft, Google or Zoom. I would without question. I have the link to last week's Proton's "Meet" announcement in the show notes for anyone who wants to follow-up: <https://proton.me/business/blog/introducing-proton-meet>

GitHub Actions to get much-needed additional security

The recent LiteLLM, Cisco and several other attacks have all been attributable to the Trivy malware scanner and the abuse of GitHub's Actions feature. In the wake of these various messes, GitHub has said it now plans to accelerate the development and rollout of some of the additional GitHub Actions security features it had originally planned to roll-out later this year. Since Actions are seeing serious abuse, there's no time like the present to improve their security!

Cloudflare says everyone's use of 1.1.1.1 remains super-private

Last week, on April 1st, Cloudflare posted:

Exactly 8 years ago today, we launched the 1.1.1.1 public DNS resolver, with the intention to build the world's fastest resolver — and the most private one. We knew that trust is everything for a service that handles the "phonebook of the Internet." That's why, at launch, we made a unique commitment to publicly confirm that we are doing what we said we would do with personal data. In 2020, we hired an independent firm to check our work, instead of just asking you to take our word for it. We shared our intention to update such examinations in the future. We also called on other providers to do the same, but, as far as we are aware, no other major public resolver has had their DNS privacy practices independently examined.

Their posting continues. But they were just audited by one of the top 4 accounting firms and again passed with flying colors. They do not want to know who uses their service nor what they look up. DNS querying source IPs are anonymized and deleted within 25 hours.

Cloudflare rewrites and replaces WordPress

Also on April 1st, Cloudflare announced their EmDash project. This caught my eye, and that of many of our listeners, since its goal is to replace Wordpress with a far more secure successor. The expression "far more secure" is not a high bar, since every Security Now! listener knows quite well what a complete security disaster Wordpress has become. It's not Wordpress' fault. The problem is its creaky old architecture which Cloudflare just replaced. The source of Wordpress' trouble has been that it promises to allow anyone to author and offer an insecure Wordpress plug-in, and many people have. Anyway, here's that we learn from Cloudflare:

The cost of building software has drastically decreased. We recently rebuilt [the most popular REACT framework] Next.js in one week using AI coding agents. But for the past two months our agents have been working on an even more ambitious project: rebuilding the WordPress open source project from the ground up.

WordPress powers over 40% of the Internet. It is a massive success that has enabled anyone to be a publisher, and created a global community of WordPress developers. But the WordPress open source project will be 24 years old this year. Hosting a website has changed dramatically during that time. When WordPress was born, AWS EC2 didn't exist. In the intervening years, that task has gone from renting virtual private servers, to uploading a JavaScript bundle to a globally distributed network at virtually no cost. It's time to upgrade the most popular CMS on the Internet to take advantage of this change.

Our name for this new CMS is EmDash. We think of it as the spiritual successor to WordPress. It's written entirely in TypeScript. It is serverless, but you can run it on your own hardware or any platform you choose. Plugins are securely sandboxed and can run in their own isolate, via Dynamic Workers, solving the fundamental security problem with the WordPress plugin architecture. And under the hood, EmDash is powered by Astro, the fastest web framework for content-driven websites.

EmDash is fully open source, MIT licensed, and available on GitHub. While EmDash aims to be compatible with WordPress functionality, no WordPress code was used to create EmDash. That allows us to license the open source project under the more permissive MIT license. We hope that allows more developers to adapt, extend, and participate in EmDash's development. You can deploy the EmDash v0.1.0 preview to your own Cloudflare account, or to any Node.js server today as part of our early developer beta.

Okay. That's all super interesting. But what about Wordpress' security? Before we get to that, Cloudflare felt the need to congratulate Wordpress and apologize for replacing it. So they wrote:

The story of WordPress is a triumph of open source that enabled publishing at a scale never before seen. Few projects have had the same recognisable impact on the generation raised on the Internet. The contributors to WordPress's core, and its many thousands of plugin and theme developers have built a platform that democratised publishing for millions; many lives and livelihoods being transformed by this ubiquitous software.

There will always be a place for WordPress, but there is also a lot more space for the world of content publishing to grow. A decade ago, people picking up a keyboard universally learned to publish their blogs with WordPress. Today it's just as likely that person picks up Astro, or another TypeScript framework to learn and build with. The ecosystem needs an option that

empowers a wide audience, in the same way it needed WordPress 23 years ago. EmDash is committed to building on what WordPress created: an open source publishing stack that anyone can install and use at little cost, while fixing the core problems that WordPress cannot solve.

And Here it come:

WordPress' plugin architecture is fundamentally insecure. 96% of security issues for WordPress sites originate in plugins. In 2025, more high severity vulnerabilities were found in the WordPress ecosystem than the previous two years combined.

Yeah In other words, things with Wordpress are getting worse, not better. They write:

Why, after over two decades, is WordPress plugin security so problematic?

A WordPress plugin is a PHP script that hooks directly into WordPress to add or modify functionality. There is no isolation: a WordPress plugin has direct access to the WordPress site's database and filesystem. When you install a WordPress plugin, you are trusting it with access to nearly everything, and trusting it to handle every malicious input or edge case perfectly.

EmDash solves this. In EmDash, each plugin runs in its own isolated sandbox: a Dynamic Worker. Rather than giving direct access to underlying data, EmDash provides the plugin with capabilities via bindings, based on what the plugin explicitly declares that it needs in its manifest. This security model has a strict guarantee: an EmDash plugin can only perform the actions explicitly declared in its manifest. You can know and trust upfront, before installing a plugin, exactly what you are granting it permission to do, similar to going through an OAuth flow and granting a 3rd party app a specific set of scoped permissions.

WordPress plugin security is such a real risk that WordPress.org manually reviews and approves each plugin in its marketplace. At the time of writing, that review queue is over 800 plugins long, and takes at least two weeks to traverse. The vulnerability surface area of WordPress plugins is so large that in practice, all parties rely on marketplace reputation, ratings and reviews. And because WordPress plugins run in the same execution context as WordPress itself and are so deeply intertwined with WordPress code, some argue they must carry forward WordPress' GPL license. These realities combine to create a chilling effect on developers building plugins, and on platforms hosting WordPress sites.

Plugin security is the root of this problem. Marketplace businesses provide trust when parties otherwise cannot easily trust each other. In the case of the WordPress marketplace, the plugin security risk is so large and probable that many of your customers can only reasonably trust your plugin via the marketplace. But in order to be part of the marketplace your code must be licensed in a way that forces you to give it away for free everywhere other than that marketplace. You are locked in.

Cloudflare's posting continues at much greater length, but everyone gets the idea. Cloudflare has leveraged AI agency to take the conceptual promise that Wordpress met and to dramatically overhaul its architecture for licensing freedom and security. **Nice going, Cloudflare!!**

The FCC Bans New Consumer Routers

Anyone encountering the news which landed two weeks ago on March 23rd would be correct in thinking that someone must have made a mistake somewhere. First of all, the reality of today's global electronics manufacturing sector is that U.S. domestically manufactured consumer-grade routers simply do not exist. All routers purchased by and available to US consumers are manufactured elsewhere – typically in China, Taiwan or Vietnam. So the FCC's surprise addition of every consumer router to the so-called "Covered List" means that the likes of Asus, Linksys, Netgear, Eero, TP-Link, D-Link and Nest have all suddenly joined the likes of previously banned non-consumer devices made by Huawei, ZTE, Hytera, and Hikvision.

The headline that appeared that appeared in the afternoon 15 days ago read: "*FCC Updates Covered List to Include Foreign-Made Consumer Routers*" The press release page explained that the full title was "*FCC Updates Covered List to Include Foreign-Made Consumer Routers, Prohibiting Approval of New Models.*" So, all of the existing, apparently attack-prone and buggy routers can still be sold. But anything that's new and hopefully improved is banned. Yay for the U.S.'s national security.

The official "Fact Sheet" that accompanied the press release included the helpful subhead: "*Update Follows Determination by Executive Branch Agencies that Consumer-Grade Routers Produced in Foreign Countries Threaten National Security*". I need to share some more of what the FCC wrote because it's not even internally consistent. The press release's Fact Sheet says:

WASHINGTON, March 23, 2026—Today, the Federal Communications Commission updated its Covered List to include all consumer-grade routers produced in foreign countries. Routers are the boxes in every home that connect computers, phones, and smart devices to the internet. This followed a determination by a White House-convened Executive Branch interagency body with appropriate national security expertise that such routers "pose unacceptable risks to the national security of the United States or the safety and security of United States persons."

The Executive Branch determination noted that foreign-produced routers (1) introduce "a supply chain vulnerability that could disrupt the U.S. economy, critical infrastructure, and national defense" and (2) pose "a severe cybersecurity risk that could be leveraged to immediately and severely disrupt U.S. critical infrastructure and directly harm U.S. persons."

President Trump's 2025 National Security Strategy stated: "the United States must never be dependent on any outside power for core components—from raw materials to parts to finished products—necessary to the nation's defense or economy. We must re-secure our own independent and reliable access to the goods we need to defend ourselves and preserve our way of life."

Malicious actors have exploited security gaps in foreign-made routers to attack American households, disrupt networks, enable espionage, and facilitate intellectual property theft. Foreign-made routers were also involved in the Volt, Flax, and Salt Typhoon cyberattacks targeting vital U.S. infrastructure.

As outlined below, today's action does not impact a consumer's continued use of routers they previously acquired. Nor does it prevent retailers from continuing to sell, import, or market router models approved previously through the FCC's equipment authorization process. By operation of the FCC's Covered List rules, the restrictions imposed today apply to new device models.

Okay. Wait. It just said “today’s action does not impact a consumer’s continued use of routers they previously acquired. Nor does it prevent retailers from continuing to sell, import, or market router models approved previously through the FCC’s equipment authorization process.”

So, in other words, every single one of the existing, apparently suddenly untrustworthy routers, that everyone in the world already has, are going to be left alone where they are. After all, what else can be done?; consumers already own those. This means that foreign manufacturers, which, again, is to say all router manufacturers, are prevented from introducing any new router models into the U.S. They’re free to keep making the existing routers. And they’re also presumably free to keep updating those routers’ firmware, which might be used to add new features or eliminate bugs. But that would mean that as WiFi technologies continue advancing and requiring support from new chipsets and radio hardware, newer routers cannot be obtained from traditional foreign suppliers.

That happened Monday afternoon 15 days ago. By the end of the week, the “Technology Policy Institute”, a Washington-based non-profit think tank, published an analysis of this action which I think is extremely useful and worth understanding because it compares what just happened to the previously enacted and outwardly similar ban on Huawei/ZTE equipment.

For the Technology Policy Institute, Scott Wallsten titled his piece: *“The FCC Got the Router Ban Wrong. It Knew Better.”* Here’s what he explained and reminds us:

On March 23, the FCC effectively banned all new foreign-made routers from the U.S. commercial market by adding them to its so-called “covered list.” The action followed a White House-convened interagency National Security Determination issued just three days earlier. The Commission took this action with no notice-and-comment proceeding, no published cost-benefit analysis, and without providing a broad transition process for the affected industry. The only path forward for manufacturers is to apply for “Conditional Approval” from the Department of Defense or the Department of Homeland Security.

I’ll note that the actual documentation about this requires this conditional approval to be obtained from the U.S. Department of **War** or the DHS. Scott appears to be choosing to use the Department's earlier name. So he continues:

The security concerns are real. Chinese state-sponsored hacking groups, including Volt Typhoon, Salt Typhoon, and Flax Typhoon, have exploited vulnerabilities in consumer routers to penetrate American networks, conduct surveillance, and build botnets for attacks on critical infrastructure.

I’m not taking issue with what Scott wrote here. But I do want to take the time to note that to the best of my knowledge, none of our **current** consumer-grade routers ship in an inherently vulnerable state. It’s true that in years past – meaning more than a decade ago – we were encountering instances where, for example, Intel’s “demonstration only” source code for their UPnP implementation was mistakenly dropped into routers. This resulted in UPnP being bound to consumer routers’ WAN-facing network interfaces. After delivering a podcast about that, the next week I announced that I had enhanced ShieldsUP! to explicitly check for public UPnP exposure. But that was all fixed back in 2013 and 2014 – 12 years ago.

And also back then, as in more than 10 years ago, we encountered instances where ISP-provided routers had open ISP administration ports. They were either using weak authentication credentials or also contained remotely exploitable weaknesses.

But for quite some time, now, it has only been when a router's user deliberately configures their router to allow external connections, and thus to implicitly solicit external attacks, that any of the various Chinese Typhoons – Volt, Salt or Flax – might have been able to get into user's networks through those routers. My point is, for quite some time now – like for the past 10 years – it's been users who have been unwittingly causing these external open-port exposure problems and none of that would be lessened by routers having domestic points of origin. Thus, nothing the FCC is attempting to do will fix anything that's broken. Scott continues:

Router security deserves serious attention. But in the past, the FCC addressed threats like these in a way that was more targeted, more precisely designed, and better built to survive a legal challenge. Comparing the FCC's handling of the Huawei and ZTE threat in 2019-2022 to the new router ban reveals what happens when an agency abandons the deliberative process that makes its expertise useful.

To respond to the national security risks posed by Huawei and ZTE, the FCC followed a deliberative process and produced a carefully constructed regulatory framework. Congress identified the specific companies as threats in Section 889 of the Fiscal Year 2019 National Defense Authorization Act. The FCC designated Huawei and ZTE as national security threats in June 2020, published its initial Covered List in March 2021, and adopted a Notice of Proposed Rulemaking and Notice of Inquiry on June 17, 2021, initiating two separate dockets and inviting public comment. The Commission then adopted a Report and Order in November 2022, with a unanimous 4-0 vote, and simultaneously issued a Further Notice of Proposed Rulemaking seeking additional comment on issues it hadn't yet resolved. That process took time. But it also produced outcomes that it could never have achieved in a weekend.

The comment process produced differentiated treatment based on actual risk. The FCC did not treat all five Chinese companies identically. It fully banned new Huawei and ZTE equipment, but took a more nuanced approach with Hikvision, Dahua, and Hytera. The FCC agreed with commenters who argued that these companies posed different levels and kinds of risk. The FCC required those three companies to document the safeguards they would put in place, and froze their applications pending that review. The router ban, by contrast, treats a Netgear router assembled in Vietnam identically to a TP-Link router designed in China.

The comment process identified a clear scope. The FCC had to define what counted as "covered" equipment. For example, it established that handset equipment designed for broadband operation with connection speeds of at least 200 kbps fell within the scope of "telecommunications equipment," while equipment below that threshold did not. That line was not in the original proposal. It emerged from the comment process, as affected companies argued that basic radio equipment should not be treated the same as broadband-capable devices. The FCC drew a principled boundary. The router ban draws no such lines. Its definition of "produced in a foreign country" encompasses "any major stage of the process through which the device is made, including manufacturing, assembly, design, and development," potentially sweeping in routers designed by American companies and assembled overseas.

The Huawei/ZTE response included transition assistance. The FCC's decision imposed real costs on carriers. Rural carriers told the FCC they couldn't afford to remove Huawei and ZTE equipment without financial help. Congress responded by creating the Secure and Trusted Communications Networks Reimbursement Program, initially funded at \$1.9 billion, which funded the removal and replacement of insecure equipment from carrier networks. The program has problems, such as a lack of evaluation and careful tracking of funds. But if the cost imposed on a company is due to a government mandate, the government should at least consider how to pay for it.

The comment process produced legal durability. During the rulemaking, commenters raised constitutional challenges, including arguments that the rules were an unconstitutional bill of attainder, violated the Equal Protection Clause, and amounted to an unconstitutional taking of property. The FCC addressed each of these arguments in its order, building a legal record. When Huawei challenged the related NDAA restrictions in court, a federal district court found the restrictions lawful because the government had demonstrated they reasonably furthered non-punitive national security goals. The router ban has no comparable record, and former FCC officials have already predicted it will face legal challenge.

Also, the process was iterative. The FCC recognized that its initial rules were a first step and continued refining them. A Second Report and Order clarified that covered equipment includes modular transmitters, proposed a definition of "critical infrastructure," and sought further comment on the scope of marketing prohibitions. The agency learned from industry input how supply chains actually work and adjusted its rules accordingly.

None of this happened with the router ban. The White House convened a panel. The panel issued a determination. Three days later the FCC implemented it.

Although the Secure Networks Act leaves the FCC little discretion over whether to add items to the Covered List once the White House makes a qualifying determination, the FCC still retains substantial leeway over how to implement the resulting equipment authorization restrictions, including its scope, transition periods, and what guidance it issues for affected parties. In the Huawei/ZTE proceeding, the Covered List addition itself was relatively quick, but the FCC spent more than a year designing the implementing rules through a public process. Nothing in the Secure Networks Act prevented the FCC from doing the same here. It chose not to.

The router ban bears all the hallmarks of a policy that never faced serious analytical scrutiny.

The stated justification is cybersecurity risk from foreign manufacturing. But the evidence the FCC itself cited undercuts the case for a country-of-manufacture approach. According to the Department of Justice, Volt Typhoon primarily targeted Cisco and Netgear routers, devices designed by American companies. The routers were vulnerable not because of where they were manufactured but because those companies had stopped providing security updates for discontinued models.

That's true. In the case of Netgear, Volt Typhoon leveraged routers whose firmware had never been updated, and was thus very old, and also exposed management interfaces with weak credentials. So, again, it's nothing about country of origin. Scott continues:

The FBI's own guidance urged router owners to replace end-of-life devices, and CISA's mitigation advice to manufacturers focused on secure design and automated updates, not supply chain origin. Salt Typhoon compromised major U.S. telecommunications carriers through network equipment made by Cisco, though Cisco's own security researchers reported that most intrusions it reviewed involved stolen credentials rather than software vulnerabilities. The national security determination includes supporting evidence from NIST, CISA, the FBI, and other agencies on router vulnerabilities generally. But none of it persuasively establishes that country of production, standing alone, is a useful proxy for cybersecurity risk.

An agency exercising careful judgment would have noticed this disconnect. If the problem is that manufacturers abandon security updates for older devices, the solution might be to mandate some kind of software maintenance or to require vulnerability disclosures, not a

blanket import ban organized around the country of manufacture. The FCC has an interdisciplinary expert staff who could have evaluated whether country of origin is actually a useful proxy for cybersecurity risk. Given the speedy timeline, it seems unlikely that they were consulted in a meaningful way.

In principle, country of manufacture could matter in hardware supply chains if a state actor could theoretically compromise hardware during production. This concern is real and deserves a serious policy response. But a blanket ban covering routers from every country on earth is not that response. A targeted action against manufacturers with documented ties to adversarial intelligence services, combined with supply chain integrity requirements for all manufacturers seeking FCC authorization, would address the hardware concern far more precisely. That is roughly what the FCC did with Huawei and ZTE. But the current ban treats a router from Finland the same as one from China.

Making the matter worse is that virtually no consumer-grade routers are manufactured in the United States. The only widely cited exception is some Starlink Wi-Fi routers that SpaceX says are made in Texas. Even major American brands including Netgear, Eero, and Google manufacture their products overseas.

The "conditional approval" process, which is the supposed escape valve, requires companies to disclose their management structure, detail their supply chain, and present a plan for onshoring manufacturing to the United States. That is not a security audit. It is industrial policy masquerading as a national security framework. No comment period helped shape it. And while there are extensive submission requirements, there appears to be no public review timeline or clear decision standard.

Meanwhile, the ban creates the very vulnerability it claims to address. Firmware and software updates for existing covered devices are permitted through at least March 2027, thanks to a blanket waiver from the FCC's Office of Engineering and Technology. But that waiver expires. A router that cannot receive security updates becomes exactly the kind of unpatched, vulnerable device that Volt Typhoon and Salt Typhoon exploited.

Some may argue that the post-Salt Typhoon threat environment necessitates faster action than the multi-year Huawei process allowed. But if that is true, it becomes hard to justify an action that does nothing about the millions of foreign-made routers already deployed in American homes and businesses, which are the actual devices that Volt Typhoon and Salt Typhoon exploited. If the threat were urgent enough to justify bypassing all deliberation, one would expect the FCC to be taking emergency action on the installed base. It is not. The ban addresses only future models, making this a forward-looking regulatory action for which a deliberative process was both feasible and appropriate.

A serious response would combine targeted restrictions on specific manufacturers with supply chain integrity and software maintenance requirements for all manufacturers seeking FCC authorization. The FCC has the expertise to design such a framework and it did exactly that with Huawei and ZTE.

In December testimony before the Senate Commerce Committee, Chairman Carr told lawmakers that the FCC "is not an independent agency, formally speaking." The router ban is a case study in what happens when that posture translates into skipping the processes that make regulation work. The comparison between the Huawei/ZTE process and the router ban is not just a story about two different policy decisions. It is a controlled experiment in what deliberative process is worth.

Same agency. Same statutory framework. Same category of threat. But the 2019-2022 process, in which the FCC used its full deliberative toolkit, produced targeted bans, differentiated treatment based on risk, precise scoping informed by industry expertise, billions in transition funding, and a legal record durable enough to survive court challenge. The 2026 process, in which the Commission used none of those tools, produced a blanket ban on an entire product category, no differentiation, no scoping analysis, no transition assistance, and a legal record so thin that former FCC officials are already predicting litigation.

The Secure Networks Act is the mechanism that enables this arrangement. Under the statute, the FCC says it cannot update the Covered List on its own but rather must implement determinations made by national security agencies. When those determinations were narrow and entity-specific, this was a manageable arrangement, and the FCC still exercised its own judgment in designing the implementing rules. Now that the determinations have expanded to cover entire product categories, and the FCC has chosen not to exercise its implementation authority, the agency is implementing sweeping trade and technology policy without the deliberation such decisions require.

*The same rapid-implementation pattern produced the December 2025 ban on foreign-made drones, which is already being challenged in court. In that case, Section 1709 of the Fiscal Year 2025 National Defense Authorization Act gave national security agencies one year to complete an evidence-based review of DJI drones, with an automatic Covered List addition as a fallback. Instead of a targeted review of DJI, the executive branch issued a broad national security determination covering **all** foreign-made drones, which the FCC implemented immediately. DJI has since sued to challenge the action.*

***"Process"** is not a bureaucratic waste of time. It is the mechanism through which an agency's expertise improves the quality of its decisions. The FCC demonstrated this in 2022 when it banned Huawei and ZTE equipment through a deliberative process that produced a more targeted, more durable, and more precisely designed result. Whatever the reasons the Commission did not follow the same approach here, the outcome speaks for itself.*

Congress should pay attention. The Secure Networks Act created a mechanism that, when combined with sweeping executive branch determinations and an FCC willing to implement them without deliberation, allows the President to ban entire categories of consumer technology without notice, without comment, without cost-benefit analysis, and without any of the procedural safeguards that normally govern consequential regulatory action. If Congress intended the Covered List to be used this way, it should say so. If it didn't, it should act before the next product category lands on the list.

So where does this leave us? This apparently arbitrary and short-sighted ban will do exactly nothing to improve the security of any existing routers whose current models may continue to be sold. Since new routers cannot now be sold, one effect may be to freeze current model numbers where they are. Unfortunately, major generational router improvements have multi-year design, development and manufacturing pipelines. This means that all of the router manufacturers will currently have their future models in the process of becoming ready for market. Except that now that market has just been killed for them.

That suggests that Scott is probably correct about future lawsuits. Under the terms of this ban, even domestic router manufacturers incorporated in the United States whose equipment is made offshore – which is to say all routers – will need to appeal to the “Conditional Approval” process which, as Scott noted *"requires companies to disclose their management structure, detail their*

supply chain, and present a plan for onshoring manufacturing to the United States.” What a mess. With any luck saner heads will prevail or competent management of the FCC will be installed. It doesn't appear that we currently have that.

I'll finish off today's look at actual consumer security with a little sanity-check reminder: Nearly everyone now has IoT network technology running inside the network security perimeter that's established and maintained by their NAT router. And many if not most of these phone home and maintain persistent connections to servers located outside the U.S. As I've been noting for many years now, these devices which we blithely invite into our homes to set up their own shops could do far more damage to consumer security than any routers designed within the last decade.

But don't tell anyone in Washington or they might become the FCC's next misguided target.

