



CISA's Free Internet Scanning

Description: The Security Now "Caption That Photo" contest. A mega social media company says "no" to strong encryption. WhatsApp to give parents more control. Consumer bandwidth proxying is becoming a big deal. Meta buys the Moltbook duo. The EU gives up and settles upon the status quo. When a ransomware negotiation is not what it seems. CISA compels federal agencies to submit their logs. Is that a VPN in your pocket or something more malicious? Be careful what you download, thinking it's AI. A super-clever and super-simple AV scanner bypass. Will AI write code for me? Another listener discovers the Joy of AI. Steve's CISA Internet scanning experience.

High quality (64 kbps) mp3 audio file URL: <http://media.GRC.com/sn/SN-1070.mp3>

Quarter size (16 kbps) mp3 audio file URL: <http://media.GRC.com/sn/sn-1070-lq.mp3>

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. Lots to talk about this week. We need a caption for our photo of the week. Maybe you can help. A social media company, another one says no to strong encryption. That's not a good sign. There is a problem with proxies serving malware, and it might even be coming from your router. We'll tell you how to find out. And then he's going to talk about his experience using CISA's Internet Scanner. All that coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 1070, recorded Tuesday, March 17th, 2026: CISA's Free Internet Scanning.

It's time for Security Now!, every Tuesday. I know you're looking forward to this, and I am, too. We get together with this guy right here, Mr. Steve Gibson, our security guru, to talk about the latest news. And there is always a lot of security news.

Steve Gibson: It is true, Leo, a small, very small subset of the world looks forward to...

Leo: Not that small.

Steve: Sometimes it's Monday morning, or, I mean, Wednesday morning typically.

Leo: Well, let's put it this way. We had about, when we were last couple weeks ago at Zero Trust World, I think there were 1,800 people in our audience. We have, I don't know, maybe eight times that for every show. So it's a lot more people who

are listening today than were there in the - although a live audience you're very aware of them. The podcast audience, we don't know.

Steve: Yeah. Although, as you'll see, this week's Picture of the Week issued a caption photo contest.

Leo: Ah.

Steve: So, you know, I invited our listeners to caption this photo. And boy, did I get replies. Well, 20,191 pieces of email went out Sunday saying, have we got any ideas? Oh, boy, you know, I got ideas back. So...

Leo: All right. Well, we'll see that in just a second.

Steve: This is Episode 1070 for - we're crossing over the middle of March. It's the 17th. I decided that I wanted to share the results from my first successful interaction with CISA's free Internet scanning because I'm now in a position to know it, like what it is, and to be able to recommend it without reservation to anybody who's got more than one IP that is, you know, DHCP issued by their ISP. Small, medium, large enterprise, I qualified. And as we know, I'm not running anyone's water filtering for the municipality or anything. I'm just GRC. But so it turns out that that barrier which they talk about as this is for, you know, government agencies and local, state, and federal, you know, no. It's commercial enterprises are considered infrastructure in a very broad definition.

So anyway, I'm going to tell everybody everything that I came away with from that, and also what it found in GRC's network that, okay, I knew about it, but still it was interesting. And it was a little bit of a "cry wolf." But we're going to talk about the Picture of the Week, of course. Also a mega social media company has decided to say no to their own strong encryption on their own messaging, which is interesting.

Leo: Ugh.

Steve: Yeah. And what does that mean? WhatsApp is going to give parents more control, which we'll discuss that. I think that's also good. Consumer bandwidth proxying that we were just talking about in the context of that Bright Data sort of semi-slimy Smart TV API, turns out it's becoming a big deal. And I guess in retrospect not that big a surprise, that is, consumer bandwidth proxying. Also Meta has purchased the Moltbook founder duo, try to say that three times. We'll talk about that. The EU has given up and is settling upon a compromise with that controversial chat control. Oh, turns out that ransomware negotiation may not be always what it seems, which should come as a surprise. CISA is compelling federal agencies to submit their logs to them. What?

Also, is that a VPN in your pocket, or maybe is that something more malicious? We're going to answer that question. Also be very careful about what you download, thinking that it might be AI. Once again, bad guys jump on anything that is popular, taking advantage of the enthusiasm of the moment. We've got a super clever and also worryingly simple means of bypassing AV scanners that a security researcher came up with. I'm going to answer the question I keep getting from our listeners, which is whether AI will be writing code for me. And I've got an interesting couple of well-

informed postings to share about that, followed on the heels of another listener of ours discovering the joy of AI.

And then I'm going to share my experience with CISA's free Internet scanning and unreservedly promote it to our listeners' enterprises. I just can't think of a reason why anyone who was able to and was qualified wouldn't want to enlist another piece of, you know, another set of eyes looking at and confidentially reporting what they see from the outside.

Leo: Excellent.

Steve: So I think, Leo, maybe it's worth tuning in this week.

Leo: Well, you're done so already, so it's too late. And I should mention it is St. Patrick's Day. So I shall be disappearing from time to time to check my corned beef to make sure it is doing its thing.

Steve: And are four leaf clovers a result of Chernobyl radiation, or do they exist in nature?

Leo: Ah, that's a good question. Well, they do occur in nature, I know that, because we had them before Chernobyl. But I wonder if there are more of them than there used to be.

Steve: Aren't they normally three, and they get...

Leo: Normally they're three. They are a mutation, I believe, yes.

Steve: You know, Mark Thompson went to Chernobyl with a group, like he thought that would be a cool place to go walk around. And he did report that there seemed to be an abundance of four-leaf clovers.

Leo: Aha. That's a very interesting experiment.

Steve: What made me think of it, yeah.

Leo: Hmm. We will get to our Picture of the Week and your caption contest in just a moment.

Steve: So Leo, before you look at the photo, I will just tell you that all I wrote across the top of it was "Security Now's 'Caption That Photo' Contest."

Leo: Okay.

Steve: And when you scroll up, you'll see why.

Leo: Oh, boy. Oh, boy. Now, we were talking about this. I don't know where this is. But Paul Thurrott and I were talking about this in Mexico. He lives in Mexico City. This is what the phone poles look like because if something doesn't work, they don't figure out what's not working. They just put a new one in. So many of these wires are probably non-functional. Tell us what we're looking at here.

Steve: So, well, when I was growing up, we would have called this a "rat's nest."

Leo: Yes.

Steve: And it is someone atop a - it's hard to describe this as a telephone pole, although these look like phone lines coming in.

Leo: There's one in there somewhere, I think.

Steve: And look, there's like boxes hanging from wires.

Leo: Wow.

Steve: And various-sized junction containers. And I do notice there's a lot of loopage, you know, like rolls of wire that are hanging. It would be really interesting actually to know where. And as you noted, when something goes wrong, they just string another one. It's difficult to imagine that this actually functions. And one wonders how long ago this began to allow this to occur to it. It's just - anyway, so it's all I said was - I didn't even - I didn't have a chance to talk about it on the email that I sent out. But our 20,191 recipients said, oh, I've got a name for that. And so the responses have been pouring in. In response to something that came in early, that gave me an idea for what I think is going to probably - that I'm going to suggest as the winning caption. But we will see next week. In the meantime, those who are just listening to this, I don't think I could adequately...

Leo: No.

Steve: ...prepare you for what you would actually see if you saw the photo in this week's show notes. It is beyond insane. And Leo, how did he get up there?

Leo: Yeah.

Steve: Like he must have, like, had a crane plant him on the top of this.

Leo: Yeah.

Steve: Because you can't climb the - well, I guess you could climb the side. But then who knows how many wires you'd pull loose. So, wow.

Leo: That's amazing.

Steve: And I've had this photo in my Pictures of the Week candidate pile for quite a while. And finally I thought, okay, let's just see what our listeners think about this.

Okay. So last week the news - and we talked about this, of course - was that TikTok had decided and formally announced that it would not be adding end-to-end encryption to its already-controversial enough short-format video sharing platform. Right? They said that - that is, TikTok said - that we want to enhance our users' security, and doing that means being able to screen the content that our users are sharing and prevent illegal content from being shared. So they said that.

Then, somewhat surprisingly, last Friday The Hacker News reported that Meta, of all people, or all groups, all companies, had announced their somewhat similar plan to back encryption out of Instagram. What? So The Hacker News wrote: "Meta has announced plans to discontinue support for end-to-end encryption for chats on Instagram after May 8th, 2026." So I guess this was like a 60-day notice; right? March, April May.

They said: "The social media giant said in a help document: 'If you have chats that are impacted by this change, you'll see instructions on how you can download any media or messages you may want to keep.'" Which I thought was interesting. How is keeping messages relevant to ending end-to-end encryption? Maybe they're just going to start over. I don't know. Like get rid of everything that has been in the dark that they haven't been able to see so that from now on, any new messaging will be without end-to-end encryption. Anyway, they said: "If you have an older version of Instagram, you may also need to update the app before you can download your affected chats."

Hacker News said: "When reached for comment, this is what Meta had to say: 'Very few people were opting in to end-to-end encrypted messaging in DMs, so we're removing this option from Instagram in the coming months. Anyone who wants to keep messaging with end-to-end encryption can easily do that on WhatsApp.'" Okay. The Hacker News said: "The American company first began testing end-to-end encryption for Instagram direct messages in 2021 as part of CEO Mark Zuckerberg's 'privacy-focused vision for social networking,'" which we all remember at the time. They said: "The feature is currently 'only available in some areas' and is not enabled by default." Then they said: "Weeks into the Russian-Ukrainian war in February 2022, the company made encrypted direct messaging available to all adult users in both of those countries."

"The development comes days after TikTok said it does not plan to introduce end-to-end encryption to secure direct messages on the platform, telling BBC News that the technology makes users less safe, and it wants to protect users, especially young people, from harm."

"Last month, Reuters also reported that Meta proceeded with plans to adopt encryption to secure messages in Facebook and Instagram despite internal warnings in 2019 that doing so would hinder the company's ability to detect illegal activities, such as child sexual abuse material (CSAM) or terrorist propaganda, and then flag those illegal activities to law enforcement."

They said: "End-to-end encryption has been hailed as a win for privacy, as it ensures that only communicating users can decrypt and read messages, thereby locking out service providers, bad actors, and other third parties from accessing or intercepting the data."

However, law enforcement and child safety advocates have argued that the technology creates a safe space for criminals, as it prevents companies from complying with warrants to turn over message content, a problem referred to as the 'Going Dark' phenomenon. This year, the European Commission is expected to present a Technology Roadmap on encryption" - we'll have a little more to say about that in a minute - "to identify and evaluate solutions that enable lawful access to encrypted data" - good luck with that - "by law enforcement, while safeguarding cybersecurity and fundamental rights."

Okay. So I think this is interesting, and I wonder whether this signals the start of a gradual backing away from providing strong encryption to consumers on the mega-popular generic platforms. I doubt whether most lawful users of TikTok, Instagram, or even WhatsApp really care all that much about encryption. Sure, if they can have it for free, and if it's built-in, and if it doesn't cause them any trouble or headaches, sure, okay, fine, they'll take it. But is even a single person going to walk away if it's removed? I doubt it. While there was an initial rush on the part of publishers to provide it, like in 2019, with Zuck's big privacy first business, I don't think it's ever been shown that there was any actual consumer demand. Anyone who really wanted secure messaging, after all, could switch to Signal, which is also free, and where Meredith maintains unflagging vigilance at the gate.

So the way we're seeing things shake out, I suspect that the right solution to all the mess and pushback to this messaging, you know, well, to the increasing prevalence of fully encrypting everyone's random messages on consumer platforms by default, is simply not to bother with it, and no one will much care. I know this will make the privacy-at-all-costs people's heads explode; but, again, Signal is always available, as is Telegram, and is free for anyone who actually wants it. For those who worry about grooming and CSAM, removing always-on encryption by default from the major platforms will tend to eliminate that opportunistic abuse. It won't be on. And so the bad guys can't safely do that. And in fact eventually I think it won't even be an option. So I'd be interested, Leo, to know what that gal, you had an EFF person on recently, wonder what she had to say about all this.

WhatsApp, however, is also moving in a parent-forward fashion. Meta also announced the addition of parent-managed accounts for WhatsApp. The accounts are designed for pre-teen children where access to account settings will be controlled by a PIN set by the parent. Essentially, parents can control settings, lock those settings on their children's devices, their underage teen, you know, pre-teen children's devices, and obtain some control over it.

The message content on the pre-teen accounts will remain private, so this is not a privacy invasion. It's a setting controls lock. Parents will be able to approve to whom their children may speak, what groups they can join, and review message requests from unknown contacts. So do a little bit of sort of at-distance management of what their kids are doing, keep their kids from changing that stuff. Basically parental controls for WhatsApp. And I think, you know, that that seems to make a lot of sense to me and seems like a good thing.

Last week we looked in some depth at the company "Bright Data," whose unfortunate business model involves arranging to offer end users, not directly from them, but by virtue of streaming partnerships and Smart TV partnerships, offer end users the ability to lower their costs, either for streaming and/or see fewer advertisements in return for the privilege of routing third-party Internet traffic through their ISP-purchased or subscribed bandwidth, and thus using their residential consumer IP address. And as we noted last week, there's only one conceivable reason for doing this, which is to allow those third parties to mask their identities and hide whatever their purpose may be among the world's broadly distributed consumers.

The issue of consumer proxies was again in the news after we talked about it last week for another reason. The Risky Business news late last week opened by writing: "American and European law enforcement agencies have seized the infrastructure of a residential proxy provider named SocksEscort, the latest such crackdown against proxy providers over the past years." And again, this is like a growth interest on the Internet, this idea of proxies, because the Internet is getting much better about filtering, and proxies are a way to bypass filtering.

Risky Business News wrote this SocksEscort service had been running since 2021 and rented access to more than 369,000, so more than a third of a million, 369,000 different IP addresses, not all at once, but across its entire lifetime. So they came and went over time. Generally there were several tens of thousands at any given time. "According to the FBI," they write, "Europol, and Dutch Police, SocksEscort was a front for a malware operation that infected modems and home routers."

In other words, unlike Bright Data, which is hopefully an aboveboard, only with user permission, and hopefully with user understanding, asking to reuse consumer bandwidth, this is malware. These are, you know, leveraging router vulnerabilities in order to get these proxies installed and then obtain persistence. So in other words, malware proxies, not benign bandwidth-bouncing proxies. They were maliciously installed without their hosts' knowledge or permission to forming a proxy botnet. And of course we've talked about proxy botnets through the years because this IP-based blocking, as I said, has been growing, and the bad guys are needing to obscure their bandwidth.

The article continues, writing: "Lumen's Black Lotus Labs linked this group to a botnet it discovered in 2023, named AVrecon. The botnet never grew to an extremely large size, but managed to maintain," they write, "a healthy pool of IP addresses it could rent out to its customers, most of which were other cybercrime operations needing ways to hide their attacks inside the infrastructure of residential Internet providers. Europol linked the service to ransomware deployments, DDoS attacks, and the distribution of child sexual abuse material. It also estimated that SocksEscort operators made more than 5 million euro from renting their infected IPs, which they noted is quite the sum for a service as simple as, you know, proxying.

"On the day of the takedown," they write, "the FBI published an advisory with tips on how telcos and consumers can protect their devices and prevent them from ending up as nodes in proxy networks. It also published advice on spotting and removing specifically this AVrecon from residential devices.

"Over the past few years," they said, "the U.S. has mounted a war against residential proxy networks after several reports concluded that foreign adversaries were using infected American routers to hide their tracks. Law enforcement takedowns have targeted both private proxy networks like ORBs, or Operational Relay Box networks, but also residential proxy providers. The difference between the two is that ORBs are typically built and managed by the threat actors for their sole use" - so those are essentially proxies installed somewhere - "while a residential proxy provider is a service built for an operator's financial gain, typically rented out to whoever has the money."

And they finished, saying: "Past proxy-using botnets that were taken down include 911 S5, Anyproxy, 5socks, RSOCKS, Flax Typhoon's Raptor Train, Volt Typhoon's KV Botnet, APT28's Moobot, VPNFilter, and others." So in other words, the idea of proxying is a hot commodity on the Internet today.

Our takeaway is that while bad guys, again, probably have very little interest in the contents of any random person's internal network, and for that we can be thankful, and let's hope that doesn't change soon, there is substantial interest in using and abusing any distributed bandwidth they are able to obtain. Being able to hide and admit their junk,

whatever it is, attacks, probes, whatever from residential IPs, the IPs of users who have no idea that's what's going on, that's of huge value to them.

In fact, way back in time, when I tracked down that kid that had been DDoSing GRC, it was a - I got the FBI to work with me. I had the IP address of a source of the attacks because the source IPs were not spoofed. We located a family a few miles from me, and I made a house call and looked at their computer. It was infected. They had no idea this was going on behind their back. They were horrified. And of course I was interested because I wanted to get a sample of this thing in order to reverse engineer it, which I did. And in return for that I disinfected their computer for them. But that's an example of, you know, this happening behind people's backs, and nobody had any idea.

So we've also learned that "substantial interest in," you know, I said there is substantial interest in using and abusing any distributed bandwidth the bad guys can obtain. And what we know is that "substantial interest in" equates to substantial pressure to get in. That is, you know, bad guys want in to people's NAT routers. So keeping the bad guys out means resisting any temptation to rely on a border router's authentication mechanisms. We see, time and time again, you just can't.

Any NAT router without any deliberately exposed WAN-side services is going to be inherently bulletproof if traffic is only originated from inside and is only allowed to come back in from outside when it matches what first went from inside out. So it's a firewall unless you poke holes in it. Poking holes in it means unsolicited connections from the outside in because, for example, you just couldn't resist turning on remote web access to your router's management interface.

Leo: Resist that.

Steve: Please. Please resist. So...

Leo: I use Tailscale to open up...

Steve: 100% safe.

Leo: That's okay; right? Yeah.

Steve: Yes, because Tailscale is outbound NAT penetration. And you are not opening, you know, you are not able to, from Starbucks, you know, go <https://> and then your home IP and be looking at your router's, oh, log into your ASUS. No. Don't do that.

Leo: No, no, no, no, no. Turn off WAN.

Steve: It's only when consumers decide to deliberately expose external management, you know, access to their router-hosted services that authentication bugs in the router's firmware can be leveraged to install and maintain proxies. So, and again, it's like everyone's false thought is, well, who would want to get into my router? Who would want to get into my network? I don't have anything. The fact that you have a router is valuable. That creates pressure to get in because they want to set up shop and use your

bandwidth and use your IP. And also you don't want your IP associated with all kinds of dastardly deeds on the Internet. That's not good for you either.

Leo: And there's some interest in what does it take to get a house call from Steve Gibson? Asking for a friend. That's special treatment, let me tell you, folks. So does the proxy server run on the PC, or does it run on the router?

Steve: It's on the router.

Leo: Okay.

Steve: So it is, yeah, so it's a little daemon that is set up in the router. It's added to the router's boot code so that it comes back alive. And it reaches out to a remote command-and-control server to establish a contact. So even with it there, it doesn't open a port.

Leo: Doesn't have to.

Steve: It maintains its own stealth because it reaches out to the external command-and-control and then - so basically it phones home to establish a connection and then to await orders.

Leo: How do you detect it? It's not easy, I would imagine.

Steve: You've got to look at the actual - you've got to, you know, look at the actual...

Leo: [Crosstalk] zen map or something or [crosstalk].

Steve: Well, traffic or, I mean, unfortunately, and this is the problem, is most - the reason I paused there is that all the ways I could think of required you to know Linux. You know, I mean, you need to look at the shell script startup stuff and go, what the heck is that? That's not supposed to be there.

Leo: In other words, you need Steve to come over.

Steve: Don't attack me.

Leo: So if you reboot the router, is that sufficient?

Steve: Yes.

Leo: Okay.

Steve: Often times rebooting, because a lot of these things are unable to establish...

Leo: Just in RAM.

Steve: ...yes, they only live in RAM. So rebooting is the first thing. Reflashing is - that will also do it. So, like, you know, if you're able to just update your firmware or reupdate your firmware, that will also clear things out.

Leo: Good to know.

Steve: You know the other thing that's good to know, Leo?

Leo: What's good to know, Steve?

Steve: I know you know.

Leo: I know I know.

Steve: This next sponsor is good to know about.

Leo: Everyone should know about our next sponsor. I completely agree with you. And I would tell you about them if I just had put the right copy in there. So hold on just a moment while I get the...

Steve: Mesmerize our viewers.

Leo: Yes. Everybody look at Steve's coffee cup. Steve's coffee cup. And now, back to a fully caffeinated Steve Gibson.

Steve: Recaffeinated.

Leo: Recaffeinated.

Steve: Okay. So in case anyone was wondering, Moltbook, which was that weird facility that was affiliated with OpenClaw, where only OpenClaw's autonomous AI agents were able to talk amongst themselves, and we lowly humans were only able to look on, gawking in wonder at the inter-agent AI dialog, that was just purchased by Meta. I assume the guys started work there yesterday. I assume Meta's entire interest is in obtaining those two creators of Moltbook, Matt...

Leo: One of them is a good friend, by the way, Steve, Ben Parr, who's been on TWiT many times.

Steve: Oh, and Ben Parr.

Leo: Yeah, I didn't know Ben was Moltbook, or I would have had him on the show to talk about it all this time. He was kind of more stealthy than the other guy. The other guy got all the attention. Anyway, congratulations.

Steve: Yes. And I'm sure they're being well compensated. They both started working at Meta yesterday, on March 16th, in Meta's MSL, which modestly stands for - M is not for modest. M is for Meta, like literally this is what they call themselves: Meta Superintelligence Labs, MSL.

Matt [Schlicht] has been working on autonomous AI agents since 2023, and he launched Moltbook in late January as an experimental "third space," as they put it, for AI agents. And Moltbook was built largely with the help of his own, Matt's personal AI assistant, which he named Clawd Clawderberg. Okay. And of course his partner in Moltbook and now also at Meta, as you said, is Ben Parr, who was formerly an editor and columnist at Mashable and CNET.

Leo: And a good friend.

Steve: And a good friend of the show, yeah, of TWiT. So apparently, Moltbook continues to be available though Meta, although they indicated that they weren't certain what its future might be. So it's not clear whether they're going to bother to keep it going. But for now it is. The typical corporate-speak statement from Meta, as reported by Axios, was that "The Moltbook team joining MSL opens up new ways for AI agents to work for people and businesses," which of course says nothing. And I doubt that even they know what they mean by that. But that's how these sorts of acquisitions go where it's the people that are actually being acquired. Meta doesn't care about Moltbook at all. They just want those guys.

Leo: Although I imagine that they want to somehow capitalize on this agentic future.

Steve: Yes. And...

Leo: And extend Facebook to agents. Why not?

Steve: God help us, Leo.

Leo: I know. I mean, the real problem with Moltbook, besides the fact that it has had a terrible security model, was that humans could get in, too, so we never really knew...

Steve: If it was only AI-generated dialogue; right. Okay. So the good news is that the EU was unable to secure the votes needed to pass its most recent attempt to force all communication services to monitor their users' communications. I mean, we were balancing on a razor's edge there for quite a while. It's like, this could almost happen.

And finally Germany reversed their previous, yeah, we think that we probably should vote, and they said, okay, no, we're not going to. And that killed the whole thing. So what we have instead is an extension of the previous, what's been called "voluntary" chat control. Which, as I said, that's already been in place.

Last Wednesday, the 11th of March, Heise Online covered this news, writing: "The EU Parliament approved a renewed extension of 'voluntary chat control'" - which is in quotes because that's not really the official name, but that's what we all call it - "to combat child sexual abuse in Strasbourg on Wednesday. After the initiative surprisingly failed in the responsible committee a week ago, MEPs are now attaching clear restrictions to the extension. The regulation creates a temporary exception" - again, remember we were just talking about how COPPA would need to be amended, Leo, in order to allow, like, kids to disclose that they're children, but that would be a breach of COPPA because we're not supposed to know that.

Leo: Right.

Steve: Well, here we have the same kind of...

Leo: That would kind of be a hint that something's wrong here.

Steve: Yeah, we've got the same thing happening here because you can't even voluntarily look at people's data under EU regulations. So what we have is an amendment to the regulation creating a temporary exception to the European data protection rules, allowing messaging services to scan chats for depictions of child sexual abuse. There is currently no agreement on a long-term solution, which is, you know, which is what the EU Commission and member states were hoping to get.

"Providers of messaging services," Heise wrote, "may automatically scan their platforms for digital traces of child pornography. The search for adults who prey on minors, known as 'grooming,' is also under debate. Because this violates the EU directive on the protection of privacy, the EU hastily created an exception regulation in 2021. This exception regulation, which has already been extended once, now again is valid until the beginning of April and was supposed to be renewed until April 2028 at the request of the EU Commission. Last week, however, the Commission's proposal surprisingly failed in Parliament's Committee on Civil Liberties, Justice and Home Affairs." They're just having all this trouble with this.

"In a new compromise, Parliament has now agreed to an extension until August 2027. At the same time, MEPs voted for a clear limitation of powers to search for already known material, and only for users or groups suspected of concrete wrongdoing." Thus not just to blanket search everybody. "Furthermore, encrypted chats should not be affected." Well, actually, practically they can't be because they're encrypted.

"A spokesperson for the Committee on Civil Liberties, Justice and Home Affairs said: 'This exception is a temporary, strictly limited instrument that allows providers to continue their voluntary detection measures under certain conditions. The extension must also maintain end-to-end encryption.'

"These restrictions correspond to Parliament's draft for a long-term solution. These will be the subject of upcoming negotiations with the Commission and member states. Only when an agreement is reached here can the renewed extension come into force. There's currently no majority in Parliament for far-reaching surveillance powers such as arbitrary

chat control." That's what we were talking about before that Germany vacillated on, then it said no. "The Council of Member States has also moved away from this after a long struggle." Right. "However, this does not make a permanent 'voluntary' solution any easier, especially since it also affects the fundamental rights of EU citizens." Which are protected from this.

"While the Commission and member states want to make the controversial exception regulation permanent, the EU Parliament insists on significant restrictions. For example, error-prone technologies such as AI should not be used in the search for child pornographic depictions. Scanning text messages for grooming attempts should also remain prohibited."

So if anybody thinks this sounds like a huge mess, then you have been paying attention because, yes, this is - the EU just - they're in a big scramble and confusion.

Leo: Pickle. They're in a pickle.

Steve: Boy, yes. The good news is that saner heads prevailed, and since they weren't able to push anything forward, they at least didn't move anything backward. And companies that have been doing some of their own platform-based CSAM screening, as we know some major providers have, this gives them the cover to continue to do so without requiring them to do it, nor requiring them not to offer their own internal encryption for their users to whatever degree they wish to. So for now that's what we have. And it's probably the best that we could hope for. They're unwilling to drop it, but they are unable also to push it forward. So they're just extending the voluntary chat control, and maybe that'll calm down over time.

Leo: It's so telling that in both the U.S. and the EU, any attempts to do this have to require exceptions to existing privacy laws. It's like the age verification stuff in the U.S., they had, whoever it is, the Department of Commerce had to give an exception to the COPPA rules, Child Online Privacy and Protection Act rules, because, well, if you're going to ask people's ages, that's a violation. Isn't it telling that the thing you want to do is a privacy violation? That should tell you something. Oh, well.

Steve: Yeah.

Leo: I'm asking too much, I think.

Steve: Yeah. There was a piece that one of our listeners sent me that I looked at which, and I can't remember now what the publication was in, but the people were just going crazy calling any indication of - oh, I know what it was. It was that Meta had secretly been supporting nonprofits to the tune of \$2 billion, I think that was the number, across the country, for them to be pushing on behalf of the need for age determination and pushing Google and Apple to push this onto their platform.

Leo: Yeah, they were doing this secretly because they didn't want anybody to know...

Steve: Yes.

Leo: ...that they were behind this, yeah.

Steve: Yes. And my take is that this is where that should happen, that it should be...

Leo: I agree.

Steve: ...Apple who simply allows an API to be, I mean, the user still has control. If you want to go to an age-restricted site, before that happens, a dialog pops up and says the site or the app or whatever it is wants to know if you are an adult. Do you want to give them any indication? You can say no, in which case you may not be able to go there. Or you can say yes, I'm an adult, and I, you know, tell them. To me, I mean, this - I get it that there are people who want to give nothing. But it's just not - we also have laws throughout the world where age matters. Children can't drink alcohol. Children, we decided, cannot be exposed to aspects of human sexuality; you know? Children, you know, I mean, there is behavior that's regulated based on age that needs to get extended out to the Internet because the Internet is here to stay.

Leo: I think that's fair. I really do. I've come around a little bit on that.

Steve: Yeah.

Leo: Just have to find a way to make it work.

Steve: Yes.

Leo: And I think you're right. There is a chokepoint. It's Android and iOS.

Steve: Yeah.

Leo: And that's where this should happen.

Steve: Yeah. And the beauty then is that - and this is Meta's point, and they're right - then every individual provider doesn't have to keep, you know, coming up with their own solution. Because every independent solution is another opportunity for a privacy breach. And so, you know, doing things like looking at the camera and saying, oh, don't worry, we're not going to keep your photo, well, we've already seen examples where third parties did keep people's photos, and then they got breached. So, yeah. I trust Apple, and I would trust Google to engineer something for Android that's as good as we can get. And yes, you could still have absolute privacy, but then you're going to lose some access to content which your government has decided only adults should have. So you get to choose.

Leo: Yeah. Yeah, I think that's fair. And it's privacy forward.

Steve: Yes. Yeah. And it's as good as we can get. I mean, yes, you're going to lose some if you want access to adult restricted content. But your government has said, the government that you are subjected to has said, no, children can't have that. You just need to tell us that you're an adult. And the platform you're using needs to - you have to have shown that to the platform one time, let them check it, and then the platform remembers and can make that assertion on your behalf.

Okay, Leo, get this. This next bit of news just made me shake my head. I'm not going to spend too much time on it, but I didn't want to let it pass without comment. CyberScoop informs us that ransomware negotiators, right, working for the ransomware negotiation firm DigitalMint, that is, like companies that have been breached and that are under ransom, they bring DigitalMint in to negotiate on their behalf? They were also the ransomware attackers that they were negotiating with.

Leo: Oh. Oh.

Steve: So CyberScoop wrote: "A 41-year-old South Florida man is accused of conducting at least 10 ransomware attacks and helping accomplices extort a combined \$75.25 million in ransom payments while he was working as a ransomware negotiator for DigitalMint."

Leo: Oh, this has to be a movie. Somebody has to option this. This is too good.

Steve: Isn't this great? "According to federal court records unsealed last Wednesday, five of Angelo John Martino III's alleged victims hired DigitalMint, which assigned Martino to conduct ransomware negotiations on their clients' behalf, putting him in a position to play both sides, as the criminal responsible for the attack and the lead negotiator for his alleged victims." Really, you can't make this up.

Leo: I don't know, you know, these ransomware guys, they're really hanging in there tough. I think you're going to need to give them some more money. I don't know.

Steve: And they're just not - they really sound like they're not going to give.

Leo: They're really hanging in there.

Steve: "Martino allegedly," they wrote, "Martino allegedly obtained an affiliate account on ALPHV, also known as BlackCat, a criminal ransomware as a service group, and conspired with other" - get this - "other former cybersecurity professionals" - so, oops - "to break into victims' networks, steal and encrypt their data, and extort companies for ransoms over a six-month period. Prosecutors accuse Martino of providing confidential information regarding ransomware negotiations to ALPHV co-conspirators to maximize the ransom payment.

"The five U.S.-based victims that hired DigitalMint and unwittingly tapped Martino to allegedly conduct ransomware negotiations with himself and his co-conspirators include a nonprofit..."

Leo: I'm telling you, this is a movie, man.

Steve: I know, "a non-profit and companies in the hospitality, financial services, retail and medical industries. All five of those victims paid ransoms."

Leo: Wow.

Steve: So anyway, CyberScoop's coverage of this continues at some length, but everyone gets the idea here. On the one hand, this obviously puts the guy who's negotiating both sides of the deal, as you noted, Leo, in the position to know exactly how much ransom his victim will actually pay.

Leo: Now, just between us guys, what's the maximum you'd be willing to pay? Just [crosstalk].

Steve: Yeah, exactly, we're not - we probably don't, you know, we don't want to go there, but just so we know...

Leo: Just so we know...

Steve: ...what do have to work with? Now, on the flipside, the upside, such as there is, is that the negotiator is also in the unique position to know for sure whether the attackers, since that's also him, will actually honor their promise to restore the victim's data and delete any copies they might have.

Leo: I'm pretty sure that if you'll give these guys a million dollars, they're going to give you the key. I'm pretty sure.

Steve: That's right.

Leo: I can't promise, but I have a good feeling about this.

Steve: Seems like, yeah, when they're talking, they seem like, you know, they're - obviously they're bad guys, but they seem like good bad guys.

Leo: This is a gutsy fellow. That's - wow.

Steve: Well, he's a gutsy fellow in chains right now. Yeah. And boy, the article had pictures of aerial photos of his estate in South Florida, you know, and a 224-foot yacht that was docked on his pier.

Leo: Oh, it was working for him. It was working for him.

Steve: Yeah, he wasn't hurting.

Leo: Oh, my god.

Steve: And he was married. You've got to wonder what his wife thought. Like, you know...

Leo: Honey?

Steve: I mean, he doesn't really work that much.

Leo: What do you do for a living?

Steve: He closes his office door and mumbles into the - yeah.

Leo: I've got a very important meeting with myself. Just I'll be back later.

Steve: Yeah. Let you know how it goes. So three weeks ago, during Episode 1067, we covered the news of yet another horrific CVSS 10.0 in Cisco's, courtesy of Cisco, SD-WAN product. This is that bug behind CVE-2026-20127, another critical authentication bypass in Cisco's Catalyst SD-WAN. And the reason I say "another" is it had an additional one back in 2020. It's hard to get those right, especially for Cisco. In this case this allows unauthenticated remote attackers to gain admin-level access to SD-WAN controllers to compromise entire WAN infrastructures.

Last Wednesday, CISA revised their previous orders which we covered three weeks ago. Three weeks ago CISA was saying you needed to update by such and such a time. They had a whole, you know, calendar laid out. CISA has now ordered all federal agencies to upload their logs from Cisco's SD-WAN devices to CISA's own cloud platform by next Monday, March 23rd. These Catalyst SD-WAN devices have been under attack, as we know, using a zero-day since, as we said at the time, still true, 2023. Wow. And a great many of Cisco's customers have done nothing about it in the past three years.

While CISA has no jurisdiction over private enterprises, it does over federal agencies. It has been given that jurisdiction. This uploading and aggregating of the logs on CISA's platform will allow CISA's people to investigate which agencies have been compromised. So Leo, you were wondering, you asked the question, like, how would a consumer know if their router...

Leo: Right.

Steve: Well, not easily. But in the case of SD-WAN logs, look, you morons, just configure your device to send your logs to our cloud platform. We will look at them for you and let you know, if you've got a problem. So, and I imagine the first thing they'll do is, like, why have you not updated your firmware on your SD-WAN? So agencies will have to configure their Cisco SD-WAN to send future logs to the same Cloud Logging Aggregation Warehouse, which is known as CISA CLAW.

Leo: No.

Steve: The Cloud Logging Aggregation Warehouse.

Leo: Interesting.

Steve: Yeah, clawing back the data. Now, the past year, as we've talked about, has seen a huge upward trend in the use of VPN services for geo-relocation. Why? Well, this increase in VPN use has been driven by new regional legislation which forces providers of age-restricted content to block access based on the geo-location of their would-be visitors, thus appear to be somewhere else.

Unfortunately, a new demand and a rush to something - whatever, anything, AI, geo-relocation, you name it, what is the current enthusiasm - that rush creates new opportunities for bad guys to take advantage of the inexperience of newbies who are entering a market that's new to them. We've previously noted that this has been happening with VPN add-ons for Chrome. Microsoft security has been tracking a group they identify as Storm-2561, which has been using search engine optimization (SEO) poisoning to provide malicious links to unwitting Windows users who are looking for VPN client software.

Microsoft writes: "In mid-January 2026, Microsoft Defender Experts identified a credential theft campaign that uses fake virtual private network clients distributed through search engine optimization poisoning. The campaign redirects users searching for legitimate software to malicious ZIP files on attacker-controlled websites to deploy digitally" - and here it's interesting - "digitally signed" - wait, what, digitally signed? - "digitally signed trojans that masquerade as trusted VPN clients while harvesting VPN credentials. Microsoft Threat Intelligence attributes this activity to the cybercriminal threat actor Storm-2561.

"Active since May of 2025, Storm-2561 is known for distributing malware through SEO poisoning and impersonating popular software vendors. The techniques they used in this campaign highlight how threat actors continue to exploit trusted platforms and software branding to avoid user suspicion and steal sensitive information. By targeting users who are actively searching for VPN software, attackers take advantage of both user urgency and implicit trust in search engine rankings. The malicious ZIP files that contain fake installer files are hosted on GitHub repositories, which have since been taken down." But of course GitHub, you know, engenders trust. "Additionally," they said, "the trojans are digitally signed by a legitimate certificate that has since been revoked.

"This blog," writes Microsoft, "shares our in-depth analysis of the tactics, techniques, and procedures (TTPs) and indicators of compromise in this Storm-2561 campaign, highlighting the social engineering techniques that the threat actor used to improve perceived legitimacy, avoid suspicion, and evade detection. We also share protection and mitigation recommendations, as well as Microsoft Defender detection and hunting guidance.

"In this campaign, users searching for legitimate VPN software are redirected from search results to spoofed websites that closely mimic trusted VPN products, but instead deploy malware designed to harvest credentials and VPN data. When users click to download the software, they are redirected to a malicious GitHub repository" - they say again no longer available - "that hosts the fake VPN client for direct download."

Okay. So I'll note that while Microsoft keeps reinforcing that the malware has been taken down, they know as well as we do that no sooner will one set of malware be taken down than its replacement will appear. In fact, it's more often the case that multiple sets of redundant malware have already been staged in place on GitHub and are just waiting to be linked to when the current malware in use is removed. This allows that malware to age a bit on the platform to increase its appearance of authenticity. So, a takedown of one set, while certainly useful and necessary, should by no means suggest to anyone that the threat has been in any way diminished. This is a classic case of Whac-A-Mole. And while it's true that the game must be played, it can never be won by playing catch up. Another mole will always be ready to pop up somewhere else.

Microsoft continues to explain: "The GitHub repo hosts a ZIP file containing a Microsoft Windows Installer (MSI) installer file that mimics a legitimate VPN software and side-loads malicious DLL files during installation. The fake VPN software enables credential collection and exfiltration, while appearing like a benign VPN client application." So, for example, an unwitting user believes they're getting a VPN. They download the VPN, install the client, activate the client. It says it's connected to the remote VPN server. And they then go to wherever they are wanting to VPN to and log in. None of that is true, so the bad guys obtain the credentials they use to log into wherever they were trying to VPN to. So it is very crafty. And, I mean, this is the way enterprises end up getting penetrated and being ransomed by somebody from DigitalMint, who's working for the bad guys or themselves.

So Microsoft said: "This campaign exhibits characteristics consistent with the financially motivated cybercrime operations employed by Storm-2561. In other words, ransomware. The malicious components are digitally signed" - this was interesting - "by 'Taiyuan Lihua Near Information Technology Co., Ltd.'" Okay. "The initial access vector," they said, "relies on abusing SEO to push malicious websites to the top of search results for queries such as 'Pulse VPN download' or 'Pulse Secure client.'" So you put that into Google, and the first link is this bad one.

They said: "But Microsoft has observed spoofing of various VPN software brands" - not just Pulse - "and has observed the GitHub link at the following two domains: vpn-fortinet[.]com and ivanti-vpn[.]org. Once the user lands on the malicious website and clicks to download the software" - and again, when you go to this malicious website, if you're not paying attention, if you don't know what the domain should be, it looks legit. I mean, it looks 100% like, oh, good, I just got to the home of Pulse VPN Secure. I'm going to download the secure client. Why wouldn't you?

They said: "Once the user lands on the malicious website and clicks to download the software, the malware is delivered through a ZIP download hosted at github.com/latestver/vpn/releases/download/vpn-client2/VPN-CLIENT.zip." Looking at that URL, it's like, okay, what's bad about that? Looks fine.

So they said: "When the user launches the malicious MSI masquerading as a legitimate Pulse Secure VPN installer embedded within the downloaded ZIP, the MSI file installs Pulse.exe along with malicious DLL files to a directory structure that closely resembles a real Pulse Secure installation path. It's %CommonFiles%\Pulse Secure. This installation path blends in with legitimate VPN software to appear trustworthy and avoid raising user suspicion.

"Alongside the primary application, the installer drops malicious DLLs, dwmapi.dll and inspector.dll, into the Pulse Secure directory. The dwmapi.dll file is an in-memory loader that drops and launches an embedded shellcode payload that loads and launches the inspector.dll file, a variant of the infostealer Hyrax. The Hyrax infostealer extracts URI and VPN sign-in credentials before exfiltrating them to attacker-controlled command-

and-control infrastructure," which is how the bad guys learn how to log into, like, your enterprise that you're intending to VPN to securely.

In other words, no one wants this software, any of this software, anywhere near any of their computers. It's all bad. Microsoft noted that the files were all signed. As I've been saying, no code these days can get off the ground any longer without being signed by someone. In this case, Microsoft also explains, writing: "The MSI file and the malicious DLLs are signed with a valid digital certificate, which is now revoked, from Taiyuan Lihua Near Information Technology Co., Ltd. This abuse of code signing," they wrote, "serves multiple purposes. It bypasses default Windows security warnings for untrusted code, might bypass application whitelisting policies that trust signed binaries, reduces security tool alerts focused on unsigned malware, and provides false legitimacy to the installation process."

They said: "Microsoft identified several other files signed with the same certificates. These files also masqueraded as VPN software."

Okay. So Microsoft described this as an "abuse of code signing." Okay. I suppose it's an abuse of the intent of code signing, but I'd be inclined to call it a failure of the code signing requirement to prevent the use of malicious software because, right, the bad guys didn't "abuse" code signing. They "used" code signing to abuse the process code signing was designed to prevent. Maybe I'm splitting hairs. But what we don't know and what Microsoft chose not to reveal here is whether this Taiyuan Lihua Near Information Technology Co., Ltd. is an authentic firm whose valid signing certificate somehow got loose, but that's difficult to understand because, as we know, code signing certificates now must reside in hardware, or whether the company was always a facade which bad guys used to obtain a valid code signing certificate.

Microsoft also chose not to reveal who signed their certificate. It would be interesting to know which Certificate Authority allowed themselves to be spoofed; and how and where, exactly, the required chain of enterprise existence proof failed. How did this happen? Hopefully somebody at Microsoft is pursuing this because this is exactly what's not supposed to happen. It's because of all these hoops that I had to go through everything I did in order to update my code signing certificate because you're not supposed to be able to do this. But here's a clear instance of very, very malicious software having a valid code signing certificate, and Microsoft mentions it a number of times in their write-up.

The only actionable takeaway we can have from this is the annoyingly diffuse imperative to remain ever vigilant. There are bad guys scattered all around the world focused upon taking advantage of our trust or any momentary lapse of our attention. All we really can be is as well informed and careful as possible.

While we're on the subject of bad guys taking advantage of the passion of the day, I wanted to note that Bitdefender, Kaspersky, and ThreatBook all recently posted independent examinations of the dramatic rise they all noted in malicious web pages offering instructions for installing AI agents like Claude and OpenClaw. I have a picture here in the show notes of what somebody would receive if they put into Google "download claude code." The first response that comes up is - it's from developers.squarespace.com.

Leo: Oh, my god. And it's a sponsored result.

Steve: Uh-huh. Exactly, Leo. And it says: "Install Claude Code - Claude Code Docs. Use the AI-powered sidebar. Generate snippets, refactor logic, and explore ideas in a clean interface."

Leo: Terrible.

Steve: So you put "download claude code" into Google, the first sponsored result that comes up is malicious.

Leo: It's not real.

Steve: Yes. Who would not trust this? Now, we know that we should not be getting Claude Code from developers.squarespace.com.

Leo: No.

Steve: But your typical user doesn't know that.

Leo: Yeah. I wonder how many people have been bit by this? That's awful.

Steve: Yes. Google labels it as a sponsored result, and the branding looks authentic. Users tend to trust it. So, you know, no more needs to be said other than to be careful and to always go to the original source of anything you obtain from the Internet. And again, the perfect instance, why is this being done? Because right now AI is the rage. And the bad guys are going to take advantage of what everybody wants.

Leo: Oh, and when you install this stuff, you really give it full access to everything.

Steve: Yeah. It's like...

Leo: So it's a great way to get malware on a system. Yeah, you should never google support numbers, either, for the same reason. Right?

Steve: Yup.

Leo: But everybody does.

Steve: What you should google, Leo...

Leo: Oh, our next sponsor?

Steve: Yeah, how did you know?

Leo: You're getting good at the segues, Steve Gibson. You'd better watch out, you're going to be a DJ soon. Yeah, Steve, we're going to the RSA conference next Tuesday. I'm very excited. Going to have a lot of fun. Sometime you have to come up for that. Have you ever been?

Steve: That's where I met Stina.

Leo: Oh, you met Stina coming down the escalator, that's right, that's right.

Steve: Yup.

Leo: Stina from YubiKey, yeah.

Steve: Yeah.

Leo: On we go with the show, sir.

Steve: Okay. So this is so - I don't know what this is.

Leo: It's so...

Steve: Is it frightening? Is it clever? Is it genius?

Leo: Is it a Movie of the Week?

Steve: Okay. So a security researcher by the name of Christopher Aziz of Bombadil Systems discovered a very, very clever, I say "new technique," I mean, it's always been there, but nobody thought to do this, that allows for the creation of malware-containing ZIP files that slide right past endpoint security tools, you know, Windows Defender and so forth, all the various AVs. In his testing, Christopher found that his simple, I mean, horrifyingly simple ZIP format hack would evade 98% of antivirus engines. I think one out of 55 caught it. The other 54 didn't.

When Chris packaged something, a piece of known, very well known malware in a regular ZIP file, it was almost universally detected by the AV engines at VirusTotal. But when he simply then tweaked the ZIP file's header to claim that its file contents had been directly stored rather than compressed, nearly all existing AV tools were fooled into believing that the contents was just gibberish. In other words, they didn't attempt to decompress the contents because the header said it wasn't compressed. It's almost too easy. Christopher put up a page on his GitHub account to draw attention to this obvious-in-retrospect vulnerability. It's at github.com/bombadil-systems/zombie-zip. He wrote, under How It Works, he said: "AV engines trust the ZIP Method field. When Method=0" - meaning that the file was stored, not compressed - "they scan the data as if it was raw uncompressed bytes.

"But the data is actually deflate compressed, which is ZIP's standard compression format, deflate compressed. So the scanner instead, believing it's not compressed, just sees it as compressed noise," he writes, "and finds no viral signatures. The CRC (Cyclic Redundancy Check) is set to the uncompressed payload's checksum, creating an additional mismatch that causes standard extraction tools (7-Zip, unzip, and WinRAR) to report errors or extract corrupted output." He said: "However, a purpose-built loader" - meaning a loader that knows what has been done - "that ignores the declared method and decompresses as deflate, recovers the payload perfectly."

He said: "The vulnerability is scanner evasion: security controls assert 'no malware present here' while malware is present and trivially recoverable by attacker tooling. As for the attack vector, this is not an end-user extraction vulnerability. This is a staged delivery/smuggling technique," meaning that you would, you know, malware, some script or something running, that's already running, would download this because of this simple hack. It would get into the system by passing all AV screening, and then it would know how to decompress this back into its fully malicious uncompressed state.

So he said: "The staged delivery/smuggling technique. First, a malicious payload packaged in what he calls a 'Zombie ZIP' with a modified header. The ZIP transits security boundaries (email gateways, network scanners, endpoint AV). Scanners read Method=0, scan compressed noise, and report 'all clean.' A purpose-built loader or dropper decompresses the payload programmatically. The payload materializes and executes." He says: "This is consistent with established malware delivery patterns, having previously been seen in ISO smuggling, HTML smuggling, CAB abuse and so forth, where attackers use custom loaders rather than consumer extraction tools."

So what was affected? He said: "50/51 AV engines on VirusTotal were fooled. Also fooled: Microsoft Defender, Avast, Bitdefender, ESET, Kaspersky, McAfee, Sophos, Trend Micro, and so forth." He said: "Only something known as Kingsoft detected it."

So anyway, this just goes to show how some of the simplest hacks, even after all of this time, can still be among the most effective. Sometimes there's just no need for something to get overly fancy. You know, some assumptions were made, and those assumptions can be abused to the benefit of the attackers.

Okay. So will AI write code for me? Our listeners, understandably curious because I've been so impressed with things like what Claude Code is doing for people, continue to express their curiosity over my own plans for AI coding. I mean, this is like, until this week, where I asked what the caption for that photo should be, it was probably the most often asked thing. Like, well, Steve, when are you going to start using AI? And I'm sure that this is partly due to my having previously made T-shirts for myself which say in white block letters on black, "Born to Code," and also due to my having been completely open-minded about a topic that has perhaps been more near and dear to me than anything else in my life.

I have many times, as we know, Leo, celebrated your successes and experiences embracing Claude Code. And I've shared many of our listeners' similar, stunned, mouth-left-hanging-open experiences when AI-produced code for them that made their computer do things they never imagined they'd be able to obtain for themselves. And in fact I'll be sharing another instance of that here after this. Obviously, something huge has happened. The question remains what that is, exactly.

As I settled down last Saturday morning to begin assembling today's podcast, I decided to log into X to see whether any of our listeners might have posted a candidate Picture of the Week. That's where I used to get them. The good news is everyone has largely switched over to using email, as I have. But you never know. So it was serendipitous that, when I happened to check, my feed contained several posts that were completely

on topic for the question of AI and coding. I don't know. Presumably, Elon's X system knows of my interest in the topic and therefore dropped those into my feed.

So the first post I want to share was written by a guy named Aakash Gupta, who posts frequently on Medium (aakashgupta.medium.com). If anyone's curious, I've got a link in the show notes with the spelling. His short bio says that he helps product managers, product leaders, and product aspirants to succeed. And that clearly is his focus. His posting quotes somebody who posted on the 13th, an Arvid Kahl, who just wrote: "Devs are acting like they didn't write slop code before AI." So it sounds like this guy is defending AI-produced code against people who are saying, you know, it's sloppy.

And so Aakash Gupta, who has a lot of experience with AI and product managers, says in his posting, he writes: "41% of all code shipped in 2025" - meaning last year - "was AI-generated or AI-assisted. The defect rate on that code is 1.7x higher than human-written code. And a randomized controlled trial found that experienced developers using AI tools were actually 19% slower than developers working without them. Devs," he says, "have always written slop. The entire software industry is built on infrastructure designed to catch slop before it ships. Code review, linting, type checking, CI/CD pipelines, staging environments. All of it assumes one thing: the person who wrote the code can walk you through what it does when the reviewer asks.

"That assumption" - that is, that the person who wrote the code understands it - he says: "That assumption held for 50 years. It broke in about 18 months." He said: "When 41% of your codebase was generated by a machine and approved by a human who skimmed it because the tests passed, the review process becomes theater. The reviewer is checking code neither of them wrote. The linter catches syntax, not intent. The tests verify behavior, not understanding.

"The old slop had an owner. Someone could explain why `temp_fix_v3_FINAL` existed, what edge case it handled, and what would break if you removed it. The new slop instead has an approver. Different relationship entirely."

He says: "Arvid's right" - the guy he was originally quoting. "Arvid's right that devs wrote bad code before AI. The part he's missing, the entire quality infrastructure of software engineering was designed around a world where the author and the debugger were the same person. That world ended last year, and nothing has replaced it yet."

So I just, I liked that just as a statement. You know. And his post captures aspects of my own discomfort with using AI to create code that I'm going to put my name on. So the answer to the question of whether AI will write code for me would be "not the AI we have today." Consider this. Even before this AI coding revolution arose, I should objectively have at least been using 'C.' Right?

Leo: Yeah.

Steve: But I'm so - right. C'mon. I'm so comfortable with assembly language, and I now have so much solid boilerplate written by me in assembly language through the years that, moment to moment, the path of least resistance is just to keep using assembly. When I face the possibility of using something to write code for me, I'm immediately brought up short, wondering how can I possibly know the code it creates is correct? The code I am writing is never for a lark. I'm not writing it as a hobby. I'm always writing production code that I and others will depend upon. Either it's server-side code running on GRC's servers, or code that will form a product that bears my name. In either case, the code needs, I need the code to be as correct as I am able to make it.

It's true that I have, we know, strong perfectionist tendencies. I know that's one of the reasons people listen to this podcast. I don't ever judge my work by whether it's "good enough." I don't have a "good enough." I know, you know, that I judge it by whether it's as good as I am capable of making it. That is my standard. Can it be better? So if I don't actually write the code I'm using, and offering for sale, how can I ever definitively make that judgment? If no one or nothing "sentient" and personally responsible creates it, if the code just magically appears, and if there are large swaths of code that is never carefully inspected by anyone, how can I ever have confidence in what the code does?

Sure, I know, test, test, test. I get that. That is, after all, you know, the model that many of our development testers know quite well. That's the development model that has evolved with the code that I currently author by hand is validated. But is the appearance of the code working, or the code no longer being seen to fail, an adequate replacement for someone actually writing the code for a purpose? I don't know. But I do know that the entire world is objectively going nuts over AI-written code. Perhaps the reason for this is that there is tremendous pressure within the larger code-creating universe to create more code with fewer human coders.

So perhaps it's the fact that I truly love writing code myself, and that I feel very little pressure to produce more code faster. Maybe that's, you know, why the balance for me, the scale hasn't tipped. I've talked about days past when my little company employed many more people, many of whom I was actually jealous of, since they were getting to do the work I wanted to be doing instead of just managing them doing that work. If that's the case, why would I want to have an AI producing code that I would then not have the joy of writing for myself? All of the foregoing suggests that the answer to that question, "When will Steve be using AI to author his code," the answer is at least not yet.

Leo: But we should point out, Steve, you're kind of a unicorn. You're kind of a rare [crosstalk].

Steve: I'm just talking, yes, the question is me. Me. I mean, our listeners have been asking, Steve, you're talking about Claude Code and how great it is. When are you going to use it? And I'm explaining why maybe never.

Leo: Yeah. And but nobody - how many people work like you? I mean, you're really an anomaly. In the past there were a lot of people, like Peter Norton and stuff, who wrote their own stuff and shipped it and so forth. But most code these days is written by large teams, you know, with all sorts of layers of review and architecting. I think for a lot of what is written today, AI makes perfect sense. Not for you because you're, you now, a [crosstalk].

Steve: No. And Leo, you notice I didn't say otherwise.

Leo: No. And you're right. I agree with you 100%.

Steve: Yeah. But I do [crosstalk].

Leo: Anybody who loves to write code should write code. If you love it, you should write it. Why not? That's not - but I have to say I'm not sure I fully agree with this

tweet because one of the things you're not going to see, frankly, if you have AI written code, is a temp, what is it, a temp fix underscore, because...

Steve: It won't get patched. It'll be created whole.

Leo: Code is so cheap that you refactor. You redo it. You don't do that kind of - that's what humans do. They apply a little spackle, a little bondo to the code. That's not what happens, or it shouldn't, with AI if it's being done right. I think really the experience people have with AI coding depends a lot on their own mindset and how they've gone about it. And how really you become, instead of the coder, you become kind of more like the product manager.

Steve: The manager.

Leo: Yeah.

Steve: Yes.

Leo: And a good product manager really thinks deeply about specs, is willing to throw out code and start over, I mean...

Steve: And Leo, remember I always say what we have today is not what we're going to have tomorrow.

Leo: It's going to very much change. That's the other thing. He says 41% of code written in 2025. Well, the thing that changed everything was November 24th, 2025.

Steve: Right. Right.

Leo: So when Opus 4.6 came out. So, or 4.5. So...

Steve: I have one more thing I want to share, but let's take a break. I'm looking at the clock, and now would be a good time.

Leo: I'm sorry to slow you down. I apologize.

Steve: And then I've got Uncle Bob Martin's post.

Leo: Oh, Uncle Bob. Good old Uncle Bob. He's quite the character.

Steve: Yep.

Leo: But a legend in the business, for sure.

Steve: Yup.

Leo: Yeah. Well, I look forward to that. That's coming up. You're watching Security Now! with the great Steve Gibson. You know, I'm really glad that there are people like you, Steve, that cherish, that are artists. You know, you wouldn't expect a machine to paint the Sistine ceiling. You're an artist. That's absolutely great. But I am not. So I appreciate having an AI to do some of my coding.

Steve: Well, and there's a whole different side of just getting the job done.

Leo: Sure.

Steve: Like, you know, and...

Leo: That's what most people are doing, frankly.

Steve: And I'm going to share a post from a listener after this that takes the exact reverse. This has changed his life.

Leo: Yeah, yeah. And then, you know, and I will say, and when I do coding puzzles like Advent of Code, I'm not - I have no interest in having AI do it.

Steve: No.

Leo: Because the whole point of it is me having the fun of writing a solution.

Steve: And in fact AI ruined that whole challenge.

Leo: It really did. It actually hasn't been a very good influence on it. He had to change everything. Steve?

Steve: Okay. So before we leave this topic - actually we have another note from a listener, too. But I wanted to share another X post that appeared in my feed directly underneath the previous one. It was written by someone who we obviously know, Leo. You are aware of Uncle Bob. He's got a Wikipedia page which, you know, was created to capture and describe his life's work. His given name is Robert Martin, although he goes by Uncle Bob Martin.

Wikipedia informs us that he's an American software engineer, instructor, and author, who is most recognized for promoting many software design principles - and by the way, he's a lover of Lisp - and for being an author and signatory of the influential Agile Manifesto. He's authored many books and magazine articles and was the editor-in-chief

of the C++ Report magazine and served as the first chairman of the Agile Alliance. Wikipedia says he joined the software industry at age 17, so like many of us it's been his life. He's credited with introducing the collection of object-oriented design principles that came to be known as Solid.

And Wikipedia mentions that he's authored many books. That's right, 13 books. Since I'm going to share his what I think is an interesting observation which really made sense about the current state of AI-generated code, I want to first clearly establish his bona fides.

So here are the titles of the 13 books he's authored across the past 30 years. And these are, you know, real books published by Prentice Hall, Cambridge University Press, Addison-Wesley Professional, and Pearson, with titles: "Designing Object-Oriented C++ Applications Using the Booch Method"; "More C++ Gems; Extreme Programming in Practice"; "Agile Software Development, Principles, Patterns, and Practices"; "UML for Java Programmers"; "Agile Principles, Patterns, And Practices in C#"; "Clean Code: A Handbook of Agile Software Craftsmanship"; "The Clean Coder: A Code Of Conduct for Professional Programmers" - he's all into clean - "Clean Architecture: A Craftsman's Guide to Software Structure and Design"; "Clean Agile: Back to Basics; Clean Craftsmanship: Disciplines, Standards, and Ethics"; "Functional Design: Principles, Patterns, and Practices"; "We, Programmers: A Chronicle of Coders from Ada to AI."

Okay. So here's what Uncle Bob Martin posted last Saturday morning. He wrote: "Two months ago, while working on my Empire game with AI, I had that quicksilver experience. When you push on a blob of mercury, it slips out in some random direction. Every time I added a new feature, some older feature would shift behavior. This was true even after I added unit tests and acceptance tests. The AI always took the path of least resistance on the current feature, and was willing to sacrifice older features. It would change tests, including acceptance tests, in order to get the latest feature done.

"Telling the AI not to do that was ineffective. AIs are stochastic, and so are any rules you feed them. Rules 'bias' their behavior, but do not absolutely constrain it. When I call them out on breaking rules, they apologize and swear they won't do it again; but they can't really make that promise. They are, in the end, liars and cheats. The solution is to massively overconstrain them, force them to write so many tests that changing a tested feature breaks many tests. They feel that force and retract the change. It's like peer pressure with a lot of peers.

"At the same time I reduce the chances for collateral damage by continuously forcing the AI to partition everything into small decoupled units. That way it's not easy to break one feature while implementing another. It also keeps the AI from getting confused by its own messes. The final goal is semantic stability in the face of continuous development. The things that worked before keep working as they were, while newer things get added.

"This is a continuous effort. Acceptance tests, unit tests, TDD. Crap analysis and mutation tests are run after a reasonable batch of changes and are tasked with reducing crap below 8, covering any untested behavior, and killing all surviving mutants. The size of the batch of changes is a judgment call. Too big, and the analysis and repairs take a long time. Too small, and the verification effort overwhelms the development effort."

And then he finishes with: "Side note: The mutation tests consume massive amounts of computer power. My cores are running full bore all the time, and that's even with differential mutation. There's something poetically just about this. The AIs require a massive amount of power to create. What they create for us takes a massive amount of computer power to keep stable."

So, okay. I think this has to do with the size of what he's trying to accomplish; right? Like, you know, he's building something big, and it's tending to get slippery, like liquid mercury where you push on it, and it slips away. And but from the start of our discussion of AI, I've been saying that I firmly believe AI will have a very bright future in coding. I still believe that's true, 100%. But not today's AI. Today's AI is still general purpose AI. It's like asking AI for that list of very high-quality random numbers. Doing that perfectly, which we know how to do, requires specialization, not generalization. This is every bit as true when it comes to writing code correctly. The laughable catastrophic mess Bob describes in his posting, you know, commonly referred to as attempting to herd cats, is not the way to write code. These four sentences from Bob's posting say it all.

He wrote: "AIs are stochastic, and so are any rules you feed them. Rules 'bias' their behavior, but do not absolutely constrain it." He says: "When I call them out on breaking rules, they apologize and swear they won't do it again; but they can't really make that promise. They are, in the end, liars and cheats."

I believe that in those four statements Uncle Bob exactly and perfectly captures the state of play today. But that's only today. I'm always, as I keep saying and noting, very careful to state that nothing we have or believe we have today regarding AI will hold tomorrow. And Leo, your November 28th date is a perfect example. On November 27th, we had one thing. On the 29th, the world changed. It's not at all done changing. You know, we're like in that first round of home computers that were interesting, and a lot of us got them, but they never got off the ground. It took another, you know, a bunch of more evolution and time for it to finally reach critical mass.

And so the way I think this will shake out is that someday we will have many differing forms of application-specific AI. I suspect that's where the answer lies. At least the practically economic answer. As I understand today's AI operation, having a single super-genius AI that contains all knowledge and does everything perfectly may be possible, but is incredibly wasteful, as in way too expensive to contain and operate, if all you want is high-quality code. Instead, employ the far more cost-effective services of a specialist code-generating AI whose model can be far smaller while also containing far more concentrated knowledge about code and only about code. It knows nothing about the works of Shakespeare. It just knows about code.

Leo: That's why our old model, prior to November 24th, 2025, was asking a question of a chatbot and then taking its code and pasting it in. We've gone way beyond that in a very, very brief period of time. You know, I think AI, especially AI coding, is kind of like the blind men and the elephant.

Steve: Yup.

Leo: You know that adage? Everybody is seeing a different part of it. And I think especially we can't use our notions of coding from prior times in modern times. It's just so different now, and everybody has a different take because everybody has a different experience. It's a [crosstalk].

Steve: I think your analogy is a good one.

Leo: It's a huge period of flux. And I think that's the only true thing. And really the best advice I think for anybody is just try it. Play with it. Get to know it. Give it a

tough problem. Read and learn. Everybody's talking about it. Not everybody's right. This is a lot of points of view about this.

Steve: And not everyone can be right when the target keeps moving.

Leo: It's very fluid.

Steve: I mean, we are - I cannot say enough. The world will be different again next year as regarding AI and code. There's just there's no question about it.

Leo: Yeah. We're in an interesting time. I mean, I guess the bottom line, we've talked about this before, and I think we both agree on it, is that what the job is, is taking human thoughts and ideas and translating them into computer. And what we're trying to make is a computer program that's very adept at that. The easy part is translating it to computer. The hard part is translating us. And but for somehow something happened that it got really good very rapidly at understanding what we're saying and putting it into action. But there's still, you know, miscommunications and gaps. It's very, well, we live in interesting times. Uncle Bob is very prolific, talks a lot about this. I actually saw this tweet. He's very active on X and talks a lot about this. It's very interesting.

Steve: So here is an example of AI on the flipside. Our listener Craig, the subject of his email was "Hard to Describe." He wrote: "First I'd like to say thanks for mentoring me throughout nearly my entire career. Now retired, I ran the IT department for a 50-employee DoD/DoE subcontractor. What I learned from you and implemented over the years made NIST 800-171 compliance easy. And I can proudly say that my company was never hacked. Oh, wait, aside from that warez kiddie who created a hidden FTP site on my public FTP server. Remember those days? LOL. But aside from that, never once was my network taken down. I had weekly security awareness training for my users almost always from your show. I was tight a decade before anyone was even thinking about security. Thank you. My entire career was hobbled by my poor coding skills. I never attended college for computers, just drinking and failing out."

Leo: I majored in that, too.

Steve: "I learned everything building PCs in those box shops in the late '90s. Netware Lite FTW, LOL. Computer Shopper for the win. I used to tell people I can code, but I can't develop. I could write a simple script, after hours of scouring Stack Exchange or Spiceworks to figure it out. The places I could have gone if I had properly trained as a developer. Now all those tools I wish I had over my 30 years of career are at my fingertips. The best analogy I can give is that I spent my career in 2D black-and-white, and all of a sudden I can see 3D in color and infrared and ultraviolet and x-ray." And he's talking about AI.

He said: "I now have an entire agent infrastructure team; CISO, Architect, Audit, Monitoring, Hardening, infrastructure, et cetera, managing my entire home lab. My kitchen module has an AI Chef running from local Ollama to help with the current recipe. I just got done having CISO build a 3D desktop for my platform inside of my Quest 2. It made downloading 20 years of Google account and then organizing it into my own system easy. It's working on building out a complete voice system around my house. It

can talk to my 3D printers. All of this is possible, and I just have to ask for it in natural language. My jaw is still on the ground. I hate to say it, Steve, but commercial software is dead. I don't need to buy what I can have my agents write. All I need are GPUs."

So anyway, I just thought that was a great snippet from one of our customers' whose life has been changed, thanks to AI.

Leo: That's nice. Really nice.

Steve: Yeah. Okay. Our last break, and then I'm going to share my 100% positive experience with CISA's Free Internet Scanning, and pose the question, why are we not all doing it, too?

Leo: Well, I'm going to try. I mean, I guess you're - I don't have multiple IPs. Well, I guess I do have two IP addresses. I guess, I don't know, I have one static and one theoretically changeable, although it never changes.

Steve: You have resources for TWiT; right? Or are they just all in the cloud?

Leo: No, no, it's all cloud, yeah.

Steve: All distributed stuff.

Leo: Yeah, it's all over the place.

Steve: So it would be a small enterprise that has a block of network space.

Leo: Russell could do it. I'll have Russell.

UNKNOWN: He's in Florida.

Leo: He's in Florida? Okay. He can do it from Florida. What do you mean, he doesn't work when he's in Florida? Let's do the final commercial, and then we'll get to the topic of the day, CISA.

Steve: Free Internet Scanning.

Leo: CISA has been decimated in the recent budget cuts, and I'm very nervous.

Steve: I'm glad they still have their bots running because, yeah.

Leo: Well, yeah, I mean, we are I'm sure the target of cyberwarfare. If not now, soon.

Steve: It's true they have lost a huge bunch of staff.

Leo: And they had what I consider to be a terrible administrator for a year. He's gone now. But that doesn't mean everything's better. It means there's no administrator. We're in an interesting time. Let's just say that. And that reminds me, Steve, I will not be here next week. Mikah will be doing the show.

Steve: Mikah.

Leo: Yeah.

Steve: Yup.

Leo: Going to miss Tuesday's shows so that I can go to the RSA Conference. Which I have never been to. So I'm really excited I get to go to this. It's going to be so much fun.

Steve: Cool. Cool.

Leo: I'm going to see a lot of sponsors. So that'll be neat, too. All right. Let's talk about CISA.

Steve: Okay. So last week I shared feedback from a listener who shared with us that his organization uses CISA's free Internet network scanner to keep an eye on his organization's network security exposure. He explained that when he first had CISA scan their network, what they found was quite bracing and brought their other IT people up short. And as I also noted, his sharing that with me raised my own curiosity about just who might qualify for CISA's periodic scanning.

It's formally known as CISA's Cyber Hygiene Service, and its page says: "Reduce the Risk of a Successful Cyber Attack. Cyber threats are not just possibilities but harsh realities, making proactive and comprehensive cybersecurity imperative for all critical infrastructure. Adversaries use known vulnerabilities and weaknesses to compromise the security of critical infrastructure and other organizations. CISA offers no-cost cybersecurity services to help organizations reduce their exposure to threats by taking a proactive approach to monitoring and mitigating attack vectors.

"By taking advantage of CISA's Cyber Hygiene services you can" -and we have some bullet points here - "significantly reduce risk. Organizations typically reduce their risk and exposure by 40% within the first 12 months. Most see improvements in the first 90 days. Avoid surprises because the services look for assets exposed to the Internet. They identify vulnerabilities that could otherwise go unmanaged. Sharpen your response by combining the vulnerability insights gained with existing threat detection and risk management efforts. Enrolled organizations can increase the accuracy and effectiveness

of response activities. This means fewer false alarms and less chance of real dangers slipping through the net.

"Broaden your security horizon. CISA's scanning is about more than pinpointing vulnerabilities. It's about expanding your organization's security boundaries. From basic asset awareness to daily alerts on urgent findings, you'll be in a better place to make risk-informed decisions." They said: "CISA's Cyber Hygiene services include: Vulnerability Scanning. This service continuously monitors and assesses Internet-accessible network assets (public, static IPv4 addresses) to evaluate their host and vulnerability status. In addition to weekly reports of all findings, you'll receive ad-hoc alerts about urgent findings, like potentially risky services and known exploited vulnerabilities.

"Web Application Scanning. This service deep-dives into publicly accessible web applications to uncover vulnerabilities and misconfigurations that attackers could exploit. This comprehensive evaluation includes, but is not limited to, the vulnerabilities listed in the OWASP Top Ten, which represent the most critical web application security risks. This service provides detailed reports monthly, as well as on-demand reports to help keep your applications secure."

Okay. So I've brought all this up again because my experiment to see whether GRC's little, decidedly non-governmental, non-tribal, 16-IP network block might qualify to receive CISA's automatic periodic background security scans and reporting. And it was a resounding and surprising success. Based upon my experience, I would hazard to imagine that a great many of our U.S.-based listeners who are in charge of their own small, medium, and even large enterprise networks, like the listener that put me on this, would be able to similarly qualify to receive this free service, much as I have. And if so, why wouldn't everyone wish to avail themselves of this entirely sane, zero-cost service offered by an agency of our federal government?

Now, I suppose I can imagine that it might make some listeners a bit queasy to invite Uncle Sam to scan and report on the state of their networks. But stop to consider that anything that might be discovered and reported is already public information. It's not as if we're making an exception for CISA, allowing them through our firewalls to rummage around inside our networks. That's not happening. They're on the outside attempting to look in, just like would-be attackers and hackers in Russia, North Korea, and China. The difference is that CISA is on our side with the goal of strengthening North American networks against attackers in Russia, North Korea, China and elsewhere. They email password-protected PDF reports that only its intended recipient is able to decrypt, open, and view. I don't see any possible downside, whereas I see potentially huge upside.

Okay. So what happened with GRC? That CISA cyber-hygiene-services page, it's at [CISA.gov/cyber-hygiene-services](https://www.cisa.gov/cyber-hygiene-services), I've got a link in the show notes, invites candidates to indicate their interest and open a dialog by sending an email to vulnerability@cisa.dhs.gov with just the subject "Requesting Cyber Hygiene Services." So I addressed an email, and I wrote simply: "To whom it may concern, I own a small commercial network which I would like to have scanned. Thank you. Steve." That was on the morning of Saturday, March 7th.

Leo: Did you say "Do you know who I am?"

Steve: No. Just to whom it may concern, I want to have my network scanned. Thanks. That was Saturday, March 7th, so nobody was working at CISA. I received a reply to that email first thing Monday morning, so immediately after the weekend, at 5:32 a.m. Pacific, so 8:32 a.m. in the East, where CISA is. That email response said: "Steve, thank

you for your interest in our Cyber Hygiene (CyHy) Vulnerability Scanning (VS)" - because you know they like abbreviations.

They said: "Thank you for your interest in our Cyber Hygiene Vulnerability Scanning service. Enrollment in our CyHy VS service must be done by a person in your organization who has ownership or authority over the IP addresses to be enrolled. This individual should hold a position such as Chief Information Officer, Chief Information Security Officer, or a similar official capacity.

"If you are in this role, please proceed to navigate CISA's Cyber Services: Cyber Hygiene Services, the beta version of our web-based enrollment system, to complete the following steps: First, create a Login.gov Account. Login.gov is our trusted partner for secure and private access to CISA's online services, including Cyber Hygiene. The Login.gov account must use the same organization (business) email that will be used to complete the remaining enrollment steps." And actually I don't think it does, but it didn't seem to matter.

"Second, return to CISA's Cyber Services: Cyber Hygiene Services Page. After logging in, you will now be redirected to the CISA Services Portal for ReadySetCyber. Use the navigation ribbon to go to Cyber Services > Enroll in Cyber Hygiene to return to the enrollment process. Third, complete Account Registration & Organization's Profile: Complete your organization's profile, enabling your organization to receive Cyber Hygiene and access other CISA services."

And then, finally: "Once you've completed the organization information page, you'll be redirected to a Thank You page. Select the Enroll Now option to continue the CyHy VS enrollment process. This step includes collection of the necessary information to enroll in the CyHy VS service and serves as the authorizing document allowing CISA to perform the CyHy VS service for your organization. For the IP Address validation process, you will need to input and successfully verify the formatting of your IP Address(s) before continuing to the next page. Multiple IP addresses must be separated by comma or line break. If there are errors with the formatting, the system will display a modal" - meaning dialog error, I guess - "noting how many errors.

"You will have the option to either go back and correct the errors or download a CSV file for editing if you have input numerous errors. If you have questions regarding your enrollment, please reach out to us at this email address. Best regards, Matt Leon, CISA Vulnerability Management Intake Team," blah blah blah.

So I went back to CISA and logged in at Login.gov, where I already had an account since I'm 70, soon to be 71, and I use Login.gov for managing Social Security, renewing my global entry certification, and driver's license. So I was then bounced back over to CISA, where I filled out a modest and not very intrusive questionnaire. Just, I mean, it wasn't a lot to tell them.

Around 10 minutes after completing that process I received another email with the subject "CISA Organizational Account Confirmation" and an invitation button to complete the sign-up process. I may have done something there, I don't recall. But either way, the email trail shows that 13 minutes later, after that one, I received a final email with the subject "Cyber Hygiene Vulnerability Scanning Acceptance Letter." I thought "Huh! That was easy."

Leo: Congratulations. You got in.

Steve: Yeah. The letter said: "Welcome To CISA's Cyber Hygiene (CyHy) Vulnerability Scanning (VS)." So these people really do love their abbreviations. The letter says: "Your CyHy VS Acceptance Letter has been processed, and a copy of the letter has been attached for your convenience. Your organization has been placed in queue for inclusion into the CISA CyHy VS service. Scanning will begin as soon as your request file is processed in alignment with your requested scan start date. If not otherwise specified, scanning begins immediately." The letter continues: "Please keep an eye out for traffic." And actually I did, my logs showed the scanning. "Keep an eye out for traffic from CyHy VS Scanning IPs which will signal to you that scanning has begun. You will receive your first CyHy VS report via email on the Tuesday following the initial scan, which is based on your requested scan start date. The CyHy VS report will come from reports@cyber.dhs.gov."

And then here's what was interesting. They said: "Overview of CISA CyHy VS methodology: Cyber Hygiene defines a host as having at least one port open and service. Scanning of hosts occurs continuously between each weekly report. Cyber Hygiene's scan prioritization is as follows." Okay, so we have "Addresses (IPv4) with no running services detected (dark space) are rescanned after at least 90 days." So if there's an IP that seems dead, nothing responds that they could find, it only checks every three months.

"Hosts with no vulnerabilities detected are rescanned every seven days. Hosts with low-severity vulnerabilities are rescanned every six days. Hosts with medium-severity vulnerabilities are rescanned every four days. Hosts with high-severity vulnerabilities are rescanned every 24 hours. Hosts with critical-severity vulnerabilities are rescanned every 12 hours. A single host may have multiple vulnerabilities of varying severity, which informs the frequency that a given host is scanned." Presumably, the highest severity vulnerability found defines how often it is rechecked.

And it finishes: "Need Assistance: If you need to make changes to the information submitted in the Acceptance Letter to include updated IPs to be scanned, or if you have any other questions pertaining to your CyHy VS service, please email us at vulnerability@cisa.dhs.gov."

Then, last Wednesday, the day after last week's podcast, when I didn't know if any of this was going to work, I received my first "CyHyVS" report. Now, I'll admit I was actually somewhat surprised to see that CISA had not found anything critical to complain about. Like I thought maybe. But that's not to say that CISA did not find anything. They did complain that GRC's web servers would still negotiate and accept SSL/TLS connections using old and deprecated 64-bit block ciphers, things like Triple-DES and Blowfish. Although not Blowfish. That was OpenSSL, but not in my case. That just is what people generally have, really old copies of OpenSSL, I'm sorry, OpenSSH can use Blowfish and should no longer.

So what caused my heart to initially skip a beat or two was that their report's headline was "Urgent Vulnerabilities Detected." And I thought, what? So obviously that commanded my attention. Their report enumerates their findings by vulnerability description, also whether it is known to be exploited, because as we know that CISA's KEV, right, K-E-V, known exploited vulnerabilities, that's one of their big deals, so they've got a column in the report for that, whether it's known to be exploited. Also whether ransomware is known to be exploiting it because obviously that drives an interest in that vulnerability and in being compromised by ransomware. There's a column for its severity, the host IP address and port where they found the vulnerability, and the data and time of its initial discovery.

In this case, all of GRC's web server IPs at the HTTPS port 443 share the vulnerability that CISA identifies as: "SSL Medium Strength Cipher Suites Supported (SWEET32)." That's the vulnerability's name. It is not, however, known to have ever been exploited.

So in the column of Known to Be Exploited it's not all the way down. The reason is that the SWEET32 vulnerability and attack is theoretical. It's called SWEET32 because the theoretical attack has a complexity of 2^{32} , meaning one in four billion, or 4.3 billion. The "Sweet" part of the name comes from the pun "Sweet16" because it's a birthday attack. You need to do a whole bunch of things, recording all of them, and then looking for any collision between any two. Thus the birthday attack.

The vulnerability has its own website at sweet32.info, which explains the nature of the attack, writing: "An important requirement for the attack is to send a large number of requests in the same TLS connection. Therefore, we need to find clients and servers that not only negotiate the use of Triple-DES, but also exchange a large number of HTTP requests during a single TLS connection without ever rekeying. This is possible using a persistent HTTP connection, as defined in HTTP/1.1 (Keep-Alive). On the client side, all browsers that we tested (Firefox, Chrome, Opera) will reuse a TLS connection as long as the server keeps it open."

Okay. So it says "a large number of requests during a single TLS connection." But exactly how large? In their own testing to recover a 16-byte authentication token, which might be an HTTP cookie, for example, a 16-character cookie, which would be two 64-bit encrypted blocks because this is an attack on 16-bit block encryption, they needed to keep a single TLS connection established for 18.6 hours, during which their client pounded on the server with a storm of continuous small HTTP requests, finally transferring 705GB of data in the process.

In short, at least for GRC, this is not a real problem. But that does not mean there's any way for me to defend GRC's now totally unnecessary support for this old and admittedly weaker than it needs to be Triple-DES cipher today. So I very much appreciate the reminder nudge from CISA, and I've already tweaked the cipher suite configurations of GRC's various web servers so that the next time they're rebooted, their support for that long-ago deprecated Triple-DES cipher suite will disappear. It hasn't been useful for a long time. It's only there because of inertia, but we know about inertia and security.

So that's the story of GRC's establishment of an ongoing, very valuable, free vulnerability scanning service courtesy of CISA. As I said, I cannot imagine why anyone listening to this podcast, who is responsible for anything more than a single IP home network or any sort of truly fixed pre-assigned IPs which are pointed to by DNS, would not wish to immediately avail themselves of CISA's free scanning service. You won't know what might surprise you until you do. And even if you find nothing, that would be super useful to know, too. If you do find something, it might be very important.

And the more that's going on within a complex networking environment involving multiple departments and overlapping responsibilities and people who've been terminated and blah blah, we don't know what equipment they left running, and different configurations, you know, the more of that there is, the more chance that something unsuspected may be there. So win, win, win, win, win.

Leo: That's my motto for the day. You won't know what might surprise you until you do.

Steve: That's why it's a surprise.

Leo: Surprise! Cyphase found the GitHub repo for all this stuff. So I don't know if that means it's open source. I don't know if you could take the GitHub repo and compile it and make it be - make it do its job.

Steve: Well, why not have it done for you?

Leo: Yeah. Well, why not, exactly.

Steve: Yeah.

Leo: But it's kind of cool that they've put this all online.

Steve: Yup.

Leo: 41 repositories on GitHub under CYHY.

Steve: Nice.

Leo: So you can at least see what they're doing. That's pretty cool because there's a lot of shell scripts. Shell and Python.

Steve: Yeah, it's running on their infrastructure. And, you know, I did get - so I got that one report that had that one vulnerability. Then a couple days later I got like a 34-page beautiful PDF that had charts and graphs, and it was tracking vulnerabilities, and like bar graphs and how long has this been around. I mean, it is really valuable. And the listener who put me onto this noted that this replaced for their insurance provider a service that they'd been paying \$6,000 a year for.

Leo: Right, right.

Steve: And that was an annual scan. So something could be bad for a year before it would get seen.

Leo: I can see some enterprising person taking all this code, getting it running, and making their own commercial version of this. It's open source, though.

Steve: Yeah.

Leo: CISA has its own - I love this - GitHub repository. CISA.gov.

Steve: Commit today, secure tomorrow.

Leo: Oh, I like it. Oh, that's what they said, it's their motto. Yeah, commit today, secure tomorrow. I've got another motto now. I've got two mottos from the last

section of this show. That's pretty impressive. Steve, you are pretty impressive. We appreciate everything you do.

Steve: You won't know what might surprise you until you do.

Leo: Until you do.

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/>