



The FCC Bans New Consumer Routers

Description: 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.

High quality (64 kbps) mp3 audio file URL: <http://media.GRC.com/sn/SN-1073.mp3>

Quarter size (16 kbps) mp3 audio file URL: <http://media.GRC.com/sn/sn-1073-lq.mp3>

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. Amazed at the large privacy-invading JavaScript blob LinkedIn is forcing on people. We'll talk about that. And what does Steve think of the FCC router ban? Hmm. That's coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 1073, recorded Tuesday, April 7th, 2026: The FCC Bans New Consumer Routers.

It's time for Security Now!, the show where we cover the latest security, privacy, and other stuff with this guy right here, Mr. Steve Gibson.

Steve Gibson: It is true, my friend, it is true. It is once again, once again Tuesday. Yeah.

Leo: How did that happen?

Steve: Yeah. Well...

Leo: Well, Leo, if 168 hours go by...

Steve: And it hasn't always been. I like being here on Tuesday. It used to annoy me when we were on, I think we were on Monday for a while, and three-day weekends would kill, cancel a podcast.

Leo: No, this is a good thing. Hey, you know what? If Wednesday would work better, you just let me know.

Steve: No.

Leo: I'm here for you, Steve.

Steve: We're good. I know, that would upset everybody else's, like they...

Leo: No, it upset Paul when we moved from Tuesday to Wednesday. He said, "What are you talking about?"

Steve: Okay. So the big topic actually from a week before from our listeners, I just didn't have a chance to dig into it and get to it and look at it and talk about it, was this bizarre sudden surprising FCC ban on new consumer routers. It's important that it's not existing consumer routers. No. It's like, anyway, we're going to talk about this. And by the time we're done with today's podcast, everybody listening will understand exactly what happened, why it doesn't really make any sense, lots of little "what I think about its" laced in there. Several of our listeners did note that I did a poor job of wording the summary because in the little summary bullet points in the email that went out Sunday and is here at the top of the show notes, I said "The FCC drops a ban on all new consumer-grade routers."

Leo: Oh, they adopted it rather than dropping it. Yes. I understand. Your thinking is...

Steve: Well, they dropped a ban on them. But of course that's got a double meaning because there could have been a ban which they then dropped, as in stopping the ban.

Leo: By the way, I remember what the "fortunately" that I left hanging before the show began has to do with this. Fortunately, anybody who watches this show knows how to make their own router, as you do.

Steve: Yes. Yes. But we're not...

Leo: But we'll I'm sure talk about that, too.

Steve: Not concerned about that. Anyway, but we have a bunch of stuff to talk about, as always. We got the fact that Apple's sort of notorious now 26.4, you know, if you just say 26.4 to an Apple person, they go, oh, yeah. It's caught many people by surprise. We've got LinkedIn's 20, okay, 2.7MB privacy-invading JavaScript blob which occurred...

Leo: What?

Steve: Yes, Microsoft has just gone off the rails with this. We also have a quick note about them forcing Windows 11 24H2 to 25H2, and the consequences of that. Cisco losing their source code to - they were another casualty of this Trivy supply chain mess that caught LiteLLM that we talked about last week. Proton, a big favorite of our listeners, has introduced what they're calling a privacy-first voice and video service known as Meet. GitHub is going to respond to the other, this whole Trivy/Cisco/LiteLLM mess by taking a closer look at the security of its actions feature, which is what was abused to make this happen, and change their rollout schedule. Cloudflare reaffirming the privacy of its DNS service. And, oh, Leo, they've recoded - Cloudflare, that is - recoded WordPress to fix its security.

Leo: Oh, yes, EmDash, yeah, yeah.

Steve: Yes, EmDash, which is a cute name for that, I think.

Leo: It is. It's a little confusing. But now, once I understood...

Steve: If you weren't into, you know, like I guess manually...

Leo: Typography, yeah.

Steve: Yeah, exactly, typography, EnDash and EmDash and so forth. So lots of stuff to talk about. We've got a great Picture of the Week which has been explained by a couple of our listeners. This was actually submitted by a listener who was walking by, saw this, and thought, okay, I've got to take a picture of this for Steve because, you know, this is wacky.

Leo: I can't wait to hear it.

Steve: So, and I love my caption, if I do say so. So anyway, we will...

Leo: We'll talk about it.

Steve: We will get there.

Leo: It is coming up next as we continue with Security Now! for this Wednesday, April 7th. Now it's time for the Picture of the Week.

Steve: Okay. So the caption, you have to have the caption first because I love this caption. I wrote: "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."

Leo: So I'm thinking it's a line with like a spiral squiggle; right? Is that a resistor? I think that's what it is.

Steve: That's a coil, actually. That's a coil.

Leo: That's a coil. What's the resistor?

Steve: It's the zigzag.

Leo: The zigzag, right, right, right.

Steve: Yup, yup.

Leo: Well, "Why?" is I guess the question.

Steve: And so this is not - I know. Clearly it's a resistor; right? I mean, that's what it is.

Leo: Yeah. Yeah.

Steve: It is a walking path resistor. Several people wrote saying, I think this is AI-generated. This is nonsense. Well, of course, in this world, unfortunately, it could well be AI-generated. I'm afraid that's just going to be the standard response to anything that looks bizarre from now on. But this was actually - this is a photo taken by a listener who, as I said, was walking by, saw this, and thought, oh, Steve's going to have fun. So, and this was Seth Smith, our listener, sent it to me. Anyway, the best explanation I've had is that a straight line may have been too steep for getting by, like a wheelchair or some...

Leo: Ah. Because it is a little bit of a grade, yeah.

Steve: It looks like there is a grade.

Leo: So this is a switchback.

Steve: And so code might require, yes, that they switch back and forth in order not to encounter a grade which is too steep in order to roll themselves if they were in some fashion handicapped up to that little table. So that, you know, sounds right.

Leo: Maybe.

Steve: But it's not very well done.

Leo: It's ugly. It feels like there's something missing, like there might have been a reason for this at one time that no longer...

Steve: And it's not rounded. It'd be nice if the corners, the points were rounded. It looks like it's kind of pinchy up there at the very far one on the far right. It's like, just doesn't look great. But anyway, it is indeed a path resistor. So...

Leo: Weird.

Steve: Yeah.

Leo: Very weird.

Steve: Okay. So last week's Apple upgrade to release 26.4 has sent age-confirming shockwaves through the UK. And in fact I even got one. I'm trying to think what it was I was doing. I was doing something, I was logging into some app. Oh, I think I might have been installing Claude on an iPhone, and I got a short little pop-up that just notified me that the app had been informed that I was over 18. And I thought, oh.

Leo: Oh, interesting. Oh.

Steve: Yeah. And, you know, I had occasion, as a consequence of all this, to go into my settings. And I, you know, I am there. It shows my name under my Apple account for Steve Gibson. Shows my name, and my birth date is, you know, it's in the phone.

Leo: Oh, interesting.

Steve: So the phone knows how old I am. So anyway, so a listener of ours, Dan Bright in Scotland, he sent an email, said: "Hi, Steve. Just FYI, although you likely already know. 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." And he sent me a screenshot. And this is, you know, clearly an iPhone. And good for you, Dan. Your phone is being kept charged. We know that our lithium-ion batteries appreciate you keeping the phone charged.

So it says: "Confirm You Are 18+. UK law requires" - and I'll just note that it actually doesn't, but okay. "UK law requires you to confirm you are an adult to change content restrictions. By continuing, your ID or credit card may be used to confirm you are an adult." And there's a big blue Click to Continue, or you can defer that and confirm later, or learn more about that.

So also, somebody using the handle "Red" over in GRC's Security Now! newsgroup posted: "I am in the UK, have had an Apple Account for more than 18 years as I had an original iPod Touch." And he said: "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.'"

Leo: That's great.

Steve: And he said: "I wonder if that's 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 that's a common check." He said: "And Apple's system doesn't accept a UK passport as proof of age." And I think that might have just changed. Apple's been iterating on this because of the problems and the feedback that they've been receiving once this went out to a much wider audience.

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 sort of call the nature of the reactions "Get off my lawn." You know, 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, certainly as we have been, they'll be understandably surprised and annoyed by this apparently sudden need for their iDevices to need proof of their ages. It's like, what? Why? What? Huh? So for those who've been paying attention, of course, or listening to this podcast, this won't be any surprise at all. One way or another, it will be coming to every device we own.

As I noted last week, even reliably reidentifying an anonymous user remotely across a network, known as authenticating, has proven to be a challenge. Now we're needing to reliably and anonymously assert anyone's age. Which is, you know, another, like, whole level. In my opinion, as I've said, I believe what Apple has done is exactly the right thing. You know, it's true this will annoy some people. 9to5Mac quoted a reader of theirs who commented on their coverage. This guy 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 license. My Apple account is about 14.5 years old," meaning not 18 like Red's was, who I shared before.

He said: "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 of those, either. 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. So as I said, get off my lawn. He said: "It needs to work for all, or not be forced on us." Right. 100%. That would be great. But there's no magic solution; right? Partly, people are freaked out over any perceived loss of their (largely fictitious) anonymity online. Partly, people are upset over the imposition of any restriction of any kind over what they can do online. They've never had any before, so why now all of a sudden? You know? Are some freedoms being taken away from us and some restrictions imposed? Yeah, they are.

But everyone should blame their democratically elected politicians. The old adage of not shooting 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 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? Uh-huh. Yes. That's going to happen. But that's the relatively small price we 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, like this 9to5Mac guy, and may not have photo IDs or not have one that Apple understands, because there were some reports of that in the UK also, it appears that it's possible to bring up a web page at privacy.apple.com. You'll be asked to login with your Apple credentials, then respond to a multi-factor prompt on one of your Apple devices. Once you've done that, you'll be taken to a "Choose the data you wish to download" page.

I have a picture of that at the bottom of page 3 in the show notes. And this is - I've never seen this page before. It's 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 the years of your account ownership with them. The one item you reportedly, because I haven't been able - I've started the process. It hasn't finished because it takes - it can take up to a week. The one item you need to select is Apple Account and Device Information. But, boy, there's a lot more, if you want more. So I selected that one and pressed Continue at the bottom of that very long and comprehensive page of things that 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 1GB, in which case it sends you however many 1GB files you need. But I chose the maximum offering of 25GB because the file isn't emailed. Once Apple has assembled the information, I then receive another email. I need to log in again to reprove my identity. Then I receive a link to download whatever Apple has to share with me. Since I initiated this just last Saturday afternoon, so three days ago, and it's expected to take as much as a week, I'm unsure, you know, when I'll have results to share. Probably by this time next week, and I'll just quickly let people know what I got.

But anyway, I just wanted to share all that in case anyone listening might also be curious to know how long they've had their account. I wasn't quick to jump on an iPhone. I don't think I ever had an early iPod. You know, I was in love with my Blackberry. I was one of those "pry this from my cold dead, you know, grip."

Leo: Do you still have some in your freezer?

Steve: No, I don't. And I actually did get - there is a physical keyboard for the iPhone.

Leo: Did you order that?

Steve: Yeah, I tried it. It's crap. It's no good.

Leo: No.

Steve: Besides, it makes your phone about - I can't show it on the screen.

Leo: Yeah, it's dorky, yeah.

Steve: But it's like it's really, like, weird. It's like, you know, a foot long and sticks out of your pocket. So, no, it's not going to go.

Leo: I cannot use the iPhone keyboard [crosstalk] terrible.

Steve: I hate it. It is the biggest tradeoff that, I mean, I get it that that's what they want to do. And remember, of course, it was meant to be a consumption device, theoretically. But no, I - now, I do have a Bluetooth keyboard that is a full-size keyboard. It's a cute little thing from Logitech, that guy.

Leo: Oh, that's nice. It's pretty.

Steve: Yeah, it is. And it allows you go associate itself with up to three different devices, and you choose which one you want. So it can appear as one of three different keyboards. And so if I know I'm going to be typing something at length. Normally, though, I'll just do - I was like, I'm so annoyed also, Leo, with this schism between iPhone and Windows. Like Apple just refuses to accept the fact that Windows owns the desktop, and they're only willing to connect to their Mac in a seamless fashion. So I'll, like, write something long. Then I'll email it to myself, get the email on my phone, select it all, drop it into Message and send it. It's like, god, really? This is what I'm being made - it's one of the reasons I'm so annoyed with Apple. But anyway...

Leo: Yeah. Apparently, you can also look at your purchases on the Apple ID and see when the oldest purchase you made...

Steve: Yes. Apparently everything, I mean, this page...

Leo: They're all there.

Steve: ...is so comprehensive. So I just thought it was cool. I didn't know, I'd never gone to privacy.apple.com and done that. But if someone wants to know how long they've had their account, I'm - I don't remember when we began the podcast. That would be an interesting - oh, I know that we began the podcast in 2005. But I don't remember whether I had an iPhone at that time.

Leo: Ah.

Steve: When I went on that...

Leo: Well, they didn't come out until 2006, so I know you didn't have an iPhone in 2005.

Steve: Oh. Oh, okay. So I was still Blackberry, happily carrying my Blackberry around.

Leo: The most you could have would be 20 years' worth of iPhone. Because actually next year's the 20th anniversary.

Steve: Oh, okay. And I probably waited a couple years because, again, I just liked, you know, my little Blackberry was...

Leo: Well, you didn't know that you had to prove you were 18 until now.

Steve: That's true. I want to - this next piece, LinkedIn and what Microsoft has done, is long. Let's take a second break.

Leo: Sure.

Steve: Then I'm going to plow into some research that a disgruntled add-on developer posted. But despite his disgruntlement, he's not wrong. And Leo, what Microsoft is doing is, you know, because they're the owner of LinkedIn, it's like, what?

Okay. So two weeks ago, while you were at RSA, Leo, Mikah and I took a look at what we might term the "super pixels" being used by Meta and TikTok now which caused their own JavaScript code to be quietly run in the browsers of anyone visiting any website that hosted those "pixels" - and I'm putting "pixels" in air quotes here because, well, I'm going to explain. I noted at the time, two weeks ago, that the use of the term "pixel" was almost catching in my throat because what has evolved over time has rendered that term laughable.

So just so that we're all starting off on the same page here, the original idea was that a so-called "tracking pixel" could be hosted by a website that a user was visiting. That "pixel" was actually an HTML URL for a single true pixel-sized dot - a 1x1 JPEG, or GIF, or PNG file. It might even be white or transparent, since it didn't want to call any attention to itself. It just wanted to be on the page.

Its entire purpose was to cause the user's browser to fetch that tiny little innocuous 1x1 image dot from some other third-party's remote server. And so 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 off-site third-party pixel. And since that pixel was referencing an image resource from another hosting domain, the user's browser would quietly be making that request to retrieve that pixel on behalf of its user.

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 third-party site in return for the addition of those simple tiny pixels. So the obvious next question is, why would some third-party site be willing to pay first-party sites whom people visit all across the Internet in return for hosting their little all-but-invisible pixels? And the answer to that, of course, is 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 third-party server, its request contained a bunch of metadata information. The request's "Referer" header would identify the entire URL of the page the site's visitor was viewing. And the request's "Cookie" header would dutifully return the unique third-party cookie that the third-party site may have previously given the user's browser to hold, assuming they encountered one of these little pixels from that same third-party site anywhere in

the past. And of course the request would come from the user's IP address. So lots of information available to some random unaffiliated, you know, not obviously affiliated third-party site. And all these little tracking beacons scattered far and wide across the Internet would be gathered by this third-party site, which could just sit back and aggregate all that 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 prizes and, like, kind of sketchy sweepstakes websites, where they would advertise across the Internet in order to draw people in. When signing up for their chance to win nonexistent or maybe prizes, unwitting users would provide a ton of personal information to that site, at least their names and email addresses and probably more. Sometimes their phone numbers. Because, hey, there's a chance to win.

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 that had been aggregated over time - every website visited, the IP address that it had been visited from, the user's IP address - would then be all deanonymized when the user provided their name and email addresses in return for essentially nothing.

Unfortunately, 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 one, 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 now, in retrospect, just a little JPEG, GIF, or PNG image pixel, into a reference to a remote host's HTML or JavaScript. Arranging to run a third-party's remotely supplied JavaScript now is the ultimate goal, and that can be done simply by directly referencing a third-party JavaScript resource in the hosting page's HTML. You know, just whatever it is .js, for JavaScript, just like a site's own provided JavaScript, the third-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 third-party at any time. And it is from such changes that the practice of web browser fingerprinting has evolved. Right?

It should now be clear why the continued use of the term "pixel," you know, for anything Meta or TikTok are doing is laughable. It's not a pixel any longer, although it still goes by that name. What we have now is essentially "hostile, uncontrolled explicitly privacy-compromising code execution by unseen third parties." That's the threat environment that users and their browsers face today. It's completely changed over the course of the last two decades, the duration of this podcast, we've seen all of this.

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. You know, 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 third-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 reining it in on a wholesale level, not just like we who care running add-ons like uBlock Origin, but like actually affecting everyone, is government legislation, which will eventually wrestle this stuff to the ground, criminalizing it, and just making it impossible to continue. And there we probably have the EU to thank because they tend to be pioneering this.

Leo: They use pixel, now pixel is used kind of generically in the ad industry as anything that's tracking. We're often asked to put pixels in our podcasts, which obviously you can't do. You can't pixels in an audio file. But they...

Steve: There were some - there have been some audio beacons, though; haven't there?

Leo: Well, and we do, actually. In the feed there are redirects, and that's as close as you can get; right.

Steve: Right, it bounces you through something, and then eventually - well, and in fact Podtrac, you were using Podtrac in order to count downloads.

Leo: That's right. That's right. That's exactly how we were counting downloads. Exactly. And we still do something similar. We don't do it with Podtrac, but yeah. That's exactly it. So they call those "pixels," too. It always puzzled me when advertisers say can you put a pixel in the podcast? And it's like, I don't think so.

Steve: Yeah, because they're not, you know...

Leo: Not a pixel. But they just need tracking of some kind.

Steve: And it's interesting, too, because it does demonstrate how much they've just come to take it for granted.

Leo: Yeah. Oh, yeah. Absolutely.

Steve: We want analytics. We want enough data, as much - yes. Well, talk about knowing everything, Leo. Wait till you hear this.

Leo: Okay.

Steve: Okay. So 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." Okay, now, this is 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 in my mind about its creator's motivations. Okay. But I'm getting ahead of the story. Let's first look at that website. It is at browsergate.eu, B-R-O-W-S-E-R-G-A-T-E dot eu.

So going there, we're first confronted with the bold black headline "LinkedIn Is Illegally Searching Your Computer." Okay. The site then elaborates, writing: "Microsoft is running one of the largest corporate espionage operations in modern history." And now that may seem like it's a little over the top, but I don't think anyone's going to think that once we're through looking at this closely. He wrote: "Every time any of LinkedIn's one billion users visits LinkedIn.com, hidden code" - okay, it's not really hidden, I mean, all code is

hidden; right? We don't look at the code. No one does. But, you know, so yes, it's de facto hidden.

"Hidden code searches their computer for installed software" - that's true - "collects the results" - that's true - "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 doesn't mention it. Because LinkedIn knows each user's real name, employer, and job title, it is not searching anonymous visitors. It is searching the computers' identified people at identified companies. Millions of companies. Every day. All over the world." He writes: "This is illegal and potentially a criminal offense in every jurisdiction we have examined."

Okay. So 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. All of the evidence that they have collected is available for download and analysis. And it has been subsequently verified by independent researchers, including by BleepingComputer, who has the advantage of objectivity and knowing their way around.

So 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's 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 its 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." Meaning he's saying that they're not acknowledging the use of this internal, this other internal API.

"At the same time," he writes, "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 of 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 offscreen, 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 download." Well, many people have that. But, he says: "All of it encrypted. None of it disclosed." And he finishes: "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."

Okay. So 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, you know, vague consent to "business purposes" language which anyone can take to mean anything, as we've seen. The good news is, European regulators are genuinely and generally unimpressed by such implied consent.

Okay. So to help us through a far less biased lens than this guy has, two days ago, on Sunday, The Next Web did some great reporting on this. 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." And just to give everyone a - so that we understand what we're talking about, it is actually reading, searching for files on a user's hard drive. It is looking through your file system when you go to a LinkedIn page, which is what BleepingComputer confirmed and showed, you know, happening in their report of this.

Leo: Is it part of its fingerprinting, do you think? Or do they want that information?

Steve: There's actually different fingerprinting than this. They apparently actually want to know what browser, Chromium browser extensions you have installed. So here's what - and The Next Web makes this clear. They said: "Every time you visit a Chromium-based browser" - and actually Firefox users apparently are subjected to far less of this because it is very browser-specific.

Leo: Ah.

Steve: Yeah, which is nice. "Every time you visit LinkedIn in a Chromium-based browser, a hidden" - again, they used the word "hidden," but okay - "JavaScript routine silently probes your browser for more than 6,000 installed extensions, collects 48 hardware and software characteristics about your device" - that's the fingerprinting part - "encrypts the resulting fingerprint, and attaches it to every API request you make during your session. The practice, labeled 'BrowserGate' by researchers, is not disclosed in LinkedIn's privacy policy," says The Next Web. You know, who doesn't have a cross to bear here, or axe to grind. LinkedIn says it's a security measure. Critics say it is covert surveillance of a billion users' browsing behavior at industrial scale.

"There's 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" - get this, Leo - "injects a 2.7MB JavaScript bundle." It's like, wow.

Leo: Geez.

Steve: 2.7MB of JavaScript "into its website that silently scans visitors' browsers for the presence" - and actually it's the visitors' PCs - "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's attached to every subsequent action taken during the session.

"The investigation," The Next Web writes, "independently confirmed by BleepingComputer, which verified the scanning behavior through its own testing, has been dubbed 'BrowserGate.' LinkedIn disputes many of the report's characterizations. The technical facts, however, 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" - on the user's file system - "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." Again, 6,222 extensions, like, that it's checking for. Why? What business is it of LinkedIn what extensions to that degree, more than 6,000 of them, that a user has installed?

"Beyond extensions, the script collects 48 distinct characteristics of the user's device: CPU core count, available memory, screen resolution, time zone, language settings, battery status, audio hardware information, and storage capacity, among others." Now, those are traditional, those are like standard fingerprinting; right? They said: "Individually, these attributes are unremarkable. Combined, they form a device fingerprint" - true - "specific enough to identify a user even after cookies are cleared." Okay. We've all seen that before.

However, they said: "Once compiled, the data is serialized 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."

Okay. Now, I'm going to pause here a minute. I haven't looked at the code, but what The Next Web described makes sense in an interesting way. 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, yet very specific data, then hash it down into an information-lossy, thus 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 third parties who wish to track browsers as they move to other sites containing the same third-party fingerprinting code. But Microsoft's LinkedIn users are already logged in with a first-party relationship to Microsoft. So why would Microsoft need to track them anywhere? Doesn't make any sense there.

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 what I will call a "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 JSON blob. That symmetric key will then be encrypted with the RSA public key contained within that massive 2.7MB of JavaScript. That means that at any later date, Microsoft, or anyone else who might have the matching RSA private key, and they alone, can decrypt the original symmetric key, then use that to decrypt and deserialize the JSON object to obtain the original 48 individual pieces of information.

Why would that be useful? Well, the problem with using a hash to fingerprint is that, thanks to the magic of cryptographic hashing, and deliberately so, if even one single bit of the hash's input data were to be 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 untrackable hash results. But Microsoft's super-fingerprint avoids the information-lossy hash. So they have presumably retained ALL of the information contained within those 48 pieces of information individually.

That means that Microsoft's super-fingerprint can retain tracking, or more likely a tight association with the user's browser, even when some of the browser's data changes. They change screen resolution. The battery status changes, right, because that's one of the things. The battery percentage change would completely result in a different hash in the old-school fingerprint mode. Here, Microsoft can examine it, see that the battery charge changed but nothing else did, and go, okay, same person. And then update the fingerprint to match the new battery status and continue tracking. So it is literally a super-fingerprint.

And 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 were to log into, since they're logging into LinkedIn. And perhaps that information could be useful for some security purpose. So that I can believe.

The other thing that would also be interesting to check would be what, exactly, those 48 pieces of information are. I didn't dig into it. You know, there could 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 that information would be lost due to the hash, okay, you know. But if we assume that Microsoft is collecting and reversibly encrypting and forwarding all of that to their mothership, it would be interesting to see exactly what they're collecting and retaining because 48 individual things, that's a lot of things. You know, more than just screen resolution and battery charge and, you know, that kind of stuff. You know, browser user-agent string and so forth.

Anyway, The Next Web does have a bit more to say. They wrote: "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" - with LinkedIn's own - "sales tools [as noted before], Apollo, Lusha, and ZoomInfo. Because LinkedIn knows the employer of each registered user, systematically scanning for the presence of a competitor's tools gives the platform [LinkedIn] visibility into which companies are evaluating or deploying rival products." Because again, they know who the LinkedIn person's employer is.

"The list also reportedly includes tools associated with neurodivergent conditions, religious practices, 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," you know, GDPR. "Knowing that a user is running a job-search extension, for instance, is a meaningful inference about their employment intentions, drawn without their consent.

"The scale of the operation has grown substantially over time." In fact, this is really what's breathtaking. "LinkedIn," they said, "began scanning for 38 specific extensions in 2017," 38 extensions being scanned for in 2017. "By 2024, that number had grown to 461." Hey, if some is good, more is better. Nobody seems to be minding or complaining, so let's just scan their hard disk more widely. "By February of 2026," meaning a month

and a half ago, "the list had reached 6,167," yes, 6,167 individual extensions files being scanned for. They wrote: "BleepingComputer's testing confirmed the scanning was active as of early April 2026."

Leo: I don't understand why they are looking for specific extensions instead of just seeing what extensions are you running. Can't you ask the browser what extensions are running as part of the fingerprinting? I think you can. It's weird.

Steve: Well, they're probably going beyond extensions. I mean, what they're doing...

Leo: They're looking for all applications.

Steve: They're looking for file names. They are doing filename queries in...

Leo: Oh, that's so weird.

Steve: It is. And Leo, I don't think I have a link here to the BleepingComputer report. But you might grab it while I'm...

Leo: I'll pull it up. Yeah, yeah.

Steve: While I'm sharing this because it's bracing to actually see what - and I think it may have been Lawrence himself who did the research. So now we're at, as of February 2026, we're at 6,167, which they note is a 1,252% increase in two years. BleepingComputer's testing confirmed the scanning was active, as I noted, as of early April 2026.

Leo: So is this - oh.

Steve: Yup, that is - and you can see the IDs and then the files that are actually being scanned for - JavaScript, AVGs, HTMLs. So those are actual file names...

Leo: I kind of blame the browser for letting it do this.

Steve: I agree. And that's where - you're going to see me reach that conclusion here in a minute, that users need control over this completely out-of-control behavior.

Leo: Yeah. Huh.

Steve: So they wrote, The Next Web wrote: "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." So they're saying they're looking at more than 6,000 extensions because they're naughty. And, okay. They said that they do not use the data to infer The company added that it does "not use the data to 'infer sensitive information about members.'"

The Next Web wrote: "LinkedIn's characterization of the source matters. Fairlinked e.V. is connected to Teamfluence Signal Systems O, an Estonian company whose managing directors are 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, 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, however, which were verified independently. It does mean the framing of those findings is contested. The 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 euros, approximately \$334 million [at the moment], for processing users' personal data for targeted advertising without a valid legal basis." So LinkedIn was found to be using personal data for targeted advertising. "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's willing to adjudge 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."

Anyway, so they talk about LinkedIn having more than one billion registered users. Oh, and they did note: "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."

Okay. So we began this topic by, you know, 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

2.7MB unsolicited blob of code that does actually, as observed, confirmed, and reported by BleepingComputer, scan the mass storage file system of its website visitors. You go to LinkedIn.com, that happens to you, looking for the files belonging to at last count, and this is rapidly increasing, 6,236 web browser extensions which are arguably 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 illegal, it seems pretty clear that things - and I agree with you, Leo, on this point - 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 some form of control over, like, what the JavaScript that the browsers are hosting is allowed to do.

It would be nice if some sort of an alert came up and said, whoa, do you realize that the website that you have just gone to has made 6,232 queries of your file system? Like for different files by name of your file system? And how do you feel about that? Anyway, this JavaScript is rummaging around inside users' computers for a while, searching for all those files. Maybe we should be asked if we consent to having it do that. That just seems like, you know, crazy behavior on Microsoft's part. And they're only getting away with it because these people use the word "hidden" or "secret," you know, it just isn't obvious. In the same that that third-party cookies have never been obvious. So, you know, out of sight, out of mind. And this has been certainly something no one would be aware of otherwise. It's crazy.

But Leo?

Leo: Yes.

Steve: I know you know what our listeners do want to know about.

Leo: They want more ads; right? Hey, I did want to mention, this broke during the show, Anthropic we know was working on a new model, and the rumors - Mythos. And that the rumors were it was going to be so good that they were holding it back for fear it would be misused. Today they announced that they are going to allow a small number of people to use Mythos for security reasons. They say it is so good, strikingly capable of computer security tasks. They've already set it against existing browsers and operating systems, and they say found tens of thousands of vulnerabilities including in operating systems that are - in every operating system and every browser that is currently in use.

So of course it could also be misused. Their fear is to find vulnerabilities that can be exploited. So they have announced something today called Project Glasswing, which they are going to - let me go full screen on this - they're going to offer to a small number of companies. It is not going to be in wide availability. But they're going to offer it to...

Steve: So offer it to the good guys.

Leo: The world's most critical software to find those vulnerabilities before the bad guys get access to this model. Now...

Steve: Leo, we are living in a science fiction world.

Leo: Well, this could be just really good marketing. I mean, understand, what better way to say, hey, our model is good. But I think it's credible. And they say they've already found thousands of high-severity vulnerabilities. They've found Linux vulnerabilities have been around for 25 years. They found a FreeBSD escalation exploit that was very severe. So they are giving this away to a broad number of companies who have, you know, mission-critical software in the enterprise. They're hoping that they'll find those vulnerabilities before it becomes - they become public. So...

Steve: Or it's going to be turned over to, you know. And I think what I read about this, a rumor at the time, because it wasn't officially disclosed, is this would be a super high-end model that would cost more to use.

Leo: Yeah. The rumor is it'll be very expensive. We don't know. That's not been confirmed yet.

Steve: Right.

Leo: They are dedicating \$100 million in token credits to these companies so they can use it, you know, kind of freely. I think this is a big security story. Again, it could be just a very good marketing ploy. But I tend to credit that this is possible, and that the scores on the software benchmarks for this thing are off the charts, much better than their current 4.6 Opus.

Steve: No kidding. Even more so than Opus.

Leo: Yeah, much more so. Like 50% better than Opus.

Steve: Wow.

Leo: So they found a 27-year-old OpenBSD bug. They found a 16-year-old ffmpeg bug, you know, I mean, we're going to see more and more of this, obviously. And the fear is that people will use it for, you know, to find those flaws and exploit them. So they want to give his away to people so that they can find them and fix them before they get exploited. So that remote code execution in the FreeBSD that's 17 years old, I mean...

Steve: Remote code execution.

Leo: Yeah.

Steve: I thought it was a privilege...

Leo: That's a different one. That's in OpenBSD. This is in FreeBSD. This allows anyone to gain root on the machine running NFS from anywhere on the Internet. So it's a pretty - and it's been around for 17 years.

Steve: Wow.

Leo: So they say fully autonomously, no human was involved either in the discovery or exploitation of this vulnerability after the initial request to find the bug. So that's the concern. A bad guy will find it, and suddenly you've got all of these FreeBSD distros completely vulnerable.

Steve: Yikes.

Leo: 4.6 was able to exploit the vulnerability but required human guidance. Mythos did not. It could be a - and that's really scary, the idea that they could autonomously be...

Steve: Well, I mean, there's no question that the bad guys will pay whatever the cost is.

Leo: Sure.

Steve: And find exploits and jump on them. This is a, you know, we talked about exactly this. For a while, remember that - I don't remember if it was OpenAI or Anthropic were saying that the good news is the AI seems better at finding problems than it is at exploiting them. That's apparently changed.

Leo: That's not true anymore. And the other thing that I thought was really telling is Mythos is able to chain exploits, as many as six exploits. So that's a big deal because, as we've talked about many, many times, it's not often that just a single exploit is enough. Often these exploits are chained.

Steve: Right, right.

Leo: And the fact that Mythos can do it autonomously, that's a little scary. So that's a big story. We'll hear a lot more about it. We'll talk about it tomorrow on Intelligent Machines, and I'm sure effectively can. We've got a really good guest on Intelligent Machines tomorrow, Daniel Miessler, who is a security researcher, many years with security, worked at security at Apple and other companies, and is an AI aficionado. So he'll have something to say about this, for sure. Should be very interesting. All right. Let's take that break that you promised. I just wanted to mention this because this just happened.

Steve: Yes, I'm glad you did. That's going to be...

Leo: A huge story.

Steve: And there will be people who didn't check their code or that weren't, I mean, that didn't qualify for Anthropic's initial offer. And bad guys, there's just no question that we're going to see, you know, on the margins there will be new exploits that are found.

Leo: And I think, again, it could just be a marketing ploy. But I think I want to give credit to Anthropic for doing the right thing, to hold this back and release it this way so that people had a chance. Because otherwise I'm afraid we don't have much of a chance in this modern world. All right. Let's get back to Steve.

Steve: Okay. So while I was over at BleepingComputer, confirming that Lawrence Abrams had independently verified Microsoft's hard-to-believe JavaScript behavior - wow - I encountered the news that Microsoft has now also begun forcing upgrades of unmanaged Windows 11 PCs from 24H2 to 25H2.

Leo: Yeah, saw that.

Steve: Yeah. 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 report, 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 200K in size.

"Microsoft said in a Monday update to the Windows release health dashboard: 'The machine learning-based intelligent rollout'" - because of course, Leo, everything has to be...

Leo: Intelligent.

Steve: ...learning intelligent AI nonsense now. Otherwise it's no good; right? Those humans don't know what they're doing.

Leo: Right. Don't want dumb. Do you have any smart?

Steve: That's right, "...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." Yeah, right. We've been there before, until they start saying, do you want to do it now or later?

They said: "Those who don't want to wait for the automatic update 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 update, 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," writes BleepingComputer, "to help users resolve problems encountered during the Windows 11 25H2 upgrade process. Since the March 2026 Patch Tuesday updates were released" - that's just last month - "Microsoft has issued several emergency updates, including ones that address 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 management tool."

So I wanted to mention this because GRC's InControl freeware can also be used to give users control over exactly this process. It configures Windows to appear as if it's under management, thus it is not "unmanaged," 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 Windows 11 25H2, control can just be as easily reversed. So I just saw that, and I wanted to give everybody a reminder. You know, as we know, we had fun with it at the time, Leo, back in the Windows 7 days where the first version of this was called Never10.

Leo: Never10, yeah.

Steve: When I vowed to never be using Windows 10. So Never10. And then there was a temptation to do Never11. But because they promised us that 10 would be the last version, I thought, no, fool me once, you know, it's Lucy pulling the football out from under Charlie Brown. So I decided instead to create InControl, which would allow us - apparently there's going to be a Windows 12; right? So okay, I'm glad I called it InControl and not Never11.

Leo: Never11, Never12. You could do it, you know what, though, if you're really thinking you could get upgrade fees every year. Just think about it. I'm just saying.

Steve: Actually it's free, so...

Leo: Oh, shoot.

Steve: No, yeah.

Leo: You see, you see, you see?

Steve: I did want to mention that it's freeware, and it was with no bugs known at the time of its release, so it's still at Release 1, and it works perfectly, and it's free, and you can use it forever.

So 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 indeed we've since learned that none other than Cisco Systems became another victim.

BleepingComputer also reported on this. They explained: "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" - steal Cisco's 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" - and here it is - "GitHub Action plugin from the recent Trivy compromise." We'll be talking about action in a minute.

They wrote: "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."

I wanted to mention that one of the things that we're seeing now is that the bad guys know how to take advantage of the things they steal. Back in the early days, you know, we were seeing, like, AWS credentials were stolen. But then we didn't immediately hear that they were used, like, to ill effect for those from whom they were stolen. Now we're always seeing that. My point is that is another thing that seems to have changed is all the bad guys know how to take advantage of what it is they steal, and they know that their window of opportunity will be very short. So they immediately jump on it; and, you know, the attacks go broad and deep and wide so that they're getting the most bang they can for the buck, to the detriment of, you know, the victims of these.

They wrote: "BleepingComputer has learned that more than 300 Cisco GitHub repositories" - 300 GitHub Cisco repositories - "were also cloned during the incident, including source code for its AI-powered products" - of course Cisco's going to have that. Why fix the old ones? - "such as AI Assistants, AI Defense, and unreleased products." Whoops. Yeah. "A portion of the stolen repositories allegedly belongs to corporate customers, including banks, BPOs, and U.S. 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" - I bet they did - "regarding the breach, but has not received a reply to our emails," they wrote.

"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."

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

Leo: Probably could, if they wanted to.

Steve: Yeah. Wow. Wow wow. So I know from seeing the domains of our Security Now! Listeners who have email subscriptions, the email domains in email subscriptions that the Proton family of products are very popular among our listeners.

Leo: Interesting, yeah.

Steve: Yeah, I see a lot of @proton.com email in our subscriber base. 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. Proton explained, writing: "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."

Okay, now I'll admit I was somewhat surprised by Proton's claim of eavesdropping, and it appears that they're mostly referring to the leakage of metadata. Not that that's not a problem, but that seems to be their focus. Also, their information may be a bit dated and skewed. Their claims included links to Zoom, Google, and Microsoft. But, for example, the Microsoft link talked about Outlook. It's like, okay. 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. Anyway, I have the link to last week's Proton's "Meet" announcement in the show notes for anyone who may be interested in following up, although I'm sure if you just go to Proton.me, which is the domain of this link, that it probably comes up on their announcements.

And I mentioned we would be talking in the future about GitHub. Well, that future is now. The recent LiteLLM and now Cisco and now several other attacks that have all been attributable to the Trivy malware scanner, also involved GitHub's Actions feature. In the wake of these various messes, GitHub has announced that 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. In other words, whoops. And maybe sooner rather than later. So since Actions are now seeing some serious abuse, there's no time like the present to improve their security. And that is happening.

Also Cloudflare said, in their posting, just sort of in passing, that the use of - they're sort of checking back in. 1.1.1.1 remains super private. And this was posted on the eighth anniversary of its launch. They said: "Exactly eight years ago today" - they posted this, I guess Sunday - "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're aware, no other major public resolver has had their DNS privacy practices independently examined."

So anyway, their posting continues at some length. But they were just audited by one - and this is the point. They just had an audit by one of the top four accounting firms, and they again passed with flying colors. They do not want to know who uses their service nor what those users look up. DNS querying source IPs are anonymized on use and deleted within 25 hours. So they are really doing everything they can to honor their privacy commitments.

And Leo, we're at an hour and a half in. Let's take a break. And then I'm going to talk about the other very cool thing that Cloudflare has just done by basically creating a WordPress replacement.

Leo: Yeah.

Steve: Very, very cool.

Leo: Yeah, I think it's very interesting, yeah. We'll talk about the EmDash in just a little bit. I think that's what Steve's aiming at.

Steve: Yup.

Leo: Yup. Okay, Steve. On we go.

Steve: Okay. So, oh, I guess it was on April 1st. Because I started this with "Also on April 1st, Cloudflare announced." So it must have been on April 1st that they were talking about their 1.1.1.1 DNS. Also on April 1st, Cloudflare announced their EmDash. For those who don't know, there's the regular hyphen-style dash, then there are dashes that are the width of an N and the width of an M. And so EmDash is what they called this. I don't really know why. But it's kind of just a cool name. Anyway, it caught my eye. And I'm sure it will - actually many of our listeners sent the announcement to me, so I knew they were aware of it. Because its goal, EmDash, the EmDash project's goal is to replace WordPress with a far more secure successor.

And, you know, the expression "far more secure," in the case of WordPress, is not a high bar. Since every Security Now! listener knows quite well what a complete security disaster WordPress has become. And it's not WordPress's fault. The problem is its creaky old architecture, which Cloudflare has just completely replaced. The source of WordPress's trouble has been that it promises to allow anyone to author and offer an insecure WordPress plugin and that its architecture doesn't allow security containment of those. Unfortunately, many people have taken them up on this, and the result is a mess.

So here's what we learned from Cloudflare. They wrote: "The cost of building software has" - oh, yeah, AI. "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."

They wrote: "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 content management system (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. 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," as opposed to GPL. "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 development beta."

Okay. So all of that is super interesting. But what about WordPress's security; right? Before we get to that, Cloudflare felt the need to congratulate WordPress and kind of apologize, I think, a little bit, 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 recognizable 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 democratized 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 upon 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 comes: "WordPress's plugin architecture," they wrote, "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 in the previous two years combined."

Okay. In other words, things with WordPress are getting worse, not better. They wrote: "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 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 third-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's 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. In other words, you are locked in."

So that's the gist of it. Their posting continues at much greater length, but everyone gets the idea; right? Cloudflare has leveraged AI agency to take the conceptual promise that WordPress met and to dramatically overhaul its architecture for licensing freedom and security. Essentially, it is a plugin architecture where these dynamic workers are running as their own execution spaces. They communicate to this new EmDash core via an API, and the API calls that they are allowed to execute are strictly controlled by the manifest which they explicitly define and export right upfront. So it is the architecture you would design today if you wanted to do what WordPress has been doing, but to do it in a secure fashion. And on GitHub, all open, all open source. MIT license means you can do much more with it if you wish. So another big home run, I think, for Cloudflare.

Leo: Yeah, it's very interesting. You know, I hope they support it. And, I mean, clearly they want to take over the 40% of the Internet that's run by WordPress. And they probably have pretty credible reason [crosstalk].

Steve: Yeah, I think if we see a core of add-ons emerge which solve the problems that people have, and you look at the security model that it offers, you know, why, if you were starting from scratch, would you use a WordPress CMS as opposed to an EmDash CMS [crosstalk] either?

Leo: There's another reason for them doing this. They want to make something that is easily manipulated by AI, that an AI could build a website, easily build a website on. And [crosstalk] certainly do that with WordPress, I think this is designed specifically with that in mind.

Steve: I bet you are exactly right.

Leo: Yeah.

Steve: Okay. We're going to talk about the FCC ban on consumer routers. Let's take our final break, and then we are going to - I'm going to take our listeners through what happened so that it's, well, as I said at the top of the show, by the time we're done here, everyone will have a very clear sense and understanding for what it means and how it may change. We'll see.

Leo: Yeah. We shall see.

Steve: And why it makes no sense.

Leo: Yeah.

Steve: I mean, it really, really, really doesn't.

Leo: Good. Well, I'm really, really interested in your take. All right. Speaking of behavior, this isn't exactly the right behavior. The FCC has banned all routers made outside the U.S.

Steve: Which is to say...

Leo: All routers.

Steve: ...as we'll see, all routers.

Leo: Yes.

Steve: So 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 do not exist. All routers purchased by and available to U.S. 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 in the afternoon 15 days ago read: "FCC Updates Covered List" - this is the FCC's own press release. So their headline is "FCC Updates Covered List to Include Foreign-Made Consumer Routers." The press release 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." Okay. So 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.'"

It continues: "The Executive Branch determination noted that foreign-produced routers, one, introduce 'a supply chain vulnerability that could disrupt the U.S. economy, critical infrastructure, and national defense'; and, two, 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," they wrote, "have exploited security gaps in foreign-made routers" - which again, routers - "to attack American households, disrupt networks, enable espionage, and facilitate intellectual property theft. Foreign-made routers" - again, all routers - "were also involved in the Volt, Flax, and Salt Typhoon cyberattacks targeting vital U.S. infrastructure."

Leo: Yeah, well, it wasn't my Linksys. Okay, sorry.

Steve: Uh-huh. "As outlined below, today's action does not impact a consumer's" - here it is - "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 because they're all foreign, 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, we would hope. But that would mean that as WiFi technologies continue advancing and requiring support from new chipsets and new radio hardware, newer routers cannot be obtained from traditional foreign suppliers. Okay.

That happened Monday afternoon 15 days ago. By the end of that 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 and 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. He wrote: "On March 23rd, 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, which I read, which requires this conditional approval to be obtained, is 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," he writes, "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."

Okay. Now, 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, more than 10 years ago, we were encountering instances, and we discussed them on the podcast, where, for example, Intel's "demonstration only" source code for their UPnP implementation was unfortunately dropped directly into routers. This resulted in UPnP being bound to consumer routers' WAN-facing network interfaces, essentially by mistake. After delivering a podcast about that, the next week I announced that I had enhanced ShieldsUP! services to explicitly check for public UPnP exposure. But all of that was fixed back in 2013 and 2014, 12 years ago. At the time, many people were exposed. We fixed that. We as an industry fixed it.

Also back then, as in more than 10 years ago, we encountered instances where ISP-provided routers had open ISP admin ports. They were either using weak authentication credentials or known authentication credentials or 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 users' 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, none of those problems would be lessened by routers having domestic points of origin. Thus, nothing the FCC is attempting to do will fix anything that is now broken.

Scott continues, writing: "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 of 2020, published its initial Covered List in March of 2021, and adopted a Notice of Proposed Rulemaking and Notice of Inquiry on June 17th, 2021, initiating two separate dockets and inviting public comment. The Commission then adopted a Report and Order in November of 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."

Now, we could argue that's bureaucracy. And bureaucracy has overhead, and it takes time. But what it does is it tends to keep it from making mistakes. And as he said, just deciding to do it over the weekend and then doing it, well, you get the kind of things that we've been seeing from this administration for the last, what, year and three months.

"The comment process," he writes, "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, as this one, 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 [that we have now] 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." As they all are.

"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 a \$1.9 billion level, 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." Okay, maybe some waste, fraud, and abuse. "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." Fortunately that doesn't apply here. Nothing really changes for consumers.

He wrote: "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 [meaning what's happened now] has no comparable record, and former FCC officials have already predicted it will face legal challenge.

"Also," he writes, "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." Routers bad. "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." Meaning it still could have done more. He writes: "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." And to those who've been watching Washington recently, what a shock. "The stated justification is cybersecurity risk from foreign manufacturing. But the evidence the FCC itself cited undercuts the case for a foreign-country-of-manufacture approach. According to the Department of Justice, Volt Typhoon primarily targeted Cisco and Netgear routers, devices designed by American companies."

Leo: Yes. Yes.

Steve: "The routers were vulnerable, not because of where they were manufactured, but because those companies had stopped providing security updates for the discontinued models." And I'll just note 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."

Basically, the White House just, you know, waved the wand, you know, waved a hand and said let's outlaw foreign-made routers. Period. All those bad consumer routers we're, you know, we're outlawing them.

"An agency," he writes, "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 [meaning three days], it seems unlikely that they were consulted in any 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's 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 WiFi 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."

Leo: That's what this really is. This is about manufacture, not security. Right?

Steve: What a shock, yes. And he writes: "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 of '27, 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, FCC 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 without 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 implementation 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 again issued a broad national security determination covering all foreign-made drones, which the FCC implemented immediately. DJI has since sued to challenge the action."

And he finishes: "'Process' is not 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" - meaning the FCC - "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."

Okay. So where does this leave us? This apparently arbitrary and short-sighted ban will do exactly nothing to actually improve the security of any existing routers whose current models may continue to be sold. Since new routers cannot 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 planned models in the process of becoming ready for market next year. Later this year, next year. 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 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 their network security perimeter that establishes and maintains, through their NAT router, persistent connections. And many, if not most, of these phone home and maintain persistent connections are 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, actual damage to consumer security than any routers designed within the last decade. We have invited them inside our network. They can see our network's management interface at 192.168.0.1 or .1.1 or wherever it is. And no one's monitoring them. For all we know, they are busy right now trying to brute force that management interface.

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

Leo: Wow.

Steve: It's just, I mean...

Leo: It's so frustrating.

Steve: I wanted to share this because it really puts in context how arbitrary this was, how it doesn't achieve anything, whereas the previous ban on Huawei and ZTE actually did. It identified a problem, substantiated the problem, and surgically excised what needed to be removed; and then helped to pay for the removal and replacement of this Chinese gear that with good reason we decided we could not trust having inside the U.S. infrastructure any longer. Instead, this just says we're banning any new foreign-made routers. Crazy. I mean, it absolutely does nothing.

Leo: Yeah.

Steve: And one hopes that it is so errant that it'll somehow get fixed. It's just hard to - it's hard to believe it will continue.

Leo: Meanwhile, I don't know what we're going to do. I guess build, you know, you've got instructions, don't you, on your website to build like a pfSense router?

Steve: Oh. It is, yes, it is just not a problem any longer...

Leo: That's the thing to do.

Steve: ...to get a little fanless PC and use - actually, OPNsense is now the one to switch to. PfSense is, you know, OPNsense grew out of it and is, I think, probably the better choice.

Leo: And you put it on - what is that little box that you put it on? I can't remember the name of that.

Steve: It's a cute little thing.

Leo: Yeah, it's a cute little thing. Somebody in the chatroom will see because I know they...

Steve: Yeah. It's called the SG-1100, although it has nothing to do with Steve Gibson. But...

Leo: Well, that's it. That's all you need is the SG-1100, and you run OPNsense. I'm sure you can run - it's from Netgate.

Steve: Yes, Netgate, that's right.

Leo: Netgate, yeah.

Steve: And that's a cute little guy.

Leo: Now, they say pfSense, but I suppose you could run OPNsense.

Steve: Yes. And they are - they are still with pfSense. And there's nothing wrong with pfSense. But if I were - there's another little platform, I bought one on Amazon, it's the one I'm using over at my place with Lorrie, it's got four network interfaces. I can't remember the name of it now. And on it I'm also using pfSense just because I know pfSense.

Leo: Right. Start with OPNsense.

Steve: But if some of you were just starting out, I would say Netgate. And this little guy, now, this is not a radio router. There is no antenna there. So you need to solve the radio connection problem separately.

Leo: Right, right.

Steve: But it is a, you know, pfSense is a very powerful little routing software.

Leo: Oh, yeah. And it has all the security features you could want, really. I mean, it's...

Steve: So, you know, we're up to WiFi 7 now. And so what this is saying is we cannot ever get a WiFi 8 router because that would require a new...

Leo: Unless Elon decides to make it.

Steve: Yeah. You know, it's just nuts.

Leo: 40% of all the routers in the United States are made by TP-Link, a Chinese company. That's the one they were going to ban initially. Then they thought, they said, well, you know, if we're going to ban them, let's ban them all. Even if they're made in Sweden. It doesn't matter. I was at RSAC, and I went over to the Ubiquiti, because I use Ubiquiti gear, I went over to the Ubiquiti booth and said, hey, I'd just like you to comment on this FCC decision. It had happened that morning. And they did not want to talk about it. There's no - and I understand why. It's no-win. What are they going to say? We don't like the idea. No. We like the idea. No.

Steve: No comment.

Leo: No comment's the only thing you can say. The only safe thing to say. Unless they have plans to build a factory in, you know, Piscataway, I don't know what else they're going to do. I think that's the plan; right? Well, everybody just has to build a factory in the U.S. And god knows, there are lots of people out there who are just dying to build routers in a router factory. They are just lining up.

Steve: Well, and all the consumers will be happy to pay an extra hundred bucks for...

Leo: Yeah, oh, that, too.

Steve: ...domestic manufacturing.

Leo: Yeah.

Steve: To get nothing in return. I mean, it's like...

Leo: Well, the Cisco routers are no more secure. Just because they're made in the U.S. doesn't mean anything.

Steve: No.

Leo: In fact, they're less, frankly. Yeah. And I'm sure this Netgate is not made in the U.S.

Steve: No. Nothing is made in the U.S.

Leo: Nothing's made in the U.S.

Steve: Nothing. Nothing. We don't make things here.

Leo: No. Well, that's a mistake. I mean, admittedly, we probably should have...

Steve: Well, the mistake is like rattling sabers between us and China. We ought to recognize that it is a good relationship, it's mutually beneficial, and we ought to just have a dtente here and say, look, you know, you want us to buy your stuff. We want you to make the stuff that we're going to buy. So come on.

Copyright (c) 2014 by Steve Gibson and Leo Laporte. SOME RIGHTS RESERVED

This work is licensed for the good of the Internet Community under the Creative Commons License v2.5. See the following Web page for details:
<http://creativecommons.org/licenses/by-nc-sa/2.5/>