



## KongTuke's CrashFix

**Description:** The lowdown on last week's "no turn" picture of the week. Is an AI-driven hacking campaign a big deal now? Clause used in multiple Mexican government attacks. Apple continues to be confronted with age restrictions. COPPA needs an exception to allow age collection. Meta swamps law enforcement with AI-slop CSAM reports. Roskomnadzor has been busy blocking VPNs. Guess how many. The UK tries to report their self-scanning success. Remember that hacker who extorted the psychotherapy patients. Scattered Lapsus\$ Hunters is actively recruiting women. Cisco lands another breathtakingly rare 10.0 CVSS. VulnCheck's report on 2025 vulnerabilities and exploits. Steve discovers a fabulous \$72 Hardware Security Module. A listener shares an interesting AI service discovery. The very potent "ClickFix" exploit evolves.

High quality (64 kbps) mp3 audio file URL: <http://media.GRC.com/sn/SN-1067.mp3>

Quarter size (16 kbps) mp3 audio file URL: <http://media.GRC.com/sn/sn-1067-lq.mp3>

---

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. A show we recorded a little bit early because we're going to Zero Trust World in Florida. We have lots to talk about, though. Jam-packed programming. We're going to talk about Scattered Lapsus\$ Hunters. They're looking for female voices for their social engineering. AI hacking. Is it here? Yes, it is. And a very potent ClickFix exploit. When you see how this works, you might wonder how you didn't get bit by it. All of that coming up next on Security Now!.

**Leo Laporte:** This is Security Now! with Steve Gibson, Episode 1067, recorded Sunday, March 1st, 2026: KongTuke's CrashFix. A weird Sunday edition.

**Steve Gibson:** A one-off.

**Leo:** Yes. Let's never do it again.

**Steve:** Let's never do it.

**Leo:** When's your flight, Mr. G?

**Steve:** We head out of San Diego tomorrow morning.

---

**Leo:** San Diego? Why that?

**Steve:** Because we have a shitty little airport here in Orange County.

**Leo:** I like the John Wayne. Oh, but you probably can't get to Orlando from there.

**Steve:** Well, remember, I couldn't get a flight to your airport, which is similarly sized.

**Leo:** Similarly shitty, yes.

**Steve:** With the Charlie Brown statues.

**Leo:** Yes.

**Steve:** I couldn't get it. And it happened it was on my birthday, of all things, that there was finally direct service for the first time from my shitty little airport to your shitty little airport.

**Leo:** Yeah. So we can come visit you now.

**Steve:** There's no excuse.

**Leo:** Yeah. Hey.

**Steve:** But anyway, if we didn't leave out of San Diego, and we left out of Orange County, an eight-hour layover was the only way to get...

**Leo:** Oh, that's absurd.

**Steve:** ...between point A and point B.

**Leo:** OMG.

**Steve:** So, yeah. So we don't get in...

**Leo:** So how long does it take you to get to San Diego?

**Steve:** Well, so it's about - and well, the problem is that we come down the 5. And San Diego is no longer a sleepy little port, you know, a little military port. It's a big deal. So now its commute traffic has gotten really bad. So but the good news is we don't actually

depart until 11:30, so we'll probably be catching the tail end of the commute traffic. Anyway, so we we're going to head out at 7:00, go down, park. I have a buddy who flies out of San Diego all the time because he's a satellite engineer.

**Leo:** For the same reason, yeah.

**Steve:** Yes. And so there's little tricks. Like you can pre-purchase your parking online and get a QR code, so you just scan yourself in, and other little goodies.

**Leo:** So how long does it take?

**Steve:** So an hour and a half.

**Leo:** Oh, it's kind of like going to San Francisco for us.

**Steve:** If there weren't traffic. Yeah, I mean, you can't get there. In fact, a nephew of mine who lives in Napa and works for Salesforce sometimes has to go into the city. And he doesn't even try. He gets on a ferry somewhere and, like...

**Leo:** Yes. You go to Larkspur, take the ferry, yeah.

**Steve:** And when my sister visits her son, my nephew, they fly into Sacramento because that's better than, I mean, anything is better than trying to get to Napa from San Francisco. It's - of course you guys have a goat trail that connects you to - which is the way you like it; right?

**Leo:** Yeah, we're going to SFO because for the same reason, it would be ridiculous to fly out of Santa Rosa.

**Steve:** So basically it's an all-day trip.

**Leo:** But the good news is we just get a car service, and we get in there, and we sit back, and we take a nap, we listen to our - it's just we don't do anything. And it is only an hour and 15 when there's no traffic, which we're leaving, our flight is at 1:30. So we're not leaving until about 10:30.

**Steve:** Oh, nice. So you'll be able to go midday or late morning.

**Leo:** So there won't be any traffic.

**Steve:** And so basically we head out at 7:00, and we arrive at 7:00. Is it 7:00 or 9:00? Its late because of the three-hour time change.

**Leo:** Yeah, yeah, we have a late flight, too, yeah.

**Steve:** Yeah, time change and so forth. And then I saw that you guys are grabbing a car. We're also going to grab a car. So we'll have wheels.

**Leo:** Yeah, because you're going to visit. You're going to see people.

**Steve:** Yeah.

**Leo:** We're going to Kennedy Space Center, which I'm very excited about.

**Steve:** The big dilemma has been resolved, and that is who's going to feed the squirrels. Because when we went off on our European trip...

**Leo:** You can go live. Did you go live already, John?

**Steve:** Oh, we've been live.

**Leo:** Oh, good, okay.

**Steve:** Yeah. When we went off on our European trip, we found this big bucket...

**Leo:** The squirrels all died.

**Steve:** No, no, which we filled with peanuts. At that time we were feeding them peanuts in the shell, which we thought was cute until the entire, I mean, there were shells everywhere after two years.

**Leo:** Yeah, well, they're not going to eat the shells.

**Steve:** Oh, my god.

**Leo:** They're not dumb.

**Steve:** But one of our security cameras showed - so the idea was - I actually have photos. They're pretty cute. We have a huge bucket with a little squirrel ladder going up the side. So they were able to climb...

**Leo:** Oh, you know what, Lorrie is just like Lisa. Lisa feeds the turkeys.

**Steve:** Well, and so...

**Leo:** And as a result the turkey population soared.

**Steve:** That's the problem. That's the problem, is first of all our squirrels are fat, and there's many more of them than there used to be because they realized, hey, we have a food source.

**Leo:** The population soars.

**Steve:** So we can expand our population. Consequently, there's a big issue with our leaving, like moving in a couple months into our new place. Because I'm determined - my problem is I effing hate crows. We have these crows the size of...

**Leo:** Oh, you're a crow hater. Oh.

**Steve:** I'm a crow hater. They're like the size of a dog.

**Leo:** And they're smart. They're so smart.

**Steve:** They're smart. They're diabolical. I hate them. And so they, like, wait till no one's watching, and then they swoop down and steal walnuts from the squirrels. Well, these are not your walnuts, you crows. Anyway, so we are determined not to do this again. When we move, nothing's being fed. Hopefully we've had our fill of it.

**Leo:** That's what we did. We moved away so we didn't have to feed the turkeys anymore. And Lisa did the exact same thing. She said never again.

**Steve:** Anyway, we have house sitters so that the animals will continue to be fed the diet to which they have become accustomed and spoiled.

**Leo:** Oh. But when you move, you're not going to do that anymore.

**Steve:** No. We're not. We're not going to make the mistake again. Anyway, what I was going to say was that, so we left this big barrel with this cute little squirrel ladder coming up the side. And I happened to check, I don't remember, we hadn't gotten very far on our trip when I checked one of our cameras, and the barrel had been knocked over, and the nuts had been spilled, and it was just a free-for-all.

**Leo:** Oh, boy.

**Steve:** Now, it wasn't ever very practical because squirrels are like monopole magnets. They repel each other.

**Leo:** Oh. Only one at a time, huh?

**Steve:** They're not herd animals. They're not like a pack of wolves.

**Leo:** Right.

**Steve:** Yes, something about they - I don't know what they're, like, if they're horny or what's going on.

**Leo:** This is fun. This is why you do this, because now you can observe their interesting behaviors.

**Steve:** Oh, yes. We are quite nature observers.

**Leo:** This is what happens if you get all - although I just saw an article that said bird watching is a good way to avoid dementia. I don't know. I am not starting to bird watch. But apparently it's good for your brain.

**Steve:** Maybe the problem is if someone is a bird watcher, you can't tell whether they're demented. It's like, well, they watch birds. Oh, that's great. No, that might be the first symptom.

**Leo:** I have your five, count them, five ads. Nothing about being on a Sunday makes any difference when it comes to that. I have your show notes, KongTuke CrashFix.

**Steve:** And you do have a deadline in three hours.

**Leo:** So we should get going.

**Steve:** I'm ready when you are.

**Leo:** Are you ready? I want to continue this story. I'm fascinated. We miss our - you know what we miss the most is Bob the peacock. But, yeah. We still have crows, but we don't feed them.

**Steve:** And I'll tell you, Leo, when they come to the glass, our back glass door, stand up like little meerkats, and then put their paw on and look in, like where's the nut lady? Where's the nut lady?

**Leo:** The nut lady's gone.

**Steve:** It's Sunday brunch. Where's the nut lady?

**Leo:** Where's the nut lady? Nut lady feeds us. Oh, that's so cute.

**Steve:** However, we do have very prolific pine trees. And so although they have to work - it's probably good for them, right, to chew on something a little harder than a - we have weaned them from the walnuts. Walnuts make them almost go into some sort of a drug haze. So peanuts. We, like, stepped them down. Also the expense. Oh my god, we were spending so much money on walnuts.

**Leo:** Okay. So we now go to the - well, we don't do this anymore. Same thing, Lisa was buying premium sunflower seeds, like for humans.

**Steve:** Organic walnuts, honey. Only the best for our squirrels.

**Leo:** So she started going to the feed store to get the food.

**Steve:** Their fur is soft and glistening. They've got a high oil content.

**Leo:** They're getting all those delicious tannins. Oh, yum yum.

**Steve:** That's right.

**Leo:** That is - you know what, Lorrie and Lisa are going to get along. That's very cute. Wow. I'll have to tell Lisa. She will like that story. She totally will get it. We don't feed the deer. We have deer, and we have crows, but that's about it. We're going to do bird feeders, I have a feeling. I think it's...

**Steve:** I have a laser that is strong enough to hurt a crow.

**Leo:** Oh, don't hurt the poor - they're smart.

**Steve:** No. I'm on the verge, but no. The other problem is we have lots of planes flying around, and they're not happy with lasers, so I have...

**Leo:** Oh. No. Especially the green ones. Stay away.

**Steve:** ...avoided - I've avoided. But, you know, a black crow, that would absorb the laser energy really well. I've spent some time fantasizing about...

**Leo:** Don't ever tell anyone. After the portable dog killer story you've got to be very careful now.

**Steve:** That is true.

**Leo:** You don't want Homeland Security knocking at your door. Actually, they probably hate crows just as much as you do. All right. Let's go. Show time, ladies and gentlemen. Thank you for joining us early. All you Club TWiT members got the notification. Ooh. Oh, you know what happened? I have a timer at 11:00 a.m. I had set it to Security Now!, and I have a timer at 11:00 a.m. that automatically changes it to TWiT. But it won't do again, I don't think.

**Steve:** That's nice. You've got automation thanks to Claude something.

**Leo:** Oh, it's totally automated, yeah. Except for that doesn't work so well when I do a show off schedule. Oh, everything's automated. In fact, because I was leaving town, I had to move all of my system D services over to the framework, which stays on. And I had to move my Claude, I had to turn on remote control on my Claude and put it in Tmux so I can log in. I am - I don't know if I can SSH, actually, come to think of it.

**Steve:** What's happening with the house remodeling project?

**Leo:** We're almost done.

**Steve:** Oh, yay.

**Leo:** The scaffolding is gone. The stucco is finished.

**Steve:** They sent that all to New York. New York got all the scaffolding.

**Leo:** You don't want to know how much we paid for that scaffolding. Actually, we caught the contractor overbilling because Lisa wisely went to the scaffolding company and said how much did you charge our contractor? And he was, like, doubling it. So we're waiting to hear back from him. So we're waiting to hear back from him. He suddenly has gone radio silent when Lisa confronted him with the evidence, saying, well, we'll pay you what you paid them plus your, you know, 20% profit. But we're not going to pay you 50% on what you...

**Steve:** We also saw a scaffolding-related cost. I didn't realize that having people up on a scaffold you pay extra for those because of...

**Leo:** Oh, because of liability, yeah.

**Steve:** Yes, hazard. Hazard pay.

**Leo:** Yeah. It's a learning process; isn't it.

**Steve:** Well, so now all of Club TWiT and our listeners are figuring out why we just wait for the recording, rather than...

**Leo:** Yes. You don't want to hear this. You shouldn't have to hear this.

**Steve:** We love our live listeners.

JOHN: I think some people would actually love to hear Steve's opinions on crows.

**Leo:** Well, if you're in the Club, we do preserve the pre-show and post-show conversations.

JOHN: Hey, this will be in the TWiT+ feed.

**Leo:** It's in the TWiT+.

JOHN: I'm going to let everyone know this will happen in the TWiT+ feed as Steve's hatred of crows.

**Steve:** Crows would tend to absorb radiation quite well.

**Leo:** Especially those green lasers. I didn't know it, but the green lasers are an order of magnitude more power than the red ones.

**Steve:** This thing will pop a balloon. It's really cool.

**Leo:** That's everywhere.

JOHN: On that note.

**Leo:** Exploding crows. Now you know why they call it a "murder of crows." All right. Here we go. It's time for Security Now!. Hello, everybody. Normally I would say you wait all week for Tuesday; but if you're watching live it's Sunday, March 1st. Steve and I are headed off to Orlando, Florida tomorrow for the incredible Zero Trust World conference put on by ThreatLocker. So we thought we'd do Security Now! a little early. Those of you who listen after the fact will get the show at the same time, so you're going, what are they talking about? But, you know, the only reason I mention this, Steve, you probably want to mention it, too, is that if anything happens on Monday, it won't be in the show until next week.

**Steve:** Well, and this has been actually a problem I've been conscious of because I've now got in the habit of preparing Tuesday's show on the previous weekend, Saturday and

Sunday. So already things are like that. And there have been a couple times where I've made notes for the following podcast, or - and I try to make a note of this. There have been edition numbers of the show notes where I've, after the mailing, which, by the way went out yesterday in the early evening, everybody got that, I've made notice that, you know, okay, I've updated the show notes because stuff has happened since. So anyway...

**Leo:** Yeah. I have to do that for our shows, too. It's just a - because we want to be up to date.

**Steve:** So March 1st, I assume your NASes have reported in, as mine have, that all looks good, nothing to see here at the first of the month.

**Leo:** On the first of the month your NAS says hello?

**Steve:** Yeah.

**Leo:** Mine probably does, too, but I don't....

**Steve:** Just checking in.

**Leo:** I have a folder where all the NAS messages go, and I don't ever check it. So I'll check it.

**Steve:** Okay. So we're going to talk about a bunch of things. This is a jam-packed news and opinion, a little editorial, which seems to be what our listeners prefer.

**Leo:** Oh, yeah, we care about what you think, for sure.

**Steve:** I called this KongTuke's CrashFix, which is a tongue-twister.

**Leo:** What's KongTuke?

**Steve:** Unless you're a Klingon, in which case KongTuke.

**Leo:** KongTuke.

**Steve:** Very much like, yeah, very much like Klingon. It's the name that is - I can't remember the name. I mean, I can't remember the security firm. We'll get to it. It's just it's a bad guy's moniker that one of the firms have...

**Leo:** They're obviously Klingon fans, or Star Trek fans; right? Yeah.

**Steve:** Yeah. I mean, where would KongTuke come from?

**Leo:** It does sound like - yeah.

**Steve:** Normally the names come from the reverse-engineered code where some reference is found to, like, you know, the KongTuke.com domain or something. Anyway, we don't know. But there is a - it's been an evolution of this problem that Microsoft is going to have to, as they used to say, I don't know when, belly up to the bar and fix.

**Leo:** Probably not the best way to fix it. Maybe that's how the bug happened in the first place.

**Steve:** Okay. So we're going to start with the lowdown on last week's no-turns-allowed Picture of the Week, which captured our audience's imagination like few others have, although we've got another good one this week. We're going to look at whether an AI-driven hacking campaign is a big deal now. And Claude used in multiple Mexican government attacks.

**Leo:** Yeah.

**Steve:** Apple continuing to confront age restriction legislation. Got some on that. Also it turns out that COPPA, the child protective act, is going to need an exception for the age collection which other legislation is now requiring.

**Leo:** That should be a hint that there's something wrong here.

**Steve:** Exactly.

**Leo:** Oh, we don't want to protect kids online when it comes to that.

**Steve:** Yeah. Exactly. Also Meta is using an AI which is - I'm noticing also, Leo, this term "AI slop" just immediately achieved traction. Like everyone knows what AI slop is. It's surprising how quick the adoption was. Anyway, we've got AI slop CSAM reports that are drowning law enforcement in false positives. We'll take a look at that. Also our favorite Internet watchdog, Roskomnadzor, has been busy blocking VPNs, but you will never believe how many. The UK makes an effort at reporting on the success of their self-scanning initiative, although there's something fishy about their report, which we're going to look at. And Leo, I knew when I saw this you would remember that hacker who was extorting psychotherapy patients whose data had been exfiltrated from their psychotherapy centers.

**Leo:** How low can you get?

**Steve:** We've actually heard back about this process. That was in 2020, so six years ago. Anyway, he's back in the news. It turns out that Scattered Lapsus\$ Hunters is actively recruiting women, and we're going to find out why. Cisco lands, boy, no one does it like Cisco, another breathtakingly rare 10.0 CVSS. Just, you know, duck and cover, as they used to say. We've got VulnCheck's report on 2025 vulnerabilities and exploits. Just a little tip of the iceberg there. That's probably going to be our topic next week because there's lots of juicy information there.

I have discovered a fabulous \$72 hardware security module that does all my code signing, multiple certificates, open source, it's fantastic. I'll be talking about that a little bit because I know that from previous feedback from our listeners, anybody who needs to sign code needs something like this. We have a listener sharing an interesting AI service discovery. And then the very potent ClickFix exploit is evolving, now being used by the Klingonese outfit KongTuke for something called CrashFix.

**Leo:** Oh, boy.

**Steve:** And of course what would a podcast be without a Picture of the Week? And I've already had a lot of feedback saying he should have used a different screw. A security screw would have been better. It's like, okay, thank you. That's true. You'll see. You'll see.

**Leo:** Yeah, for last week's picture we got a confirmation from a number of people that that is a real picture from Canada, and that really did happen.

**Steve:** And the local government was embarrassed. I have a link to the actual story saying, uh, we're sorry.

**Leo:** That was a dumb thing to do. Hey, it's a very special Sunday Security Now! with Mr. - I didn't introduce you, but I think everybody knows that's Steve Gibson.

**Steve:** If they're here, they're like, okay.

**Leo:** They know who you are.

**Steve:** Move on.

**Leo:** Yes. And we are so glad to have you on the show this week. Glad you're watching, those of you who are live and figured out we were going to be doing this early. We're glad you're here. Our show today brought to you by...

**Steve:** Those of you who are dead, you know...

**Leo:** No, yeah, you wouldn't know.

**Steve:** You've not watching anyway.

**Leo:** You're not missing a thing, though. Well, maybe. I don't know. I think if you've passed, you don't have to worry about security so much.

**Steve:** Opinions vary on that topic because...

**Leo:** Yes. We don't know. We don't know. We'll just have to wait and find out. Okay. I'm ready for the Picture of the Week.

**Steve:** So our caption on this photo, this is Dad saying to, because he's a dad, one of his kids: "This is the last time I'm going to tell you to turn down the volume of what you call music."

**Leo:** Oh, Dad. Dad found a solution.

**Steve:** Yes, he did. And given the location of the little volume indicator dot on the volume control, which is like right at minimum it looks like now, doesn't look like Junior gets to turn this up very high.

**Leo:** OMG.

**Steve:** And now you can see why one of our listeners said he should have used a security screw, you know, where you can only screw it in, but the Phillips head is unable to get a grip when you're trying to go in the other direction.

**Leo:** It does look like somebody has tried to unscrew it, actually.

**Steve:** I think that the drill skittered a little bit on...

**Leo:** Oh, maybe that's it, yeah.

**Steve:** On, yeah. So for those who are not looking at the video, or don't have the show notes, I'm sorry. What we have is what we would call an old-school volume limiter. The problem, of course, is that the kids have a stereo system which they just are unable not to turn up so that it's bugging Mom and Dad who can't think, not only due to the nature of the music, but its volume. So finally, at the end of his rope, Dad has come up with a solution. He's drilled a hole in the side of the volume knob with a screw sticking out of it about an inch. And then another screw in the faceplate of this stereo such that the screw that rotates as you try to turn the volume up will hit the limiting screw, preventing it from going, looks like maybe more than maybe level 2 or 3.

**Leo:** Not much, yeah.

**Steve:** Yeah, not very much. So clearly, regardless of the backstory here, this is obviously someone's determined effort to prevent the volume control from being turned up very far.

**Leo:** Burke points out this would be good in a nightclub, where patrons tend to go over to the sound system and crank it; you know?

**Steve:** Yup.

**Leo:** Or your neighbor snuck into your apartment.

**Steve:** As a frequent patron of restaurants, I've had the experience where I'm in early for dinner, and the crew of workers have been there. They turn the volume up and then leave it up. And they, you know, they just forget. And it's like, god, can you turn the volume down? I can't think.

**Leo:** Don't you know it's early bird time. You've got to turn it down.

**Steve:** So a solution has been found. I of course would have put a small resistor network in line with the speakers in order to take the energy out of the speaker line. And then Junior would think, god, what happened? Did I blow the amp? It's not working the way it used to.

**Leo:** It's so quiet. Maybe I'm deaf.

**Steve:** Okay. So I wanted to thank many listeners who were made curious by last week's Picture of the Week. And Leo, you heard from many people, too.

**Leo:** I did.

**Steve:** Just to remind people, that was the street which was the stem of a T intersection. So the street we were seeing was leading up to a T intersection in the distance. And signage which would be encountered as the driver was driving toward that T intersection indicated that neither turning left nor right would be legal. Thus I gave the caption, "But Officer..." to the picture. Thanks to listener research, of which there was much, and some used AI, you know, asked AI to track this down, we now know that the photo, first of all, was not synthetic. That was my, you know, a common thought was, oh, come on, that was just Photoshopped. It was bizarre, but authentic.

And after the photo went viral a few years ago, it became a significant embarrassment to the local government who was responsible for its emplacement. The location was a town called Simcoe in Norfolk County, Canada. And a news report that one of our intrepid listeners found and shared explained that "Drivers please note that signs were installed this week which restrict left and right-hand turns at the intersection of Crescent Boulevard and Queensway in Simcoe. The intent of the new signs was to make Crescent Boulevard a dead-end street. The signs have been removed." So anyway, in other words, the signage was technically correct, and it was like up to you to come to this stop sign,

having seen that you can't turn left, can't turn right and, what, do a U-turn, as if it dead-ended at that point, rather than allowed you to cross into the cross street.

**Leo:** That's just crazy.

**Steve:** So anyway, my favorite quip about last week's photo was provided to us by a listener, Joseph Rork, who noted: "Despite the presence of the Tim Horton's in the background, we know this cannot be Canada. Otherwise there'd be a line of cars sitting at the stop sign." So many thanks to our listeners among the more than 20,000 who received a weekly mailing and whose imaginations were captured and took time to do the research and/or comment. Anyway, and also a big thanks to whoever it was who sent that to me in the first place. You know who you are. And I ask our listeners to keep them coming because they're fun to share.

Okay. So the headline in the news last week was - this is the headline: "AI-driven hacking campaign breaches 600+ Fortinet devices." Now, I'm going to first share the news report. Then I have a few things to say about it. The reporting says: "A Russian-speaking financially motivated threat actor used commercial AI toolkits to hack more than 600 Fortinet firewalls. The campaign began at the start of the year, around January 11th, according to the AWS security team. The attacker did not exploit zero-days or older vulnerabilities. Instead, they targeted FortiGate devices that had their management ports" - oh lord - "exposed online, used weak passwords, and didn't have multifactor authentication enabled." Okay. So to interrupt here, FortiGate devices, publicly exposed management ports, weak passwords, and no other authentication required. So no flaws were used, just very poor configuration hygiene.

The story continues: "Once inside, the attacker employed a collection of scripts that AWS says were written by AI tools. While AWS did not name products, researchers from Cyber & Ramen and Ctrl-Alt-Int3l identified them as being Claude and DeepSeek. DeepSeek was used to create scripts to perform reconnaissance and extract configurations from the hacked devices, while Claude was used to generate scripts for vulnerability assessments and to run offensive tools against the networks.

"Since this is the intersection of AI and infosec," writes this story, "the report generated a tornado of feedback and opinions on social media. The general consensus was that the threat actor wasn't particularly sophisticated, which AWS also believes. AWS CISO CJ Moses said the attacker was more interested in scale than value. Every time they encountered errors caused by hardened or non-standard internal networks, the attacker just moved on to a softer target.

"Once they did move laterally from the Fortinet device, the attacker compromised the victim's Active Directory environment, extracted database credentials, and tried to gain access to backup infrastructure. This led everyone to believe the threat actor was a relatively low-skilled initial access broker, right, an IAB that gain initial footholds on corporate environments and then sell their access to the hacked network on underground portals."

Okay. So I think it's entirely expected that anyone who has any need for any sort of code or scripting for any purpose whatsoever will increasingly be using AI. That's just today's reality. Good guys are doing it, and bad guys are doing it, and there's no reason to expect AI to be able to discriminate between the two. A high-level language compiler doesn't know or care who's using it, or to what purpose the code it's helping to produce will be put. Right? That's not its job. So the fact that we have now chosen to give consciousness-emulating large language models the marketing label of "Artificial

Intelligence" should not and does not automatically mean that these new tools somehow have responsibility for what they're being asked to produce.

So okay. But don't these AI tools make attackers more powerful? Yes, they do. And they also make the good guys more productive. That's why everyone, both good and evil, is now using them. In the current instance, there's nothing inherently wrong with a script that performs a vulnerability assessment. White hat security researchers employ such tools to aid their beneficial research much as bad guys may use the same tools to perform pre-attack vulnerability assessments.

My point is that any social media hysteria arising from the fact that AI "was involved" is now ridiculous. If you encounter it online I would recommend meeting it with a shrug and clicking on the "thumbs down" button. This is just the way the future is going to look now. It may have surprised us a few years ago, but it should surprise us no longer. And "AI" should not receive any of the blame for the way its creators, we humans, choose to use it. It's a tool and nothing more. It has no social obligations or responsibilities. It's not accountable. We are.

**Leo:** I like that. Because that eliminates that whole issue of AI safety, to be honest.

**Steve:** Yes, yes. Which, as I said, we might as well give up because we're not going to get it. And again, you know, we called it "artificial intelligence." It's not intelligent. It doesn't know anything. It's a very powerful new tool. But it's still a tool, and it's not responsible for the way we use it.

**Leo:** As usual, it's the humans. We're the problem.

**Steve:** Exactly. Okay. Now I'm going to give everyone a quick self-test to see whether the point I hope I've just made has had the chance to sink in. Perform a self-assessment to see how you feel about this next piece of news. It reads: "A hacker has stolen more than 150 gigabytes of data from multiple Mexican government agencies. The attacker allegedly used Claude to assemble scripts to gain access to government networks. According to Bloomberg, the attacker breached and stole data from Mexico's tax authority, national electoral institute, and several state water utilities. The stolen data covers 195 million taxpayer and voter records, government employee credentials, and civil registry files."

Okay. Should we care at all that AI was employed in these attacks? No. The fact that Claude was used in these attacks appears to be the highlight of Bloomberg's piece because they've got, you know, they're looking for clickbait; right? You know? It was certainly the headline which they attempted to make inflammatory. Eventually the world will get used to this, and it will just be assumed. And I hope everybody listening to this podcast will be in the lead on that because, again, that's the technical reality here.

Another technical reality is that Apple appears to be feeling the pressure to respond to the growing legislation-driven need for the providers of Internet services and online apps and app apps, you know, Apple Store apps, to know and to respond to the age of their users. Last Tuesday, Apple posted an update to their developer portal, addressed to their app developers. So this was written - when you see the word like "your app." So this is written to app developers.

They said: "Today we're providing an update on the tools available to developers to meet their age assurance obligations under upcoming U.S. and regional laws, including in

Brazil, Australia, Singapore, Utah, and Louisiana. Updates to the Declared Age Range API are now available in beta for testing. For Brazil: Developers who are distributing apps in Brazil can use the updated Declared Age Range API to obtain a user's age category. Age categories for users in Brazil will be shared when the user or a parent or guardian (where relevant) agrees to share the age category with you. The API will also return a signal from the user's device about the method of age assurance. For developers distributing their apps in Brazil, if you identify that your app contains loot boxes through the age rating questionnaire, the age rating of your app on the Brazil storefront will be updated to 18+.

"For apps rated 18+ in Australia, Singapore, and Brazil." And if this is all seeming like a big mess, you're getting the clue here, yes. They say: "Starting February 24th" - which was last Tuesday, the date of this announcement, in other words, you know, why this was posted - "Apple will block users in Australia, Brazil, and Singapore from downloading apps rated 18+ unless they have been confirmed to be adults through reasonable methods." And, boy, I hate that kind of language. Like, okay, you know, it's like any legislation that is written that isn't airtight is subject to interpretation. And it's like, oh, let's let the attorneys sort this out. Oh, god. Through reasonable methods. Whatever that is.

They say: "The App Store will perform this confirmation automatically." Oh, that's good. "However, developers may have separate obligations to independently confirm that their users are adults. To assist with this, the Declared Age Range API, available on iOS, iPadOS, and macOS, provides developers with a helpful signal about a user's age." Okay, so they're being helpful. "For apps rated 18+, Australia, Singapore, Brazil. However, for Utah and Louisiana" - oh, but not yet, wait for it.

"For users with new Apple accounts in Utah as of May 6th" - okay, so okay, wait, fresh accounts? How new do they have to be? Are they accounts created after May 6th? We don't know. "For users with new Apple accounts in Utah as of May 6th, 2026" - so a couple months from now - "and in Louisiana as of July 1st, 2026" - what a mess - "age categories will be shared with the developer's app when requested through the Declared Age Range API. The tools we previously announced have been expanded to help developers meet compliance obligations for Louisiana and Utah, including: Declared Age Range API, Significant Change API under PermissionKit" - that's that thing where if your app undergoes a significant change, you need to declare that because then that makes it potentially subject to all kinds of reevaluation. Then there's the "New age rating property type in StoreKit, and App Store Server Notifications."

They said: "New signals are now available through the Declared Age Range API, including whether age-related regulatory requirements apply to the user" - what a mess - "and if the user is required to share their age range. The API will also let you know if you need to get a parent or guardian's permission for significant app updates for a child.

"Developers can use the Declared Age Range API to present significant update notifications to adults in these states through the Significant Update Action, now in beta. When releasing a significant update, developers must follow the Human Interface Guidelines and provide users with a meaningful description of the update."

Leo, you know, on one hand, I would be - I'm a little tempted to feel some empathy and a little sorrow for Apple. At the same time I would say, guys, you brought this on yourself by refusing to do this five years ago.

**Leo:** Right.

**Steve:** They could have so easily put a far simpler system in place that would have satisfied people, that would have solved this problem, and prevented all of this ridiculous fragmentation. I mean, you're going to need a whole new building at Apple in order to, like, figure out what to do for whom on what day, depending upon whether they're, you know, oh, lord, what a mess.

**Leo:** This is what Meta wanted, by the way. They didn't want to do it, so they said, "Make Apple do this, please." By the way, there is in California...

**Steve:** Yes.

**Leo:** ...a law that goes in effect - you know about this.

**Steve:** Yes. Yes.

**Leo:** January 1st.

**Steve:** On January 1st of '27.

**Leo:** That will require operating systems, all operating systems to do this. And the Linux community is a little worried about it because nobody - the real issue is it's unenforceable. California can't make Linux do this. They can make Apple do it. They can make Google do it because they're gatekeepers. They can go after the companies.

**Steve:** And this is part of the larger plot; right? Like the 3D printer restriction is also unenforceable.

**Leo:** Right. Unenforceable.

**Steve:** It doesn't work.

**Leo:** Right.

**Steve:** You can write a law. It doesn't mean you can get what you want.

**Leo:** Yeah. Yeah.

**Steve:** But Leo, we're going to let our listeners get what they want.

**Leo:** What do they want? Do they want another commercial?

**Steve:** I want some coffee.

**Leo:** Whatever Steve wants, Steve gets. We'll be back with more Security Now!. I know you really want that. I should tell you, though, if you're in IT, if you're responsible for the security of your company, our advertisers here at Security Now! are always something you should be interested in. We have people, I think companies have realized, if you want to reach these IT decision-makers, you come to Steve because of who...

**Steve:** I'm so impressed by who our listeners are, Leo. When I hear from them, it's just like, I mean, I'm embarrassed that they're listening to me. It's like, what, me?

**Leo:** I know. Oh. I know, I have the same reaction constantly. Oh, you listen. Oh, it makes me a little nervous. We're going to meet a lot of our listeners in Florida, by the way. I'm very excited. Steve and I are headed to Zero Trust World. We'll tell you more about that in a little bit. And Steve's giving a presentation on Wednesday. Usually when we do these things, Steve, we've done them a couple of times before, there's a long line out the door to get a selfie with Steve Gibson. So we're going to have to...

**Steve:** They wanted Leo, too.

**Leo:** Oh, no, no. They wanted you. I usually jumped in just so they had me. In case they went home and said, oh, where's Leo? Oh, well.

**Steve:** Well, for what it's worth, I'm happy to, you know, smile into your phone, all of you listeners.

**Leo:** It'll be fun. It'll be fun. We'll line up a photographer.

**Steve:** If it doesn't break your camera, that's good.

**Leo:** Steve, now fully caffeinated, will continue.

**Steve:** As if I need more caffeine. And speaking of online web-based services, there has apparently been some concern, I would say justified, you know, if you want to follow the rules, over the intersection of child privacy enforcement and the apparent explicit need to violate that very privacy for the sake of complying with legislated age determination. Last Wednesday, on the heels of Apple's begrudging update to their age-related APIs and their app download enforcement, the U.S. Federal Trade Commission, our FTC, issued a formal policy statement with the headline: "FTC Issues COPPA Policy Statement to Incentivize the Use of Age Verification Technologies to Protect Children Online."

They wrote: "The Federal Trade Commission issued a policy statement today announcing that the Commission will not bring an enforcement action" - I don't know if I would call that incentivizing. It's like dethreatenizing - "will not bring an enforcement action under the Children's Online Privacy Protection Rule (COPPA) against website and online service

operators that collect, use, and disclose personal information for the sole purpose of determining a user's age via age verification technologies.

"The COPPA Rule requires operators of commercial websites or online services directed to children under 13, and operators with actual knowledge that they are collecting personal information from a child, to provide notice of their information practices to parents and to obtain verifiable parental consent before collecting, using, or disclosing personal information collected from a child under 13." And what a pain in the butt it is to actually do that; right?

So we see the problem here; right? The emerging age restriction regulations are placing the burden upon online services to whatever they must do to determine their visitors' ages. But doing this could force the site to run afoul of other regulations, specifically COPPA, which are already in place to protect the privacy of their underage visitors and users. In this instance, it's necessary to carve out an explicit privacy exception so that online services will be able to collect the data that they must without fear of tripping over COPPA's restrictions.

So the FTC explains: "Age verification technologies play a critical role in helping parents as they monitor their child's online activities. Since COPPA was enacted in 1998" - so it's been around for a while - "there's been an explosion in the use of Internet-connected technologies by children. To help parents navigate the challenges associated with their child's online activities, some states have started requiring some websites and online services to use age verification mechanisms to help determine the age of users. But as noted at the FTC's recent workshop on age verification technologies, some age verification technologies may require the collection of personal information from children, prompting questions about whether such activities could violate the COPPA Rule.

"Christopher Mufarrige, Director of the FTC's Bureau of Consumer Protection, said: 'Age verification technologies are some of the most child-protective technologies to emerge in decades. Our statement incentivizes operators to use these innovative tools'" - again, I would say, doesn't, you know, suspends disincentivizing them because it's the threat of action under COPPA that is causing them to say, wait a minute - "'which empowers parents to protect their children online.'

"The policy statement states that the Commission will not bring" - this is the statement from the FTC - "will not bring an enforcement action under COPPA Rule against operators of general audience sites and services and mixed audience sites and services that collect, use, or disclose personal information for the sole purpose of determining a user's age without first obtaining verifiable parental consent if they comply with certain conditions, specifically that they" - and we've got six bullet points.

"Do not use or disclose information collected for age verification purposes for any purpose except to determine a user's age; two, do not retain this information longer than necessary to fulfill the age verification purposes, and delete such information promptly thereafter; three, disclose information collected for age verification purposes only to those third parties the operator has taken reasonable steps" - and here again, I hate that kind of language, but okay - "to determine are capable of maintaining the confidentiality, security, and integrity of the information, including by obtaining certain written assurances from those third parties." Okay. So at least transferring responsibility, hopefully legally enforceable.

"Fourth, provide clear notice to parents and children of the information collected for age verification purposes; fifth, employ reasonable security safeguards for information collected for age verification purposes; and, finally, sixth, take reasonable steps to determine that any product, service, method, or third party utilized for age verification purposes is likely to provide reasonably accurate results as to the user's age." Again,

does that mean facial recognition, which we know is really prone to error, whatever. Finally, they say: "The policy statement indicates that the Commission intends to initiate a review of the COPPA rule to address age verification mechanisms. The policy statement will remain effective until the Commission publishes final rule amendments on this issue in the Federal Register, or until otherwise withdrawn."

Okay. So this policy statement is intended essentially to provide interim cover for online sites and services that do need to enforce privacy-breaching age-restriction measures today which would otherwise expose the site to COPPA infringement. This suggests that COPPA itself, as they said here toward the end of this FDC announcement, COPPA itself will require amending to provide a permanent and clear path for privacy-respecting age verification for minors. So again, one piece of legislation colliding with another. Surprise.

The Guardian reports that Meta's CSAM-detection AI is flooding law enforcement with low-quality unactionable - as we'll see here, it's like it's really sad - false positive reports of online child sexual abuse that are seriously hampering law enforcement's ability to function. Under The Guardian's headline "Meta's AI sending 'junk' tips to DoJ, U.S. child abuse investigators say," here's what The Guardian reported.

They said: "Officers from the U.S. Internet Crimes Against Children (ICAC) taskforce said that Meta's use of artificial intelligence to moderate its social media platforms is generating large volumes of useless reports about cases of child sexual abuse, which are draining resources and hindering investigations. Benjamin Zwiebel, a special agent with the ICAC taskforce in New Mexico, said last week during his testimony in the state's trial against Meta" - so this is New Mexico versus Meta. He said: "'We get a lot of tips from Meta that are just junk.' The state's attorney general alleges the company's platforms are putting profits over child safety."

Okay, now, at first I have say, I'll take a break here from this to say I was puzzled by that. But what I believe New Mexico's attorney general is saying is that rather than employing humans who would be able to usefully discriminate between what is and is not actual child exploitation and abuse, Meta is endeavoring, they allege, to save money by using AI, which is not actually doing the job. So Meta is failing in their obligation, but they're failing in a way that's causing lots of trouble.

The report continues, saying: "Meta disputes these allegations, citing changes it has introduced on its platforms, such as teen accounts with default protections. The ICAC taskforce is a nationwide network of law enforcement agencies coordinated with the U.S. Department of Justice to investigate and prosecute online child exploitation and abuse cases. Another ICAC officer, speaking on the condition of anonymity to discuss internal matters, said: 'Meta is providing thousands of tips each month. It's pretty overwhelming because we're getting so many reports, but the quality of the reports is really lacking in terms of our ability to take serious action.' The ICAC officer added that the total number of cybertips their department had received doubled from 2024 to 2025.

"Both Zwiebel and two ICAC officers said that unviable tips from Instagram, Facebook, and WhatsApp often contain information that's not criminal. The anonymous officers added that, in other cases, tips sometimes contain information indicating that a crime may have occurred, yet vital images, videos, or text are missing or redacted. The ICAC officer added: '[Unviable tips from] Instagram have really skyrocketed recently, especially in the last couple of months, and that's one of the biggest places where we're seeing important information not being provided. In those cases,' he said, 'we don't have the information to further the investigation. It weighs on you to know that this crime occurred, but we can't identify the perpetrator.'"

So just to clarify that point, you know, these investigators are saying that what they see are clearly crimes which Meta's use of AI happened to have found, so not a false positive,

it's true, but that the evidence that's needed to take any action about it is missing, which would not normally be the case if it were a human-driven investigation. So Meta's use of AI is not only flooding law enforcement with crap, but it's also serving to obscure the necessary details of actual crimes it detects. If we didn't know better, we'd be inclined to think this had been deliberately designed by criminals for criminals. It wasn't, I'm not suggesting that, but it's having that effect; right?

The story says: "Asked about Zwiebel's testimony and the ICAC officers' remarks, a Meta spokesperson said: 'We've supported law enforcement to prosecute criminals for years. The DoJ has repeatedly praised our fast cooperation that has helped lead to arrests, and NCMEC has praised our streamlined and improved tip reporting process. In 2024, we received over 9,000 emergency requests from U.S. authorities and resolved them within an average of 67 minutes, and even more quickly for cases involving child safety and suicide. Consistent with applicable law, we have reported apparent child sexual exploitation imagery to NCMEC and support them to prioritize reports, from helping build their case management tool to labeling cybertips so they know which are urgent.'"

Okay, so I'll just note that, while this sounds great, it doesn't appear to be responsive to the question of AI's use. That Meta spokesperson appears to be referring to the work of humans employed by Meta, not their cost-saving AI. The Guardian's reporting then shifts gears to provide some background on NCMEC, which is the National Center for Missing and Exploited Children.

The Guardian writes: "By law, social media companies based in the United States are required to report any detected child sexual abuse material (CSAM) on their platforms to the National Center for Missing and Exploited Children (NCMEC). It serves as a national clearinghouse for reports, which it forwards to the appropriate law enforcement agencies across the United States and internationally. NCMEC does not have the authority to filter out any tips that may be unviable before they're sent to the relevant law enforcement agencies." So 100% has to flow through.

"Meta is by far the largest reporter to NCMEC. In its data report for 2024, NCMEC said Meta made 13.8 million reports across Facebook, Instagram, and WhatsApp." Okay. So, you know, 13.8 million; right? Well, we have 12 months in a year. So simple math tells us that's over a million reports per month is coming from Facebook, Instagram, and WhatsApp. And that 13.8 million is out of a total of 20.5 million tips that NCMEC received in total. So well over half.

"NCMEC said that in 2024, more than one million CyberTipline reports were linkable to a specific U.S. state, and those reports were made available to the ICAC taskforces around the country, as well as other federal, state, and local law enforcement agencies, for investigation.

"Meta and other social media companies use AI to detect and report suspicious material on their sites and employ human moderators to review some of the flagged content before sending it to law enforcement. The Guardian has previously reported that tips generated by AI that have not also been reviewed by a social media company employee often cannot be opened by a law enforcement officer without a warrant because of Fourth Amendment protections. This extra step also slows investigations of potential crimes, lawyers involved in such cases have said.

"A Meta spokesperson said: 'It's unfortunate that court rulings have increased the burden on law enforcement by requiring search warrants to open identical copies of content we've already reviewed and reported. Our image-matching system finds copies of known child exploitation at scale that would be impossible to do manually, and we work to detect new child exploitation content through technology, reports from our community, and investigations by our specialist child safety teams.'

"Under the 'Report' Act, where REPORT is an acronym for Revising Existing Procedures On Reporting via Technology (REPORT), which came into force in November 2024, online service providers must broaden and strengthen their reporting obligations by notifying NCMC's CyberTipline, not only about child sexual abuse material, but also about planned or imminent abuse, child sex trafficking, and related exploitation. They must also preserve evidence for a longer period and face higher penalties if they knowingly fail to comply.

"Since the act passed, the number of unviable tips supplied by Meta has increased dramatically, which could be because the company is acting to ensure it is not falling foul of the law, two ICAC officers said. So in other words, Meta is complying because they're being forced to comply. The result, however, is a lot more noise among the signal. They said many of these tips could not be construed as a crime, such as adolescent girls talking about which celebrity they find most attractive.

"Special agent Benjamin Zwiebel said in court: 'Based on my training and experience, it appears that they are being submitted through the use of AI, as these are common mistakes that an AI would make that a human observer would not. Zwiebel added that his department receives significantly fewer tips on legitimate cases of child sexual abuse material distribution from Meta than in previous years.'" So in other words, not only has the noise gone up, but the signal, the quality, has gone down.

"Every tip that reaches an ICAC division must be reviewed, and the influx of unviable tips is taking time and resources away from investigating legitimate cases of child abuse, said two officers. One ICAC officer said: 'It's killing morale. We're drowning in tips, and we want to get out there and do this work. We don't have the personnel to sustain that. There's no way that we can keep up with the flood that's now coming in.'"

So I want to chalk this up less to Meta being evil, which I don't think is the case, than to the growing pains of effective AI deployment. We're still very much learning how to best use the new and surprising capabilities of large language model networks. And I suspect that a strong case could be made for there truly being far too much content for humans to manually inspect. In other words, we've talked about this, right, with the legislation that the UK keeps circulating and trying to make happen? Where it's just, you know, how are we going to do this? Apple has proposed doing on-device CSAM image comparison, and nobody wanted that. I mean, the actual volume of content is beyond human management.

So, you know, although the specter of having overlord AIs examining everything that's transacted over social media feels very Orwellian, our legislators are requiring a level of oversight from social media companies that likely has no other workable solution. AI it will be. We just need to continue figuring out how to best use it. And again, all evidence is we're making headway. We're going to get a lot better than we are. We can clearly see how much better we are now in using AI for code than we were a couple years ago. This is going to get better. And I think we're just going to, in the future, the legislators are going to force it to be the case that some machine intelligence is going to be watching dialogs, and we're just, you know, users are going to have to put up with that as a cost of the privilege of being able to communicate with encryption.

I just saw a short mention blurb that surprised me. The news was just that Russia's wonderfully named Internet watchdog, Roskomnadzor, has now blocked Russian citizens' access to - you're not going to believe how many - 469, Leo, individual VPN services inside Russia.

**Leo:** Of course. All of them, in other words. All of them they could find.

**Steve:** Yes. Yes. I mean, which means, but none of the ones that have sprung up since then; right?

**Leo:** Right, right.

**Steve:** It seems to me that the fact that there are 469 VPN services inside, you know, discrete individual VPN services inside Russia to be blocked in the first place, that's the real story here. You know; right? Talk about a citizenry that's desperate to escape the shackles of their own state's filtering and tampering and management. This is a citizenry that is desperate for contact with the outside world, and a repressive government that's doing everything it can to prevent that. It's becoming increasingly clear why Russia has been experimenting with completely disconnecting from the global Internet. They want the ability to just go internal sovereign and cut off all outside contact.

In other Russian news, I saw a report that indicated that the Kremlin had decided to fully block Telegram starting in April of this year. Right? Okay, next month. That puzzled me, since I thought that Telegram was already being fully blocked. We talked about that just recently. But this reporting stated that Telegram was currently only 55% blocked. Okay. It's not clear to me what a 55% block might mean. The only thing I can figure is that perhaps access to Telegram is currently being limited to specific regions or sectors or industries, and that additional regions are being added to the master block list so that by the end of this month of March, nothing will be left. Okay. Whatever the case, Russia appears to quite intent upon controlling its citizens' access to information. Good luck. Information...

**Leo:** It's inevitable. If you want to do that, you've got to get rid of VPNs. That's the next step.

**Steve:** Yes. It is. Yes. And as we know, information wants to be free.

**Leo:** Yeah. It's pretty hard.

**Steve:** As has been said, it's very difficult. I mean, you know, we've got satellite now, too. Okay. This one. Oh. About 14 months ago, in January '25, we reported that the UK was launching a plan to begin continuously and proactively scanning its own national public-facing network segments for the purpose of preemptively detecting vulnerabilities and alerting those owners of the IP addresses where vulnerabilities were found. Our listeners may also recall that I was jumping up and down over how much I thought this made sense, and suggesting that this was something every nation should be doing to its own public-facing Internet address ranges in its own self-interest. I think this is just a great idea.

So we're talking about this again today, 14 months later, because last Thursday the UK, out of a celebratory press release, used the headline "Government cuts cyber-attack fix times by 84% and launches new profession to protect public services." A new profession. Huh. Okay. The press release led with three summary bullet points. They said: "Critical cyber weaknesses across the public sector will now be fixed six times faster than before. Ministers are determined to go further, with first-ever dedicated government Cyber Profession" - that's in caps, capital C, capital P, Cyber Profession - "to give the state the skilled staff it needs to protect UK's key services from cyber threats. And finally, the number of serious unresolved cyber security weaknesses across government cut by three

quarters as part of government-wide efforts to strengthen Britain's digital defenses." Wow. Sounds great.

Before I share what the press office of the UK said, allow me to preface this by noting that we're going to encounter something that makes no sense whatsoever to me. But regardless, here's what they wrote. They first said: "Public services millions of people depend on, from the NHS to the Legal Aid Agency, are becoming significantly safer and more resilient thanks to major improvements by the government to identify and fix cyber threats." Great.

"A specialist government monitoring service, introduced as part of the Blueprint for Modern Digital Government, published in January '25, means serious security weaknesses in public sector websites are fixed six times faster, cutting the average time from nearly two months to just over one week." Okay. So far so good. But then this appears to go off the rails.

The release next says: "The vulnerabilities are in the Domain Name System (DNS), the Internet's address book that turns website names into the numbers computers use to find them. Weaknesses in DNS can allow attackers to redirect users to fraudulent sites, steal sensitive data, or take services offline entirely, with potentially serious consequences for anyone relying on government services." Okay.

The press release says: "Before this service was in place, a weakness in a government DNS record could go unnoticed for nearly two months, long enough for a hostile actor to redirect someone trying to access a government service to a fake site designed to steal their personal details, intercept sensitive communications, or disrupt services that people rely on. The vulnerability monitoring service has closed this window down to eight days. It alerts the right people with clear, practical guidance on how to fix the problem, and tracks progress until each issue is resolved."

Okay. What the hell are they talking about? What is a "weakness in a government DNS record" What? In this day and age, when I see something that sounds entirely plausible and reasonable to a lay person, but which is actual nonsense, the first thing I think is that some AI somewhere was having a bad day. The press release said: "Before this service was in place, a weakness in a government DNS record could go unnoticed for nearly two months." Again. What? What is a weakness in a, like, it makes no sense at all. There's no such thing. Okay. So let's just play along and see what else happens.

The release continues: "Speaking at the annual Government Cyber Security and Digital Resilience conference, Digital Government Minister Ian Murray will outline how this will sharply reduce," right, the reduction in weak government DNS records, apparently. What? Will sharply reduce something. Oh, the risk of hackers targeting essential services like the NHS. Well, that's good. It's got a weak DNS record. You don't want that. So by all means, reducing its effect somehow from almost two months of weakness down to just eight days, that's a big improvement. No one would argue.

"He'll also outline how the government has reduced its backlog of these weak DNS vulnerability records, okay, by 75%, significantly shrinking the window for cybercriminals to target essential government services due to weak DNS records, okay, from GP surgeries and ambulance trusts to hospitals and social care providers. Today's announcement marks a decisive step in closing the door on such threats, whatever they are, with the government going even further with the launch of the first-ever dedicated government Cyber Profession. Apparently we're going to have a cyber profession, capital C, capital P, that focuses on the weakness, I don't know of what. DNS? What are they going to do, DNS monitoring? What are they - okay.

So the press release says: "This program will recruit and train the top-tier cyber experts needed to keep public services safe." Oh, good. "Minister for Digital Government Ian Murray said: 'Cyber-attacks aren't abstract ideas.'" Oh, no. We know that. "They delay NHS appointments, disrupt essential services" - almost put Jaguar out of business. That's me, not him - "and put people's most sensitive data at risk. When public services struggle, it's families, patients, and frontline workers that feel it.

"The vulnerability monitoring service has transformed how quickly we can spot and fix weaknesses before they're exploited so we can protect against that. We've cut cyber-attack fix times by 84% and reduced the backlog of critical issues by three quarters. And as the service expands to cover more types of cyber threats" - what, beyond weak DNS records, whatever those are? - "fix times are falling there, too. But technologies alone aren't enough. Today," he says, "I'm launching a new government Cyber Profession" - capital C, capital P - "to attract and develop the talented people we need to stay ahead of increasingly sophisticated threats - making government a destination of choice" - that's right, baby, government is a destination of choice for cyber professionals worldwide - "who want to protect the services that matter most to people's lives."

Dr Richard Horne, CEO of NCSC, said: "Cybersecurity is more consequential than ever today with attacks" - it does sound like maybe some good AI wrote this part - "ever today..."

**Leo:** Are there bullet points?

**Steve:** "...with attacks in the headlines showing the profound impacts they can have on people's everyday lives and livelihoods. As our public services continue to innovate, it's vital that they remain resilient to evolving threats and," okay, blah blah blah blah. So they finally said: "The VMS" - that's this new system that's been online for 14 months - "continuously scans 6,000 UK public sector bodies, detecting around 1,000 different types of cyber vulnerabilities. When a weakness is identified, the service alerts the relevant organization with specific, actionable guidance, and tracks progress until the issue is resolved."

Okay. Now, THAT finally makes sense. THAT is what we would expect. They have a continuously running Internet scanner that's scanning 6,000 UK public sector agencies and entities looking for 1,000 different types of cyber vulnerabilities at each of the IPs of the configured targets. Yay. Unfortunately, the presence of that looney tunes nonsense about weakness in government DNS records casts the entire announcement into question. Just where does the AI brain fart that apparently occurred end in this announcement and reality begin? If that's in there, it's hard to know what else is just fuzzy. But we do now appear to be, you know, back on track.

The release finishes up, writing: "By automating and detecting and streamlining remediation, the service has [bullet point] reduced median time to fix domain-related vulnerabilities from 50 days to eight days, an 84% improvement." Okay, now, we're back to crazy town there. What is a 'domain-related vulnerability,' and how can it have been reduced from taking 50 days to fix down to just eight days? How can it take any days? You know, it really does seem as though an AI had a hand in the preparation of this release. Which is too bad. The other three bullet points seem more reasonable. They are: "Reduced median time to fix other cyber vulnerabilities from 53 days down to 32." Okay, not great, but better. "Cut the backlog of critical open domain-related vulnerabilities" - whatever that is again - "by 75%, processed and resolved around 400 confirmed vulnerabilities each month."

So the press release finishes, saying: "The new government Cyber Profession is co-branded with the Department for Science, Innovation, and Technology and the National Cyber Security Centre. It will introduce a competitive total employee offer, establish a dedicated Cyber Resourcing Hub to streamline recruitment, and create a clear career framework aligned with UK Cyber Security Council professional standards.

"It will also include a government Cyber Academy for training and deployment, a new apprenticeship scheme to build future talent, and structured career pathways to strengthen long-term capability across the public sector. The North West will serve as a primary hub for the profession, building on Manchester's growing digital ecosystem and the forthcoming government Digital Campus."

So all that sounds great and reasonable, too. The UK has clearly implemented, although they seem unable to describe what it is they have, an extremely useful service. And I do seriously hope that other nations pick up on this idea and put it into practice, the idea of a country scanning its own Internet infrastructure preemptively for known problems. I mean, this is what CISA should be doing. And then finding out who owns those IPs and letting them know they've got problems there. That's a win-win. I don't know what a soft government DNS record is. Wow. And I don't think anybody else does either because, you know, we would know what that was. Right? We understand this stuff, and, like, what? What are you talking about? Really, it's just a mystery.

Leo, let's take a break.

**Leo:** Okay.

**Steve:** For a sponsor who's not a mystery.

**Leo:** No mystery to you or me because we're about to head to Orlando for Zero Trust World, ThreatLocker's big security conference. Steve's going to give a presentation Wednesday, last event of the day. So it's right before the cocktail party. In fact, it might even overlap a little bit. But it'd be worth sticking around. And Steve and I will stick around afterwards to talk to you.

**Steve:** And you're going to be in costume; right?

**Leo:** Not for this.

**Steve:** Oh.

**Leo:** Thursday, they're very famous for every year ThreatLocker has a costume party. And I think the theme this year is Sixties Space.

**Steve:** Oh, thank goodness. I thought I was going to be the only person not in costume. But for the Wednesday evening cocktail party, no costumes.

**Leo:** No costumes. Just be normal.

**Steve:** Okay. Okay.

**Leo:** Which is black; right? You're going to wear black of some kind.

**Steve:** I'll be wearing black.

**Leo:** Me, too.

**Steve:** Even though Orlando is hot, and black absorbs heat. Just like it does for the crows, Leo.

**Leo:** Oh, yes. They absorb the energy focused upon them, whether by the sun or some other third party. On we go with the show, Mr. Gibson.

**Steve:** Okay. So I mentioned this at the top as somebody I knew, Leo, you would remember. I was just scanning the news, and I encountered a piece of news declaring that "Vastaamo hacker disappears." And I thought, okay. I have no idea what that is. But then reading a bit into the story it mentions that a Finnish hacker lost his appeal and will have to go back to prison after a court increased his original sentence. Okay.

So again, like, okay, nothing stands out there. But we'll recall this event from six years ago. The report explains that this Finnish hacker was sentenced to six years and three months for hacking the Vastaamo Psychotherapy Centre in 2020 and then extorting its patients, which is what made that stand out, both as I was reading this and of course I remembered we talked about this at the time. This creep obtained the psychotherapy center's very personal and highly confidential medical psychotherapy records.

**Leo:** That's just awful.

**Steve:** Including, of course, the contact information that would be needed for them to be contacted for the sake of extorting them. Which he then did. He threatened them with public exposure of their mental illnesses unless they paid up.

So beyond this, as I also recall, Leo, you and I were shocked when we saw the sheer number of patient records that that Vastaamo Psychotherapy Centre had maintained online which were stolen. That was the other part of the scandal. We noted that not only were they at fault for not better protecting their data, but they should not, you know, they should not have had that much old patient data around. They should be held accountable for leaving the data of years and years of previous patients in hot storage, online and readily accessible. You know, I understand they might have felt they needed to retain records for some possible future need, but those could be archived offline for retrieval on demand, not sitting on the same server with all of the current records, all of which this hacker sucked up.

**Leo:** I agree 100%, yeah.

**Steve:** So anyway, just a weird little aside. I mean, I'm like, remember that guy? Yeah, we talked about him. Funny how we seem to catch all the important bits. I'm happy about that.

So in their Cyber Intel Brief, the cyber intelligence firm Dataminr - they left the "e" out, so it's D-A-T-A-M-I-N-R, Dataminr - reports that the Scattered Lapsus\$ Hunters, which we're now abbreviating SLH, although I don't know if anyone's going to remember what SLH is, so I'm going to keep saying Scattered Lapsus\$ Hunters because it's fun, that they've begun recruiting female individuals for their voice phishing campaign. SLH is offering upfront payments - big ones - for social engineering calls targeting IT help desks. Dataminr's report offered three key takeaways.

They said under Tactical Evolution: "SLH is diversifying its social engineering pool by specifically recruiting women to conduct voice phishing attacks, likely to increase the success rate of help desk impersonation." Under Large Incentives they said: "The group is offering significant financial incentives, between \$500 and \$1,000 upfront per call" - which stuns me - "and providing pre-written scripts to their recruits." And High-Profile Risk. They said: "SLH is a 'supergroup' alliance of Lapsus\$, Scattered Spider, and ShinyHunters, known for compromising major global corporations and stealing over 1.5 billion records." So far and counting.

The Dataminr posting then walks us through their discovery of SLH's online recruitment postings and ends with some useful advice to any potential enterprise targets. Under their heading "Organizations should adopt a heightened defensive posture against social engineering," they enumerate four points. "First, Help Desk Training: Immediately brief IT help desk and support personnel on this specific recruitment trend. Emphasize that attackers may be using pre-written scripts and polished voice impersonation." And the fact that it's a girl on the phone doesn't mean it's not your typical hacker attacker guy. So don't be fooled by that.

"Strict Identity Verification: Enforce out-of-band identity verification, you know, a video call or a secondary internal verification or some sort." You know, it's like when you receive email that says "Phone this number if you'd like more information," pretending to be your bank. You need to go look up your bank's phone number yourself rather than using the number that came in the email, that kind of thing.

"Harden MFA Policies." They said: "Move away from SMS or push-based MFA (multifactor authentication), both of which are vulnerable to SLH's known TTPs like SIM swapping and fatigue bombing. Implement FIDO2-compliant hardware security keys wherever possible." And finally, "Monitor Anomalous Access: Audit logs for new user creation or administrative privilege escalation immediately following all help desk interactions." Meaning, you know, check your logs after a help desk interaction to see whether there might be anything going on that the bad guys immediately launched into following that interaction. So the point being you really do need to be proactive. And it's interesting...

**Leo:** I remember a phishing attack some years ago where a woman called a customer service rep. Remember that customer service is the first two words in their title. They want to help customers. So the way this phishing attack, this social engineering attack worked, the woman was frantic, saying my husband left his phone at home, and he's on a business trip, and he's going to desperately need it, and I need to reach him. And they played a baby crying in the background on a recorder. I mean, it was this whole scenario.

**Steve:** So you would really get sucked in and believe.

---

**Leo:** Yeah. And of course you can't do that with a guy. So yes, a woman's voice is going to in some cases really be more effective because I think you're right, people don't expect a woman to be social engineering them.

**Steve:** Yeah. So again, it just knocks your guard down a notch.

**Leo:** Yeah. Yeah. And of course it was a simjacking attempt. All they wanted to do is get the phone number transferred so that they could get those SMS, you know, text messages.

**Steve:** Yup. And then goodnight.

**Leo:** Goodnight, yup.

**Steve:** So last Wednesday, Cisco released the news of CVE-2026-20127, once again achieving that rarest of rare CVSS 10.0 scores.

**Leo:** Yikes. Good ol' Cisco. You know what, they always come in strong.

**Steve:** Used to be, "Oh, Newman." Now it's "Oh, Cisco." This one is an actively exploited zero-day, first discovered while it was being abused in the wild. The title Cisco gave their disclosure was "Cisco Catalyst SD-WAN Controller Authentication Bypass Vulnerability." Yep, you heard it right. Surprise, surprise, an authentication bypass vulnerability.

Cisco wrote: "A vulnerability in the peering authentication in Cisco Catalyst SD-WAN Controller, formerly SD-WAN vSmart, and Cisco Catalyst SD-WAN Manager, formerly SD-WAN vManage, could allow" - right, could - "could allow an unauthenticated remote attacker to bypass authentication" - that pesky authentication - "and obtain administrative privileges on an affected system." They said: "This vulnerability exists because the peering authentication mechanism in an affected system is not working properly."

Huh. "Not working properly." Okay, no one would disagree with that, although calling it "catastrophically defective" might be more accurate. Okay, this one is so bad that both the U.S. NSA and CISA here in the U.S., the Australian Signals Directorate's Australian Cyber Security Centre, the Canadian Centre for Cyber Security, New Zealand's National Cyber Security Centre, and the UK's National Cyber Security Centre all published "patch or perish" announcements in a desperate attempt to bring the need to patch all systems to the attention of their owners. The "SD" in SD-WAN stands for Software Defined. So it's a software-based networking platform that connects branch offices, data centers, and cloud environments together through a centrally managed system. It uses a controller to "securely" route traffic - securely in quotes, of course, air quotes, between sites over encrypted connections.

This is another instance where any company that recognized that simple authentication can never be relied upon for security, and had therefore taken the trouble to preemptively separately restrict, for example, incoming SD-WAN connections to only known endpoint peers, well, they'd never have anything to worry about. They wouldn't have anything to fear from these authentication failures and, (a) would not have suffered

a potentially devastating network compromise; and, (b) could therefore update their SD-WAN instances with something less than pants-on-fire emergency at their leisure.

Once again, Cisco's own announcement moderately underplayed the consequences. They wrote: "An attacker could exploit this vulnerability by sending crafted requests to an affected system." Of course all systems are affected. "A successful exploit could allow the attacker to log in to an affected Cisco Catalyst SD-WAN Controller as an internal, high-privileged, non-root user account. Using this account, the attacker could access NETCONF (net configuration), which would then allow the attacker to manipulate network configuration for the SD-WAN fabric."

It was the Australian Signals Directorate, I'll note, who first discovered and reported these attacks being used in the wild. Not surprisingly, they paint a somewhat less rosy picture of the consequences, writing - this is Australia - "Malicious cyber threat actors are targeting SD-WANs of organizations, globally. These actors exploited a Cisco Catalyst SD-WAN controller authentication bypass vulnerability, CVE-2026-20127. After exploitation of this vulnerability, the malicious actors add a rogue peer, and eventually gain root access to establish long-term persistence in SD-WANs." So, sorry, Cisco. Not just non-root user accounts.

**Leo:** I like that new term "rogue peer." I'm going to keep that around.

**Steve:** Rogue peer.

**Leo:** Yeah.

**Steve:** Yeah. Again, it's one of our main themes here. You cannot rely upon authentication. And more importantly, you don't need to. You can apply additional factors of authentication, not allow somebody to get to a port. You know, you have a bunch of offices scattered around, what, the world. They have their own networks. You know what their IPs are, even their IP blocks. Probably the specific IP of your peer SD-WAN. Why not take the time to put a rule in the firewall so that you only accept incoming traffic from that IP to your SD-WAN. Why not?

**Leo:** Can you spoof an incoming IP?

**Steve:** No.

**Leo:** How interesting.

**Steve:** No, because it requires a connection.

**Leo:** It's a conversation, yeah, right.

**Steve:** Yes. Yes. And so it was like, all anybody has to do is not assume that, I mean, first of all, I was about to say not assume that Cisco is perfect. Who would do that?

**Leo:** Well, that's a good thing to not assume, yes.

**Steve:** Please.

**Leo:** Safe bet.

**Steve:** So, you know, protect yourself. Put firewall rules in because you're talking to fixed endpoint IPs. Only allow the conversations from them. Why would you ever want China or Russia or North Korea to connect to your SD-LAN? You don't.

**Leo:** Right. I mean, I do that with my freaking Synology. It's not - how hard could it be?

**Steve:** Exactly. Exactly. I mean, you know, yes.

**Leo:** Yeah.

**Steve:** Okay. So VulnCheck's annual report on the in-the-wild use of known security CVEs like, you know, CVEs about security breaches, is interesting. All I have, I have the entire 41-page report. It will probably be next week's topic because it looked like in a quick glance through it there was just so much juicy stuff there. But the teaser summary, which is all you get until you give it your name and email address so they can market to you for the rest of your life, it was interesting, too.

They said: "In 2025, barely 1%" - this is what was interesting. "Barely 1% of disclosed vulnerabilities were exploited in the wild." Which might not be what we think. It means that the distribution of exploits is not uniform. It is very peaky. Of course it's the juicy exploits which were exploited; right? They said: "Yet those that were exploited were operationalized quickly, attracted diverse threat actors, and often caused outsized damage before organizations had a chance to respond." Just like this SD-WAN nightmare. They said: "This report identifies which vulnerabilities mattered, why attackers targeted them, and where timing failures left organizations exposed." Like I said, that's going to be fun to talk about, to look at this analysis.

They said: "VulnCheck tracked exploitation patterns, threat actor behavior, and weaponization timelines across hundreds of thousands of vulnerabilities in 2025. The data revealed how quickly new vulnerabilities became bona fide threats, how AI proof-of-concept code is polluting risk assessment pipelines" - interesting - "and which threat actors ramped up vulnerability exploitation amid geopolitical tension."

Then we have three bullet points. "VulnCheck identified 50 routinely targeted vulnerabilities from 2025 that had elevated risk profiles by the end of the year, drawing interest from ransomware, threat actors, botnets, and researchers, often all at once. Second, proof-of-concept exploits for new CVEs increased 16.5% in 2025, inundating organizations with 'risk' signals that often turned out to be false or misleading AI-generated slop." Again, AI slop is a term which has taken hold. "And finally, China-nexus threat actor attributions increased 52% year over year, while ransomware groups shifted toward zero-day exploitation at accelerating rates, with 56.4% of ransomware CVEs discovered through zero-day activity." So the landscape is changing. These guys have

analyzed everything that happened in 2025, produced a 41-page report full of information. And I suspect that's how we're going to start next week.

We're going to start our listener feedback, Leo.

**Leo:** Yeah.

**Steve:** And I think we should take one more break because then we'll have one before our final coverage.

**Leo:** Okay. That sounds good to me. You're watching Security Now! Special Edition, in a sense, because we are doing this on a Sunday. Steve and I, as I mentioned, are going to Florida tomorrow, and we'll be gone all week. I've got people covering the shows for me. We will put this show out in the normal Tuesday timeslot. And if you're a Security Now! fan, good news because this week we'll have two Security Nows, a second show which will be the presentation Steve's giving at Zero Trust World. What's it called?

**Steve:** "The Call Is Coming From Inside the House."

**Leo:** I'll leave it to you to speculate as to what...

**Steve:** I was just thinking, you have people covering your shows, and we've got people covering our squirrels' need to continue being fed while we're gone.

**Leo:** Yes. I have Mikah. You have a squirrel sitter. It makes sense.

**Steve:** We have house sitters that are going to keep the squirrels fed because Lorrie said, "What about the squirrels?" It's like, okay.

**Leo:** Anyway, we're glad, if you're watching live, we're so glad that you figured that out.

**Steve:** A listener, David Benedict, said: "Hi, Steve. Not to pull you back" - but he's going to - "into the whole Code Signing discussion again" - it's a lot of interest, a lot of our listeners have expressed a strong interest. He said: "But what if WE simply don't buy those code signing certs? What if we simply start self-signing code? Is there anything to stop us from self-signing and building our own reputation that way? Thanks, Dave Benedict."

So, okay. That's an interesting idea. The moves that the CA/Browser Forum have been making on the code signing front feel entirely different from their earlier squeezing on the TLS certificate side. The reason Let's Encrypt was able to effectively replace and displace the traditional certificate authorities for the world's web server domain validation certificates is that Let's Encrypt is only providing what its name suggests: encryption. Let's Encrypt. It's making no assertion of any kind about the reputation of the domain name holder.

And when you think about it, where strong assertions of identity are needed and are being made about the owner of a certificate, whether for a web domain, maybe the digital signer of a document - that matters, too - or the authorship of code, we do need entities such as certificate authorities standing by to do the necessary work of verifying identity and carefully issuing certificates which attest to what their research has found.

Unfortunately, while we've been going about our lives, the certificate authority business has been quietly consolidating. This has sometimes been triggered, as we've covered on the podcast, when an irresponsible certificate authority so flagrantly abuses its position of trust that the various root programs are finally forced to revoke their trust. In those cases, the disgraced CA is forced to sell off its Certificate Authority business assets to another certificate authority. In other cases, it's just a bigger fish swallowing up a smaller fish, reducing the competition. While I was scouting around for a new code signing certificate authority, I noticed that many of the smaller-looking companies had exactly the same pricing as DigiCert. It turned out that many of them have simply become fronts for DigiCert's products. They're just resellers.

The upshot of many years of CA industry consolidation is that the world no longer has a - this is sad, but true. The world no longer has a competitive certificate authority industry. We are watching the formation of a monopoly that has the gall to charge its customers per signature. We can see the writing on the wall. There are already plans like that happening. It's where we're headed.

Dave began his note, writing: "Hi, Steve. Not to pull you back into the whole code signing discussion again." It's not your fault, Dave. This whole thing obviously rubs me the wrong way. One of my personality hot buttons happens to be bullying. I have never been okay with the abuse of power which is what, I believe, anyone observing the actions of the CA industry would conclude is happening. I don't see any way out of this, but I will gladly share any solutions I find.

To that end, during this research I discovered that all of the various CAs (certificate authorities) who offer code signing certificates - remember that now any code signing certificate must be in hardware, you no longer get a software certificate - all of the code signing offering CAs provide the option of installing certificates into a customer-provided HSM (hardware security module) rather than selling the certificate pre-installed in their own dongle token. You know, typically they charge another \$100 for that. But that's it. That's all it can do. Period.

The reason I'm mentioning it is I found a gorgeous \$72 form factor USB-A HSM dongle that I love. It's called the "SmartCard-HSM 4K," 4K because it can handle 4096-bit RSA keys, which is now what's necessary. It also does elliptic curve keys, which can be much smaller. I have a link to this device in the show notes to one particular retailer of this device. It's got its own website at [smartcard-hsm.com](http://smartcard-hsm.com). And most significantly, all of the card's multi-platform support is open source. So this is a fully open source \$72 beautiful little hardware security module. I've got a link to its GitHub page in the show notes.

One of the very cool features of this for me is that HSM, you know, having a hardware security module enables secure and encrypted cross-HSM private key and certificate transfers. In other words, I have multiple machines where I want to be able to sign code. I've got two working locations and GRC's servers in the Level 3 data center. So I first had the first HSM securely generate a 4096-bit RSA key pair. The private key never leaves the device, which is what the certificate authorities require. But the public key is exported in a CSR, a Certificate Signing Request. I uploaded that CSR to IdenTrust for it to receive their signature. They promptly returned the resulting certificate, which is then used to verify any signatures that the HSM generates for my code, since it'll be doing the code signing.

One of the many cool things about this solution is that each of these HSMs includes its own permanent device certificates that enable it to establish a secure key sharing key among others of its kind. This allows one HSM's private keys to be securely duplicated across many other devices, as many as you may wish, as well as being externally backed up for export without ever being able to expose its private key. So it meets all the requirements for us for security, yet gives us as HSM users and code signers way more flexibility.

Each HSM also has a large amount of storage with room for hundreds of keys and certificates and whatever you want to put in there. PGP, GPG, all of that stuff is supported. All of the platforms are supported, and everything is open source on GitHub. So, for example, if an enterprise might have a number of trusted developers, work-from-home, satellite offices or whatever, for the price of \$72 each, as many developers can be given the ability to securely sign code as needed.

Anyway, there's much more than what I've shared. So if you have an interest or need, check out [smartcard-hsm.com](http://smartcard-hsm.com). The retailer I used, [CardLogix.com](http://CardLogix.com), I've got a link in the show notes, they're the retailer which I found happened to be near me in the U.S. The "Where to Buy" page at [smartcard-hsm.com](http://smartcard-hsm.com) also lists a German and a Taiwanese reseller. So if you're over in Europe you can find someone near you, or in Taiwan.

**Leo:** I have a Nitro, I have a bunch of Nitrokeys. It works on that, too, which I didn't realize.

**Steve:** Yes, yes. Nitrokey is also supported by all of this software, yeah. Since my original DigiCert EV code signing certificate does not expire until November 20th, as it happens, of this year, but I wanted to, remember that I wanted to obtain a three-year certificate before they were no longer available, my plan has been to see whether I can pre-establish a reputation for the new, now three-year duration IdenTrust certificate by having it cosign GRC's freeware. That's now in place. GRC's most popular freeware, like for example ValiDrive, which is now being downloaded 1,000 times a day, is now cosigned, both with DigiCert's original certificate and the new IdenTrust certificate.

So I'm hoping that, once we get to November, I'll be able to drop the DigiCert signatures, because my DigiCert certificate, code signing certificate, will have expired, and that my newer IdenTrust certificate, which will by then have 10 months of exposure to Windows Defender and smart whatever the hell, you know, Microsoft will have seen this and gotten used to it. And I'm hopeful that it will be able to stand alone and keep Windows' trigger-happy gatekeepers happy.

Okay. So, and then finally, just to see whether I could because I had so much fun playing with this new technology, last week, as I mentioned talking to DigiCert, I also reissued my existing DigiCert certificate in, instead of in they provide me with a dongle, which they did the first time, 2.5 years ago, I did it in this customer-provided HSM mode. That allowed me to add DigiCert's certificate into my new HSMs, alongside the newly minted IdenTrust certificate. It all worked perfectly. Now I have HSMs containing both the existing, expiring in November DigiCert code signing certificate and the new IdenTrust code signing cert, which goes for three years.

So, okay. Believe it or not, I haven't forgotten about David. He started me off on all this by asking about the possibility of coders sidestepping all this nonsense by using self-signed certificates. The use of self-signed certificates has been common practice for web developers for many years. I have a self-signed certificate for "localhost" sitting in the trusted root stores of my various workstations. I run a web server on those machines which uses that certificate, and I use it for local web development. Having a self-signed

certificate for "localhost" allows me to use HTTPS URLs starting with `https://localhost/` and then whatever, without the web browsers that I'm using pitching a fit. So it's just very handy.

Okay. So let's explore how this might be extended for code signing. If, rather than having DigiCert or IdenTrust or whomever sign my code signing request, if I could instead use my private key to sign the certificate, creating a self-signed certificate which would then be installed into the system's trusted root store, how would that work? Well, from that point on, any code I signed would carry that root certificate's matching public key, and any check on the validity of my code's signature on this machine would show its signature to be valid.

But the reason this is not a useful solution for a software publisher, unfortunately, such as GRC, is that these code signatures would only be valid on machines that had previously installed its matching root certificate. What DigiCert, IdenTrust, and all the other CAs have going for them is that their root certificates are already pre-installed wherever any certificates they have signed might need to be trusted. Since Windows has now developed the practice of deleting on sight any executable that appears on its drive without a valid and trusted signature, especially one downloaded from the Internet, and that's probably why people are able to compile their own code is it came from them, although I've compiled my own code, too, and Windows has immediately stomped on it.

It would be necessary for GRC, if I was using a self-signed cert, it would be necessary for GRC to tell its customers that before attempting to download, let alone run, any of our software, they must first download and install GRC's own trusted root certificate. Well, since I would never install anyone else's root certificate into my machine's root store, I would never ask anyone else to do that for us in order to download and run my code. The burden of making my code acceptable to someone else's machine should be on me, not on them. So while signing one's own code for use on our own machines would work, just like using a self-signed web certificate for local use of a web browser and web server, I don't see any way to break our Certificate Authorities' stranglehold on developers for code signing that needs to be universally trusted.

And as I said before, I get the need for certificate authority. Just for encrypting web domains, we don't need them. That's why Let's Encrypt is a viable alternative, and that's why it works, because all we're saying is encrypt this traffic to wherever I'm going. And I'm not sure where I'm going. I'm going to this domain name. Certificate authorities, we need a third party like a CA when we need to have the ability to digitally sign a document and have that signature mean something because we proved who we were to them, or to sign our code. Or if we want an organization validation certificate for TLS web servers, not just a domain validation certificate. So I'm not saying that certificate authorities don't have a place, and we don't need them. I've got a problem with that abuse.

Now, there is one place where self-signing could make sense because everything I said about individual developers like me, that does not apply to enterprises; right? Enterprises might choose to use - they could buy a cert from DigiCert. They could use a publicly trusted code signing certificate for their internal use. But within an enterprise it might also work to sign code with a certificate that is only trusted within the enterprise's own enterprise machines. Remember that many enterprises already install their own TLS web root certificates on all internal workstations so that their networking middleboxes are able to intercept, decrypt, and scan everything that enters and exits their network.

You can't get on the enterprise LAN and get out to the outside world unless you have one of, you know, their own TLS cert in your enterprise workstation machine. So I could see that it would make sense to add a private code signing root certificate to all enterprise machines for their own internal use. On the other hand, if you're an enterprise, you may

not care that much about what the various CAs have now chosen to charge for the privilege of signing code. Although it does appear to be escalating over time.

DGC wrote: "Hey, Steve. Long-time listener, but I'm a few episodes behind right now. In Episode 1062 you said: 'We also see employees in positions of trust on internal enterprise networks being tricked into clicking malicious links and inviting malware inside the house. No form of fancy coding AI is going to fix any of those things.'" Then he writes: "That may not be entirely true. I recently came across a new solution, 'Charlemagne,' which runs on a desktop and monitors (privately, locally) what the user is doing. It uses an SLM, a Small Language Model, to detect potentially malicious actions like lookalike websites, malicious links the user might click on, et cetera, and then warns the user not to do those things. So an AI agent helping a user avoid accidental bad actions could be helpful; no?"

To which I say, could be helpful, yes. And I very much like the idea. As we were saying, the talk that Leo and I will be holding during the Zero Trust World event is titled "The Call is Coming From Inside the House." That is obviously a metaphor for what I believe to be the biggest and most intractable problem facing today's enterprises. You know, stated as succinctly as possible, the problem is "Overprivileged users making mistakes." Though the term "overprivileged" is typically used in a derogatory context, I don't mean it that way at all. I'm using it in its strict computer science context, where "overprivileged" is the result of not following a strict "least privilege" model.

The beauty of describing the problem as "overprivileged users who make mistakes" is that it points us toward two solutions to the problem: Either we no longer overprivilege our users, or we somehow arrange to prevent them from making mistakes. Whereas it might be possible to constrain what the employees of an enterprise might be able to do on the enterprise's network, a large part of the joy of using a personal computer is that its use is "personal," which is to say almost entirely unconstrained. We can do anything we want with our own machines. Since operating within a "least privileged" environment is no fun, and would not be tolerated by personal computer users, that suggests that the solution for personal computer users lies in somehow arranging to prevent their mistakes, something previously not thought possible.

So to that end, I love the idea of some form of active client-side local AI agent that carefully scrutinizes everything the user is seeing and doing and interposes itself between the user's actions and the computer system. Leo, I, and our listeners know how to examine a URL to detect trickery. But even the best of us are not always paying 100% attention. And even when we are, we might miss the use of embedded Unicode characters to create a lookalike URL. Or in our haste we might click a link without first carefully checking all the way back to the right of its domain portion to make sure that its top-level domain is what we expect.

So by all means, the idea of having an AI agent peeking over our shoulder and watching our back to help detect the things we might well miss, I think it makes all kinds of sense. And needless to say, I hope that it would totally freak out if its user were getting ready to paste the contents of the clipboard which was pasted from their web browser into their Windows "Run" dialog. So, you know, if Microsoft wants to deploy AI, Leo, I would so much - instead of having something recording everything I do, I would much rather have something watching, you know, running locally, not phoning home, but making sure I don't click a link in email that could get me in trouble.

**Leo:** Yeah.

**Steve:** I am 100% bullish. And I'll bet you we're going to end up seeing that.

**Leo:** Yeah.

**Steve:** You know, you and I have complained for years that antivirus software has essentially become pass, obsolete, you know, I don't know anybody who would install it now except that they have a loyalty to packages that they were using in the past, and so that has endured. I don't run any. I just, you know, I'm careful about what I do. And I assume that Windows is going to find something. And it never has, except it's found my own code, which is really annoying because, you know, that's just what it does. So I've had to isolate a whole tree of my directory system in order to say "leave it alone."

And in fact, I discovered that in Windows 11, there's something coming called a "dev drive" because their own AV has become so intrusive and such a problem that they said, okay, we're going to create a thing in Windows 11 called a "dev drive" where we're going to give you responsibility for what's there because they've had no choice. They're driving anyone developing on Windows crazy by their - because, I mean, in order to protect them they have to delete anything onsite.

**Leo:** Right. I mean, it's becoming universal. Apple's going in that direction. Google's going in that direction. Everybody's doing that. Code signing is the future, I think, unfortunately.

**Steve:** Yeah.

**Leo:** I use, by the way, I use Claude now to do security audits on everything. And it's very good at finding security flaws and fixing them.

**Steve:** And we covered a couple instances of that last week where there was a guy who was running Claude, he had a WordPress site and server and had a bunch of WordPress add-ons and had Claude checking them before he put them online. And in one case it found many problems. And in some it was a really bad problem that he was like, you know, really glad for. So we're going to see a lot of cleanup on aisle nine, I think.

**Leo:** Claude, cleanup on aisle nine. Bring the mop. All right. We're going to take a break, and then we're going to go KongTuke, do a little Klingon in just a bit. You're watching Security Now!, the early edition of Security Now!. Don't get your hopes up. We're going to go back to Tuesday after this week. But for those of you who have a free Sunday and can watch the show, it's great. We're glad you're watching live. We do this stream on YouTube and Twitch and X and Facebook and LinkedIn and Kick and of course in our Club. Lots of people do like to watch live. But you can always download copies from a variety of places. I'll tell you where at the end of the show. And now, KongTuke.

**Steve:** So, yes. I wanted to finish today's podcast by sharing a newly arrived spin on the ClickFix attack which we've discussed previously, and which has me really worried. Remember, that's the attack where the familiar CAPTCHA "prove you're human" test is maliciously extended to ask its victim to please paste the contents of the Windows system clipboard into the "Run" dialog and just press Enter. Just trust us. Prove you're human by doing that. Right. In the newer form of this, which its discoverers, Huntress

Labs, that's the name I couldn't remember at the top of the show, Huntress Labs, they dubbed this "CrashFix" because the victim's web browser is made to appear hung, broken, or defective, thus crashed. And as for the Klingonesque "KongTuke," it's the name Huntress Labs has given to this threat actor, which they've been tracking for the past year.

So Huntress wrote: "In January '26, Huntress Senior Security Operations Analyst Tanner Filip observed threat actors using a malicious browser extension to display a fake security warning, claiming the browser had 'stopped abnormally' and prompting users to run a 'scan' to remediate the threats."

**Leo:** That looks very credible. I would fall for that.

**Steve:** Yes. That is why this is so compelling. This could come up, and you would think, oops, okay. It's exactly what a Microsoft popup looks like.

**Leo:** Yeah.

**Steve:** It says "Microsoft Edge stopped abnormally." Then it says: "Microsoft Edge has detected potential security threats that may compromise your browsing data." Oh. That's not good. You would believe that. And then there's a "Run Scan" button. And you'd think, oh, scanning is good. And then down below there's a checkbox, checked by default, "Help make Microsoft Edge better by reporting current system information." And of course you would think, oh, I've got to prevent this from getting other people. So, I mean, again...

**Leo:** Is this all it takes? If you hit that button you're done?

**Steve:** No. No. Not yet.

**Leo:** No, okay.

**Steve:** So that's the good news. But it does get you involved; right? They said: "Our analysis revealed this campaign is the work of KongTuke, a threat actor we've been tracking since the beginning of 2025. In this latest operation, we identified several new developments: a malicious browser extension called NexShield that impersonates" - get this, Leo - "the legitimate uBlock Origin Lite ad blocker" - that impersonates it by stealing its source code - "a new ClickFix variant we've dubbed CrashFix that intentionally crashes the browser, then baits users into running malicious commands." I forgot to mention, it doesn't make it that clear here, they have script which does just bring the browser down.

**Leo:** So that was a real crash.

**Steve:** It was a, yeah, well, no, it was - no, because it invokes their dialog next. But they do crash the browser.

**Leo:** So your browser is dead.

**Steve:** Your browser's dead.

**Leo:** And you get their dialog, and it's credible because your browser's dead.

**Steve:** Exactly.

**Leo:** Yeah.

**Steve:** Exactly. So they said: "Ironically, the victim" - the victim who got infected by this. "The victim," they wrote, "was searching for an ad blocker when they encountered a malicious advertisement. The ad directed users to the official Chrome Web Store's NexShield Advanced Web filter."

Then they said: "The deliberate targeting of domain-joined machines" - which is what this thing ends up doing - "suggests KongTuke is after corporate environments where a foothold means access to Active Directory, internal systems, and lateral movement opportunities. Home users..."

**Leo:** This is terrifying.

**Steve:** Yes, it is. And look at this next page where you see the next stage of this attack, which is what you get after you click the scan button. Then you get the familiar open Win, you know, press WIN+R, press CTRL+V, press ENTER, bing bing bing.

**Leo:** Oh, that's interesting. So they put on your clipboard the malicious code.

**Steve:** The attack, yes.

**Leo:** So you don't even see that text.

**Steve:** Nope. All you do is follow the instructions.

**Leo:** Oh, and but here it is again. Edge fix browser hash.

**Steve:** Yup. Yup. So they said: "Home users on a standalone workstation receive a separate infection chain." So they have an enterprise infection chain and a home user infection chain.

**Leo:** Wow.

**Steve:** They said they receive an infection chain.

---

**Leo:** This is sophisticated as hell.

**Steve:** Yes.

**Leo:** Good lord.

**Steve:** It appears - oh, and they said that the home infection chain appears to still be in testing. They said: "When we got through all the layers, the C2 (command and control infrastructure) responded on a home environment with 'TEST PAYLOAD.'" Meaning it didn't do anything yet.

**Leo:** Oh, wow. We're just testing.

**Steve:** They said: "Whether this means non-domain targets are lower priority or the campaign is still being built out, one thing is clear: KongTuke is evolving their operations and showing increased interest in enterprise networks." They said: "So what are CrashFix and NexShield? The malicious NexShield app is all the more diabolical by being a fully functioning working clone of the authentic open source uBlock Origin Lite browser extension. So its users will be pleased with the ad and annoyance blocking behavior of the extension they have just found and installed.

"But after using their browser for a while, the bogus 'Microsoft Edge stopped abnormally' display will appear with its 'Run Scan' button. Upon pressing that, the user will be presented with a fake 'Security issues detected' alert and instructed to manually 'fix' the issue by opening the Windows Run dialog (WIN+R), pasting from their clipboard (CTRL+V), and pressing ENTER. The malicious extension silently copies a PowerShell command to the clipboard, disguised as a legitimate repair command. When the user follows these steps, they unknowingly execute the malicious command."

They said: "We were not about to blindly paste from the clipboard, so we tried copying the displayed command" - which starts with `edge.exe -fix-browser -hash= blah blah blah` - "like civilized malware analysts." That's what they're calling themselves, and of course they are.

They said: "The browser's response? A complete freeze. When your 'fix' causes crashes, the name writes itself. Thus they named this CrashFix. Before we go deep diving into how we ended up with a malicious pop-up message, let's take a step back and look at how it got delivered. You've probably heard the recommendation to install an ad blocker to protect yourself from malvertising, malicious advertisements that deliver malware through legitimate ad networks. Our victim likely just wanted to get rid of annoying ads. Instead, they downloaded a malicious one (NexShield) while searching for an ad blocker for Chrome.

"This header falsely attributes the code to Raymond Hill, the legitimate developer, as we know, of uBlock Origin, and references a non-existent GitHub repository. This tactic exploits the trust users place in well-known open-source projects. The actual uBlock Origin Lite repository is located at `github.com/uBlockOrigin/uBOL-home`, not the URL referenced in this malicious extension. The NexShield extension is almost entirely," they write, "a clone of uBlock Origin Lite, a legitimate extension by Raymond Hill. The threat actor likely ran a few find-and-replaces to replace every instance of 'uBlock' with 'nexshield.'"

Okay. So then Huntress continues with their analysis of this latest discovery of theirs. But for us, the takeaway is that the malware community at large has stumbled upon a fundamental security weakness of Microsoft Windows, which is its users' comparatively script-following level of understanding of Windows, when set against Windows' increasing power and sophistication. It's no longer useful to ask what can be done with Powershell and Dotnet. The question is, what cannot be done? That pairing, you know, Powershell and Dotnet comes to mind because, you know, while I was assembling today's podcast, I encountered other exploits which did exactly that. And this one is also using a Powershell command. It used native users' invocation of Powershell with a command they did not understand. They're just following instructions.

Now that we're encountering a proliferation of similar abuses of powerful commands escaping the browser with unwitting users blindly pasting and executing these commands that they did not write and do not understand, it should be clear that this story only has one ending. Sooner or later, Microsoft will need to step up to protect users from themselves, much as they did with the Mark of the Web, which flags files that were downloaded across a network. Files containing the Mark of the Web are treated much more cautiously and with skepticism by Windows. You know, you're asked, are you sure you want to run this? This was downloaded from the Internet. The system's clipboard needs to be handled similarly. Contents that were sourced by any web browser need to be quarantined.

Like I said, there's only one possible ending to this trouble. This problem is not going to go away because users are not going to get better or smarter suddenly. Let's hope Microsoft does not wait too long before updating Windows with this change. I wish I believed they would act responsibly here. You know, we can hope. And I'll just note that creating a third-party utility to fix this, because I kind of thought, well, maybe this - I should fix this. That won't help.

**Leo:** No.

**Steve:** Since it's all the people who would never know about such a utility who need it the most.

**Leo:** Right. Right.

**Steve:** You know, we don't need it, we listeners of the podcast. The only fix for this is to come - it's got to come from Microsoft as an integral part of Windows.

**Leo:** Yeah. It's got to be built in.

**Steve:** Yeah.

**Leo:** Yeah, or it's not going to happen. Give your folks Chromebooks, kids.

**Steve:** I just had a neighbor, as a matter of fact, and Lorrie and I encountered him while we were taking a walk yesterday. He was - he's an ex-engineer. He said: "I just got a Chromebook." He says: "I love it."

---

**Leo:** Yeah.

**Steve:** And he's: "I can't believe how everything imported into it." I mean, he was just blown away by it.

**Leo:** It's all most people need, honestly.

**Steve:** He is an Android user. And so when he explained that he connected his Android phone, I said, okay, well, that helps to explain its importation, at least. But...

**Leo:** Yeah, Microsoft considered this with Windows S. They really - and I wish they'd followed through. I think Apple, Microsoft should both offer Chrome OS-like very restricted environments. And then they can let those of us who know what we're doing use the less-restricted environments. Would be a real - a good solution.

**Steve:** Windows is way too powerful for most people. They don't need all of this. They get lost in directory hierarchies and directory privileges and, you know, basically you're running a machine you don't understand the operation of at all. And really, today, who among us does? We have a deeper understanding, but I remember the day, Leo...

**Leo:** I fall for stuff.

**Steve:** ...when we actually knew what the files were on our own hard drive.

**Leo:** We were editing out autoexec.bats, our config.syses.

**Steve:** And you remember my very tech-y friend Bob, you know, he was like, he was complaining, he, like, "I still know what everything does." And I said, "Well, Bob, good luck with that."

**Leo:** Not for long. Not for long, yeah. More, I mean, and basically that's what mobile OSes are, are highly restricted operating systems. They're not truly general purpose operating systems. Anything that's general purpose is going to be able to do anything, including run a network.

**Steve:** I can't even use my iPhone anymore, Leo.

**Leo:** No.

**Steve:** It's got, you know, like...

**Leo:** It's locked down.

**Steve:** Oh. Well, there's just so much in there. You hold this, and you...

**Leo:** Oh, it is, it's too complicated, yes.

**Steve:** You, like, slide something. You go three times to the right and click your heels, and then you get a magic dialog.

**Leo:** Too many hidden things, yeah. I spend many, many hours of my life looking through the settings, trying to find the setting that I want to change. And, you know, it's just so hard.

**Steve:** And remember, that was the breakthrough from Xerox PARC of the menu.

**Leo:** Right.

**Steve:** The commands were discoverable. You could browser around and see everything. And in fact that's one of the big changes coming in the next version of the DNS Benchmark that everybody will get for free is I put a traditional Windows menu on it instead of just overloading the system menu underneath the icon in the upper left. It's so much better. It's like, Gibson, come on, why did that take so long?

**Leo:** Nice. Well, we'll look forward to that.

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/>