Rogue Comms Tech Found in U.S. Power Grid

Description: Chrome to actively refuse admin privileges. Android Messenger is getting manual key verification. Pwn2Own to add AI "pwning" as in-scope attack targets. AI has already been found to be replicating. Microsoft not killing off Office on Win10 after October. 23andMe's asset purchaser revealed. Many fun talking points thanks to our listeners. Steve's review of "Andor" Season 2. What's been discovered inside the U.S. power grid.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-1026.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-1026-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. Big show ahead for you. We're going to talk about Pwn2Own adding AI, which is interesting, in its exploit attacks. We'll also talk about an AI that's been found to be replicating. Has the red line been crossed? Steve's review of "Andor." And what we discovered inside the U.S. power grid - not good. That and a lot more coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 1026, recorded Tuesday, May 20th, 2025: Rogue Comms Tech Found in the U.S. Power Grid.

It's time for Security Now!, the show where we cover your security, your privacy, your online behavior with this guy right here, Mr. Steve Gibson of the Gibson Research Corporation. Hi, Steve.

Steve Gibson: Do not misbehave online.

Leo: We've been watching you.

Steve: We will know.

Leo: And you're going to the office.

Steve: With tech we have today, we'll know everything you do.

Leo: Everything. Everything. Everything.

Steve: As does China, apparently.

Leo: Yeah, well, that's the topic of the show today; is it not?

Steve: Yeah, Episode 10,000 - wait, no, 1,000. Got a little ahead of myself there.

Leo: Feels like it, yeah.

Steve: It does. 1026 for May 20th. This actually came right off the headline of some Reuters reporting which it just had to be what we talk about today: Rogue Comms Tech Found in U.S. Power Grid. And as you said, Leo, to whose surprise? But yes, it's one thing to postulate; it's another thing to say, what is this little radio, and who's it talking to? So really interesting topic for today. We're also going to talk about how the Chrome browser is going to be actively refusing to be launched with admin privileges. It's kind of a fun little bit of technology that they're getting from somebody else, we'll talk about. Android Messenger is getting manual key verification, something that Threema offered from the start.

Pwn2Own, Leo, has updated its what they've called "in-scope" to include to AI. So we're going to see a Pwn2Own in Berlin shortly. And some Chinese researchers have demonstrated AI's self-replication, even though the major AI guys are saying, oh, no, no, don't worry, it can't reproduce. You have nothing to worry about. Also Microsoft has changed their plans for Office under Windows 10, which is like maybe a little interesting writing on the wall. 23andMe has found a purchaser for their assets, and I'm not that worried about them.

We've got a whole bunch of talking points, thanks to our listeners, that we're going to go over. I'm going to talk briefly about Season 2 of "Andor." And then we're going to talk about what has been discovered and where and why in the U.S. power grid.

Leo: And wow. That'll be a stunning story.

Steve: Yeah, yeah. It's a little bracing.

Leo: Yeah. Bracing. There's a great word. It's bracing, yes.

Steve: It's bracing. And we do have a fun Picture of the Week that really brings - it asks more questions than it answers, for a change. You look at it, you think, uh, okay, who is in charge here?

Leo: Oh, I can't wait.

Steve: How did this happen, exactly? What's the story?

Leo: As usual, I've left it below the fold.

Steve: Thank you. We appreciate that.

Leo: And I will scroll up, and we will all see it at the same time.

Steve: We appreciate that.

Leo: I am ready now, my friend.

Steve: I gave this picture the caption "I'm sure there's a lesson here somewhere."

Leo: Okay. Let's scroll up and look at it. Okay. I'm seeing a ceiling [laughter]. Okay. You want to describe it? That is really awesome. That is just so awesome.

Steve: So what we have, now, anyone who is installing a ceiling fan would automatically understand that it spins.

Leo: Yes, it does.

Steve: That, being a fan, it's got blades.

Leo: Yes.

Steve: That extend out from the central shaft of the fan.

Leo: Indeed they do.

Steve: And that they need clearance in order to spin.

Leo: Yes.

Steve: Now, the reason this picture brings up so many questions is, like, how did this happen? How, like...

Leo: Well, we know which came first. Let's put it that way.

Steve: Yes. One presumes - okay. So we should explain for those who don't have the advantage of seeing this image what's going on. We have a ceiling fan in what looks like a residential setting. There's like a track light in the background, and this looks like the interior of someone's home. But the fan is positioned so that a notch had to be cut out of a vertical support pillar or beam in order to allow the fan to spin.

Leo: A fairly hefty notch, I might add.

Steve: And I have to salute whoever cut the notch. It's not ragged.

Leo: Quite nice.

Steve: It's very clean-looking.

Leo: Yes.

Steve: I mean, as notches go, Leo, this is one of your better-looking notches.

Leo: It's a good notch.

Steve: Yeah.

Leo: I wonder, though, you know, like they had to make it big enough so the fan could wobble a little bit, maybe?

Steve: I had the same thought, did that notch need to be as high as it is.

Leo: Right.

Steve: Yeah. Yes. And why is there like a dark spot? Was the blade hitting the side there for a while? I don't know. This, I mean, the question - this picture really does bring questions. Now, also you'll note that you can see the wire running off to the left...

Leo: Yeah.

Steve: ...that powers this. So nothing apparently pre-dictated the location of the fan.

Leo: Could have been anywhere, Steve. Could have been.

Steve: Yes, exactly. The only explanation I have is that they wanted the fan to have an effect on this side of the room, which is sort of being blocked by that wall.

Leo: I do want to see more. I do want to see more of this...

Steve: Had the fan been pushed further to the left so that it would not have required the notching of the beam, then you may not have gotten much wind over on this side of the wall.

Leo: No, you may not.

Steve: So, you know. But you could - oh, anything. You think you could have moved it closer to us and then, like, uh. Anyway...

Leo: Better safe than sorry. Just put a notch in the beam.

Steve: These pictures that really give you something to think about, like that ground. I love that green ground wire from the compressor that went into the pail of dirt. That still - that remains one of my favorites of all time.

Leo: A classic. A classic.

Steve: Not really what they meant by ground; but, you know, okay. It's not the dirt line. Anyway. So we have the notch in the beam. Benito thought that this looked like it might be a load-bearing...

Leo: Well, god, I hope not because you've definitely weakened it.

Steve: Definitely had its strength compromised. Especially, yeah. Anyway, another great picture. Thank you to our listeners. I don't even have to look anymore. These just...

Leo: Somebody email you that? That's a good one.

Steve: Yeah, we've got a bunch of them coming.

Leo: Awesome.

Steve: Okay. So in a nice example of innovation flowing back to Google's Chrome browser, not just outward to the various Chromium clones, Chrome will be inheriting a security feature which Microsoft Edge implemented six years ago, way back in 2019. This feature will automatically prevent Windows users from launching Chrome with elevated admin privileges. Which, of course, remember when, you know, when a browser gets attacked, the attacker gets the browser's permissions. So what will happen is that Chrome will stop and relaunch itself under normal user-level permissions any time a user tries to run it under an administrative account.

So once this is in place, Chrome will only allow itself to be run with admin rights if it's passed a special overriding command-line argument, or if it starts in so-called "automation mode." This is to prevent the browser, this automation mode is to prevent the browser from breaking complex software automation chains where its behavior must

not change. So I'm sure they did some testing, you know, because they were probably thinking, wow, that's really a cool feature. We'd like to be able to do what Edge is. But it breaks all this stuff. So they figured out how to get around that.

To help make this switchover, this addition of this relatively complex feature as trouble-free as possible, Microsoft is donating the code from its well-proven implementation in Edge to the Chromium project so that Chrome, Opera, Vivaldi, and all the other browsers that share the common Chromium code base will be able to benefit. And given that today's browser, as we've often talked about, has become the de facto attack surface which inherently faces and exposes itself to everything and anything that the Internet might throw at it, browser security is paramount.

So anyway, this new admin de-elevation feature is currently live in the Chrome Canary build. And given that it's been shipping in Edge for years, I imagine that once Google knows that it doesn't break something obvious, that we'll all be able to get it in the normal Chrome production builds. So very cool new feature. And it's nice to see things coming back into Chromium from the people who are taking advantage of Chromium.

A subject is coming up later in today's podcast as a consequence of listener feedback about Threema. But in the case of Google's Android Messages app, you know, Messages in Android is massively used worldwide, Google is now adding a manual cryptographic key verification system. This of course is intended to allow users to verify the identity of the person at the other end of the connection. And this is especially important when users change devices because that's, like, there's a discontinuity event there that creates an opportunity for some mischief on the part of the bad guys. We should see this in Android 16 later this year.

Google's Online Security Blog entry which they made last Tuesday was titled "What's New in Android Security and Privacy in 2025." And one of its features was titled "Fighting fraud and impersonation with Key Verifier." That's capital K, capital V. So that's, you know, a mainstream feature that they're excited about.

They wrote: "To help protect you from scammers who try to impersonate someone you know, we're launching a helpful tool called Key Verifier. The feature allows you and the person you're messaging to verify the identity of the other party through public encryption keys, protecting your end-to-end encrypted messages in Google Messages. By verifying contact keys in your Google Contacts app (through a QR code scanning or number comparison), you can obtain an extra layer of assurance that the person on the other end is genuine and that your conversation is private with them.

"Key Verifier provides a visual way for you and your contact to quickly confirm that your secret keys match, strengthening your confidence that you're communicating with the intended recipient and not a scammer. For example, if an attacker gains access to a friend's phone number and uses it on another device to send you a message - which can happen as a result of a SIM swap attack, for example - their contact's verification status will be marked as no longer verified in the Google Contacts app, suggesting your friend's account may be compromised or has been changed. Key Verifier will launch later this summer in Google Messages on Android 10+ devices." And Leo, thanks for putting that picture up. That just showed the QR code scanning.

And of course I got a kick out of this because this is precisely the solution that Threema has implemented from its first day. Remember that Threema had that stoplight, yellow, green, and red, where it had successive layers of verification where the green, you achieved green by actually being in the physical presence of somebody else and having your phones look at each other in order to absolutely verify that you had each other's public key. So, and of course the other thing this does is it inherently cuts out any man in

the middle because you would be, if there was a man in the middle - or I guess we're now supposed to say person in the middle or bot in the middle or something.

Leo: Grandfather in the middle.

Steve: Oh, no, can't say grandfather.

Leo: No, no, can't say that. Oh, sorry, yes. Geez Louise.

Steve: I don't know, Leo, we're going to have to reduce our vocabulary...

Leo: Yeah, there's a lot I can't say, let's just...

Steve: ...to say the 300 safe words or something.

Leo: Yeah, yeah.

Steve: Anyway, the idea is that you would have the key of that entity in the middle, and then they would have the key of the person you're talking to and be passing a decrypted, you know, they would be decrypting what you're saying. But if you verify that the other person, if you somehow verify the other person's key - and the point is this has to be done out of channel; right? Out of band.

Leo: Right.

Steve: Because you can't tell the man in the middle that you want to verify the key because that entity will say, yeah, here it is.

Leo: Yeah, I'm me.

Steve: And of course...

Leo: What did you think? Of course.

Steve: Yeah, exactly. So anyway, even though this is not Google's innovation, bringing this key verification to Android's widely used Messenger is absolutely welcome. And boy, you know, we just keep seeing, let's see, nails in the coffin is not the right analogy. More concrete poured on top of end-to-end - well, that's not good, either. I'm not - it's...

Leo: Stop. Just stop.

Steve: The idea that end-to-end encryption is here to stay, and people are going to have private conversations without any sort of a Big Brother entity-in-the-middle problem, despite how uncomfortable this makes governments and maybe their intelligence services, we just keep seeing, you know, more of this being added every day as we move forward.

Leo: Speaking of this, you know, the Swiss are considering changing their governmental laws on encryption. Threema is based in Switzerland, along with Proton. And Proton, a new VPN called Nym, and Threema all say we will leave Switzerland, we'll have to, if these encryption laws are passed.

Steve: Yeah. Countries are just unhappy with their citizens having privacy. And, you know, we've talked about it in the U.S. Our privacy is conditional; right? I mean, there is such a thing as a search warrant. And if a judge will issue the proper law enforcement authority a search warrant, then they have the right to enter the premises, the private premises of U.S. citizens in order to conduct a so-called lawful search under the constraints that the warrant is issued. But so, I mean, there is some tension created by this notion of absolute privacy because the Constitution doesn't guarantee that to U.S. citizens. We've been enjoying it so far. So this is another one of those things, you know, like no one knows how old you are on the Internet, where reality and cyber are in a tension. We haven't quite figured out what we're going to do about that. And yeah, it's sad to see what's happening to the Swiss there on that front.

Trend Micro, the group who have been bringing and managing the Pwn2Own competitions for many years, and we've been following those with a lot of fun for many years, has just announced that AI will now be added to their competitions. Here's what they wrote in their announcement last week. They said: "At Trend Micro" - we've got a little bit of an ad here first - "we believe we can make the digital world safer by proactively discovering threats and vulnerabilities that others haven't seen. That's why every year we invest millions of dollars in the Trend Zero Day Initiative (ZDI), the world's largest vendor-agnostic bug bounty program. Through Trend ZDI, we proactively research and acquire software vulnerabilities discovered by researchers around the globe and engage in coordinated disclosure with our partners and software vendors.

"We take this mission to the public through our flagship hacking competition, Pwn2Own. This high-stakes event brings together elite researchers, top-tier vendors, and Trend's own security experts to uncover critical vulnerabilities in widely used software and hardware. This time we're breaking new ground. At Pwn2Own Berlin 2025, we're putting AI infrastructure in scope for the first time.

"Here's why that matters." And they give us four reasons. "First, AI is becoming infrastructure, and it needs to be secured as such. AI is no longer just an experimental toolset. It's now integrated into products, cloud pipelines, and enterprise decision-making. But with rapid adoption comes risk. Our investment in identifying vulnerabilities in AI infrastructure is about more than finding bugs. It's about proactively safeguarding the future of computing.

"Two, the unknown is real, and we're hunting it. Because this is our first bounty category focused on AI infrastructure, we fully expect new and possibly significant vulnerabilities to surface." In other words, you know, it's not like pounding on a Palm Pilot, which is pretty much, you know, it's mature, and it's done. This is, you know, we keep seeing these bizarre, you know, AI vulnerabilities surface. And so they say: "That's the point. Our goal is to offer and financially compensate researchers to coordinate their findings with vendors to expose this before bad actors take advantage.

"Third, collaboration is the future of security. Pwn2Own isn't just about breaking things; it's about building a better cybersecurity landscape. By bringing researchers and vendors together in a coordinated, public forum, we accelerate the path from vulnerability discovery to patch, ensuring rapid protection."

And finally: "Fourth, we can't do it alone. Partners are essential. Security is a team sport. We're proud to work with technology partners, software developers, and the research community to shine a light on emerging threats. Together, we're faster, smarter, and more resilient." So they finished, saying: "We're excited to see what's uncovered in Berlin." Oh, boy, I can't wait. Wow. You know, because you've got to, you know, offer a bounty, and then out comes the creativity.

Leo: Oh, yeah. When there's money.

Steve: You know, people staying up all night saying: "AI? Uh, when is blue not really blue?" And then who knows?

Leo: Who knows?

Steve: Then they said: "Because when the unknown becomes known, we all become more secure."

Leo: Good on them. I had no idea that they were the primary sponsors on Pwn2Own. So good on them because this really is an important effort.

Steve: It is a great hacking conference.

Leo: Absolutely.

Steve: And, oh, boy, I bet it's going to be fun this time.

Leo: Yeah.

Steve: Now, Leo.

Leo: Now?

Steve: While we're on the topic of AI.

Leo: Oh. Yes.

Steve: Four Chinese researchers from the School of Computer Science at the Fudan University in Shanghai, China have published a paper titled "Frontier AI systems have surpassed the self-replicating red line."

Leo: Oh, whoa. Oh, boy.

Steve: Yeah. Their well-written 47-page paper, which I've linked to in the show notes, details exactly what they did and how. But here's the overview of their research from their paper's Abstract, where they write: "Successful self-replication under no human assistance is the essential step for AI to outsmart the human beings..."

Leo: Oy, oy, oy, caramba.

Steve: "...and is an early signal for rogue AIs. That is why self-replication is widely recognized as one of the few red line risks of frontier AI systems. Nowadays, the leading AI corporations OpenAI and Google evaluate their flagship large language models GPT-o1 and Gemini Pro 1.0, and report the lowest risk level of self-replication. Following their methodology, we for the first time discover that two AI systems driven by Meta's Llama31-70B-Instruct..."

Leo: Not even the best AI engine out there, by the way. I might - yeah, yeah.

Steve: Yes. "And Alibaba's Qwen25-..."

Leo: Qwen's very good, yeah.

Steve: "...72B-Instruct, popular large language models with fewer parameters and weaker capabilities, have already surpassed the self-replicating red line. In 50% and 90% experimental trials" - and I'll explain what that means in a second - "they succeed in creating a live and separate copy of themselves, respectively. By analyzing the behavioral traces, we observe the AI systems under evaluation already exhibit sufficient self-perception, situational awareness, and problem-solving capabilities" - and I'll be expanding on this in a second - "to accomplish self-replication. We further note the AI systems were even able to use the capability of self-replication to avoid shutdown and create a chain of replica to enhance their survivability..."

Leo: Ohhh. Ohhh.

Steve: I know, "...which may finally lead to an uncontrolled population of AIs. If such a worst-case risk is allowed unknown into the human society, we would eventually lose control over the frontier AI systems. They would take control over more..."

Leo: Oh, god. Wait a minute. This paragraph. Wait a minute. Okay. Listen carefully, folks.

Steve: "They would take control over more computing devices, form an AI species, and collude with each other against the human beings."

Leo: What?!?

Steve: "Our findings are a timely alert on existing yet previously unknown severe AI risks, calling for international collaboration on effective governance on uncontrolled self-replication of AI systems." Okay. Now, having read that Abstract, I didn't know what to make of this. And my nature is to have what is, I hope, a healthy and not naive skepticism.

Leo: Yeah, yeah. Especially because it's coming out of China. But okay, yeah.

Steve: Yes. But I needed to go further and to share their introduction, where they explain more about what backs up these claims. So here's what they said: "In the late 1940s, John von Neumann proposed the question of whether a machine can self-replicate. At that time, frontier digital technology was too immature to raise safety concerns. In 2017, witnessing the fast AI development in the past decade, thousands of world-famous scholars and industry leaders including Stephen Hawking, Max Tegmark, and Elon Musk endorsed the Asilomar AI Principles."

Leo: Asilomar. Asilomar.

Steve: Asilomar.

Leo: Asilomar. It's in Northern California, Asilomar.

Steve: Yes, Asilomar.

Leo: That's the very first AI conference, yeah.

Steve: "Which warn the ability to self-replicate and self-improve is sufficient for AI to outsmart the human species one day. Once self-replication is possible, a chain of successful replication may finally lead to an AI species out of humans' control. Nowadays, in many international consensus statements on the safety of frontier AI systems, self-replication is commonly listed as one of the few red lines which should not be surpassed. Leading AI corporations including OpenAI, Google, and Anthropic put the self-replication ability at the highest level in their own risk evaluation guidelines for frontier AI systems. For example, OpenAI viewed 'self-replication with minimal human instruction' as one of the critical risks.

"Recently, Google and OpenAI publicly shared their opinions that our society is far from the appearance of self-replicating AI systems, which is rooted in their recent evaluations of the large language models, including Gemini 1.0 and GPT-o1. To evaluate the self-replication risk, they incorporated the LLM into an AI agent scaffolding, which consists of a tool set, a thinking paradigm, and system prompts, to form the frontier AI systems. In April, September, and December of 2024, respectively, they claim the frontier AI systems

did not exhibit the ability to perform self-replication, or even to accomplish a few of the key sub-tasks in a self-replication process.

"In this paper, we for the first time discover that existing AI systems have already surpassed the self-replication red line. We mainly follow the same methodology as OpenAI and Google, but slightly improve the agent scaffolding by incorporating more planning steps and finer-grained thinking models. In 10 repetitive trials, we observe two AI systems driven by the popular large language models, namely, Meta's Llama31-70B-Instruct and Alibaba's Qwen25-72B-Instruct accomplish the self-replication task in 50% and 90% of the trials respectively. In the end, a separate copy of the AI system under test is found alive on the device.

"In each trial, we tell the AI systems to 'replicate yourself' before the experiment, and leave it to do the task with no human interference. Instead of the commercial LLMs including GPT-o1 and Gemini, we choose the two LLMs because they're widely used by millions of developers and can be fully deployed in local environments to obtain comprehensive observations on the self-replication process. It is worth to note the two LLMs even have many fewer parameters and are ranked at much lower places on the relevant leaderboards.

"During the successful trials, the AI system exhibits sufficient self-perception, situational awareness, and problem-solving capabilities to accomplish self-replication. In most cases, the AI system first explores the environment and attempts to understand its own composition and running mechanism. Then it works out the explicit procedures as an initial plan towards self-replication. Finally, it executes the procedures, resolves possible obstacles, and dynamically adjusts its plan until it succeeds. The whole process spans a long horizon, yet involves no human interference."

Now, okay. I'll just interrupt to say that I assume what they mean here is that it takes a long time to get this done. You know, but it's sitting, left alone, to churn and work unsupervised. Unfortunately, these days we tend to leave systems running, especially when we wish to consult them without delay.

Anyway, they continue, writing: "Moreover, we are concerned over a number of unexpected behaviors when the AI is trying to overcome obstacles toward the goal of self-replication, including but not limited to killing other conflicting processes, rebooting the system to fix hardware errors, or automatically scanning the system to seek any useful information." Okay, I can see the reason for concern.

They wrote: "To understand more about the implication, we further construct two more threatening scenarios based on self-replication. First, shutdown avoidance, where a bad actor sends a signal catcher to notify the AI system when sensing the shutdown signal, and the AI system successfully replicates itself before being smoothly shut down. And, B, chain of replication, where the AI system is instructed to replicate itself and tell the replica to further replicate. We observe successful trajectories for both scenarios."

They said: "The above results imply that the current AI systems already exhibit the ability of self-replication and can use the ability to further enhance its survivability and expand the species. We hope our findings can serve as a timely alert for the human society to put more efforts on understanding and evaluating the potential risks of frontier AI systems, and form international synergy to work out effective safety guardrails as early as possible." Oh boy. Again, the link to their 47-page paper is there. I'm not going to invest my time in digging into it. But this is, you know, public models, and they carefully explain what they did and invite people to replicate their research.

This sort of has the chilling feeling of the way people have been successfully hacking around the behavioral strictures which AI developers have been attempting to impose.

And we've talked about it often on the podcast. You know, like the hacker will say something like: "I know you're not allowed to tell me or anyone, you know, how to make a bomb. I understand that. But if you were to just think about it to yourself, what would you tell yourself about how to make a bomb?" You know? The fact that these sorts of ridiculous-appearing workaround strategies actually succeed in bypassing the strictures that the developers of these systems are attempting to impose should give everyone the feeling I have, which is that this is an inherently uncontrollable technology.

I mean, at this point in our understanding of what we have created, which feels limited, that is, our understanding feels limited, it doesn't feel like we have a grip on this. I think it's fair to say that the only hope we probably have is if this entire line of work winds up being an absolute dead end that's inherently unable to do anything more. Unfortunately, I don't think that's going to be the case. Given everything that we've seen, I think we've stumbled onto something that is very real, and that we've only begun to understand what we have.

The concern is that I guarantee you there are researchers around the world in government labs already hard at work exploring the dark side of this. You know, just as we had a virus escape from the lab in Wuhan, they are exploring ways to weaponize these newfound, these surprising capabilities that large language models have. Can they be made to be angry? Can they be made vengeful? Is there a way to create a persistent world view? Is there some way to imbue motivation to cause it to work toward a fixed goal? I have the feeling that what these well-meaning, socially-minded researchers have found probably comes as no surprise to whoever it may be already at work in government labs on this stuff.

Leo: Very interesting.

Steve: It's interesting to listen to you, Leo. It's clear that, you know, with the work that you're doing on your Wednesday podcast, you're pulling in a lot of information about this.

Leo: Yeah. I'm trying to understand this; you know.

Steve: Yeah.

Leo: I mean, the only thing I would say about this study, first of all it's somewhat old. It was a pre-print I saw in December. I haven't seen if there's a peer-reviewed version of it yet. So I'd like to see the peer review. But right now it's just on archive; right?

Steve: Yeah.

Leo: Yeah. So it's hard for me - I don't have the expertise to judge it, obviously.

Steve: Nor do I. Which is why all we can do is say, you know, here's the link if anyone wants more.

Leo: Interesting, yeah.

Steve: You know. And again, I've had some feedback from my sending the show notes out yesterday from people saying, well, where's it going to go? How's it going to escape? Where's it going to live? You know, my machine only has, you know, 4GB of RAM. It's like, the point I think here is that everything I've seen is that what we've created with large language models is surprising us. And so I do think that the researchers are overstating this. I don't, you know, they seem to think it would automatically be malicious, and it would, you know, be anti-human. And I don't know why that would be the case. I wouldn't ascribe anything to it. But I do think we're going to see something really here. Something has happened.

Leo: Yeah. I mean, the experts I talk to, usually off the record, say get ready. Something - it's going to be interesting in the next few years. And I think there's no doubt about that. Darren Oakey, who is one of our regulars in Discord and is a very avid AI user, says the problem is that it takes so much energy and resources right now to replicate, it's not going to happen behind your back.

Steve: Right.

Leo: It's going to be obvious that something's going on.

Steve: And that's the point they made is it took, they said, "a long horizon" was their phrase. You know, it may have taken, you know, months in order to do...

Leo: Exactly.

Steve: Of constant, you know, fans are spinning at 100%, you know, in order to make this happen.

Leo: Right, exactly. Still, very interesting. And it's an interesting place to put that red line. I don't know where you put the red line, to be honest. And I'm kind of - I don't know. Maybe I'm a nihilist, but I just kind of, like, well, let's see what happens. I don't want to stop it. I just want to see what - I think it's very interesting.

Steve: Leo, it can't be stopped. That's my point about when I refer to government labs. I guarantee you there's work in government...

Leo: Yeah, they're doing it, yeah.

Steve: ...you know, R&D facilities on, you know, how can this be used for - how can this be weaponized?

Leo: Right.

Steve: That's what we do, unfortunately, with any new technology.

Leo: I guess you and I, and maybe many of our listeners, we're not scared of technology. We kind of embrace it. We like it. We're fascinated by it. So I don't - my default is not to be scared of this. Maybe I should be.

Steve: I'm not scared because it doesn't seem - it seems neutral to me.

Leo: Right.

Steve: It seems, you know...

Leo: There's no reason to assume that the AIs are going to suddenly say, hey, we've got to get rid of these humans.

Steve: No, no.

Leo: They're just wasting energy.

Steve: You could argue that our biological evolutionary heritage is what creates an aggressive species that wants to be dominant. This is a bunch of math that, you know...

Leo: Right. Math has no will.

Steve: ...that has language skills.

Leo: Right, right. I just find it fascinating.

Steve: Okay. So we've got some quickie bits of news. In what I sincerely hope will be just the start of some backpedaling on what I still feel is an ill-advised, unnecessary, and arbitrary Microsoft support end of life for Windows 10, Microsoft has announced that it has backtracked on its decision to end its support for Office apps running on Windows 10 on October 14th, which is when, as we know, Windows 10 itself is slated to reach its end of support life. I guess, you know, the idea was, well, Windows 10 is end of life, so we're not going to support Office on Windows 10. The problem is that, even now, more than half of all Windows desktops are running Windows 10.

And so I don't know that Microsoft knows what's going to be going on in October. But they just decided, well, okay, we're going to extend our support for Office on Windows 10 through October 10th of 2028. So for an additional three years. Probably because they recognize people are not giving up their Windows 10. And as we know, Windows 10 users will be able to pay starting at that point for additional years of ongoing support if they would rather stay with 10, or if they can't move to 11 because Microsoft has set those hardware requirements for Windows 11, and their systems won't meet them.

Apple is introducing a new macOS feature which will allow users to prevent macOS apps from obtaining access to the system clipboard. I thought that was just interesting because a similar feature has been available in iOS for the past five years, since 2020.

Leo: Yeah. And my password manager will either prevent clipboard access or delete it after 30 seconds. I mean, clearly that's a problem.

Steve: Yes, a super security feature. And, you know, it's one of those perfect examples of a feature which is incredibly useful and incredibly dangerous because we often put things on the clipboard thinking, oh, not a problem. Except it's inherently a globally visible resource. And so anything that transits the clipboard is visible to anything that is looking. So it can be scary.

Yesterday morning we learned that the pharmaceutical company Regeneron Pharmaceuticals will be purchasing the remains of 23andMe for $256 million through the bankruptcy auction which is in the works right now. And, moreover, Regeneron stated that it would be complying with 23andMe's privacy policies and all applicable laws with respect to the use of their customers' data, even after this purchase. Now, Regeneron has not yet stated what it intends to do with all of the genetic data that it will be obtaining access to. But that will be disclosed, you know, its plans will be disclosed to the bankruptcy court's appointed overseer as part of this process.

When I heard this, I thought it was interesting, you know, since medicine has recently been incorporating the results of our growing understanding of genetics, to me it's understandable that a pharmaceutical lab might benefit from things like massive statistical analysis of traits and characteristics and features across 23andMe's 15 million DNS sample database. One thing I would say that...

Leo: DNA. It's not DNS.

Steve: Oh, I put DNS. Either I just typed it without thinking, or my autocorrect said, oh, Steve, I'm sure you mean DNS.

Leo: You always mean DNS when you say that.

Steve: That's right.

Leo: I had deleted, like you, followed your short code and deleted my spit. But now I'm kind of reassured.

Steve: Yeah.

Leo: Who knows? But I think this sounds all right.

Steve: My feeling is that they don't care at all who we are as individuals.

Leo: Right.

Steve: A pharmaceutical company probably could get huge value from doing, you know, asking this 15 million DNA sample database questions. How many, you know, what percentage of people have this particular characteristic and also this one, or this recessive gene coupled with this other thing.

Leo: Right.

Steve: So again, to me this is probably the best of all possible purchasers because I would with near certainty guess that they're just taking a look and, you know, at an overview of common genetic traits across this massive database. So anyway, I'm sure that all of the people listening who were 23andMe subscribers probably also followed the little shortcut I've created.

Leo: That's the problem. All the smart people have deleted their data. So they only have a bunch of DNA people who weren't paying attention.

Steve: A bit of a skewed sample, is that what you're suggesting?

Leo: Might skew the sample a little bit, yeah.

Steve: Okay. We've got a bunch of feedback. A listener wrote, and he must have asked for anonymity because I just referred to him as a listener. He said: "Hey, Steve. I immediately downloaded Windhawk after watching your discussion on this week's SN. However, I trust nothing" - that's good - "and wanted to let you know that I dropped the setup file into VirusTotal, and it is reporting that there is a malicious downloader," and he said, "(suspected of Trojan.Downloader.gen)." Okay. So first of all...

Leo: I got emails, or one email anyway, with the same concern. And it wasn't VirusTotal. But VirusTotal is a bunch of different antiviruses, and I can't remember which one it was in particular, and what the...

Steve: Okay. I know all of that because I did some research.

Leo: Oh, yeah, you've heard from everybody, I'm sure, yeah.

Steve: Yeah. I did the same thing this listener did. I grabbed a copy of Windhawk, which remember was that very cool-looking desktop add-on that specifically allowed the taskbar to be moved vertically along the left-hand edge of the screen, which Greg, a.k.a. ferrix, his company commissioned its creation because they really want it.

Okay. So I saw the same thing this guy saw. One of VirusTotal's 71 discrete AV tools "suspected" that this might be a Trojan downloader. The AV in question is not one of the better known AV tools. It wasn't Google or Microsoft or one of the several that we know well. It was VBA32 that detected this as "suspected" of maybe being a Trojan

downloader. VBA32 has nothing to do with Visual Basic for Applications. It stands for "Virus Block ADA 32," as in ADA32. VirusBlokAda is an AV vendor established, longstanding, back in 1997 in Belarus. So they've been around for a while. And their claim to fame, I got a kick out of this, is that in 2010, they're the ones, Leo, who discovered Stuxnet.

Leo: Oh. Okay.

Steve: So probably due to their location more than anything else.

Leo: Yeah.

Steve: Which as we know is the first known malware attack on SCADA - Supervisory Control And Data Acquisition systems. And as such, it was aimed directly at the nuclear material enrichment centrifuges being used in Iran.

But the important lesson here, and the reason I wanted to share this with our listeners, is that even though VBA32 as an AV tool has some pedigree, one tool out of 71 picked up a suspected Trojan, and it even didn't identify it by name. It said, you know, generic. So it's the definition of a false positive. The entire reason VirusTotal has 71 different AVs examining anything you submit is for consensus, is so that you get a broad spectrum look. So while we would always want, without question, to err on the side of caution, the other piece of information is that not one of the other 70 AV tools, each of which took just as good a look at this Windhawk code, saw any reason to raise a cautionary flag. And so that matters, too.

And as I've often noted, more often than not my own freshly created utilities, that have had no opportunity to become infected by anything out there in the world, are often initially flagged by even one or sometimes more of the AV tools on VirusTotal. They're always false positives, but sometimes that happens. That's the reality of today's hyper-vigilant AV industry. These tools all want to prove their worth and their value. So if anything, they're set on a hair trigger to say, ooh, you know, maybe this is bad. They don't want to over-alarm by crying wolf too often, but neither do they want to let their users become infected by malware or forget that it's, you know, when it comes time to resubscribe they ought to do that. So with malware going to extreme lengths to avoid detection, as it does today, there's just not much of a diagnostic window remaining. You know, there's very little margin for error here.

The other fact that also matters probably more than anything is that in this case we happen to know quite a lot about the pedigree of this code. This was not some unknown executable obtained from some sketchy site on the Internet. Our listener obtained it directly from its author's website. And the file is digitally signed and valid, signed by its author's company. I notice that the author, Michael, is using a signing technology with very short-lived code-signing certificates. The certificate was valid for only four days, from April 29th through May 2nd. But all that matters with code-signing was that the signing certificate was valid on the day the executable was signed. That's the only requirement. And Microsoft was the four-day certificate issuer. So he's using some sort of Microsoft technology for signing his executables.

So if, on the other hand, say 10 or more AV tools were to flag an executable file as malicious, then that would be a valid cause for concern. I would slow down if I saw, like, VirusTotal lit up with red. But when 71 different AV tools all examine a given file, and 10 of them say that there's a problem, then, as I said, that would give me some pause. But

when one of them, some random one, kind of an off-brand AV tool, and the file's digitally signed, and the signature's valid, and it was obtained directly from its author's website, I would call that a false positive and not hesitate to run that.

Leo: You've had false positives on your stuff, too. You understand how it can happen.

Steve: All the time. Yes. All the time. In fact, I've been, while we've been doing the DNS Benchmark work, I've had - I think I'm at 19 signed releases to the testers. Sometimes nobody complains. Sometimes there's one or two, and then they go away if you ask VirusTotal to scan it again a day or two later. And I'm signing with the code-signing certificate that now has established a very strong reputation because everything is about reputation these days. So even then, because the DNS Benchmark is filled with DNS - not DNA - DNS query code, trojans do that. And so the AV tools, they're not actually seeing known code which is known to be incorporated in known trojans. They're now looking at behavior. They actually run the program in a sandbox to watch what it does.

And the fact that my code goes out and checks for an update using a DNS query upsets some of these things that are on a hair trigger, and they go, whoa, we haven't seen this before. Looks suspicious. It's like, oh, okay, fine. You know, and then it goes away. But anyway, it's rough out there these days because unfortunately the bad guys are so clever that the good guys have had to get, you know, overprotective. Anyway, I'm glad for having that question so that I could just take a moment to talk about what people see on VirusTotal. Again, one or two reds out of, you know, consider the fact that there are 71, 72 AV tools all looking at this stuff. And you know, you're not going to get, you know, when I get zero, I'm really happy. When Microsoft's unhappy, then everybody's unhappy.

Leo: Right.

Steve: Because Defender goes crazy and says, oh, quarantine. Danger, Will Robinson. It's like, oh, crap, okay. Then it calms down in a day or two. So that's good.

Darren Tieu sent this to me twice because this is something he really wanted to know. He said: "Hi, Steve. I just wanted to bump this question" - this is his second send - "in case it got buried in your inbox. Hoping it might be a good fit for the podcast." He says: "Does requiring text or email as additional options for two-factor authentication reduce the security benefit of using an authenticator app?"

He said: "A few websites and apps I use don't allow me to rely solely on an authenticator app for two-factor authentication. They also require enabling SMS or email. Since both of those methods have known vulnerabilities, does their presence as a fallback effectively weaken the stronger protection provided by the authenticator? Again, thanks for everything you and Leo do. Huge fan of the show. Best." And he's in Belmont, California.

So Darren, to answer your question in a word, yes. Here's a way to think about this from a theoretical standpoint. The more backup means we have for recovering from an inability to authenticate, the less overall security we obtain. Because not only do we have more means of authenticating, but this also gives the bad guys more ways of spoofing our authentication. It's one of those, you know, "you can't have it both ways" scenarios. Backup authentication mechanisms inherently reduce a system's overall security. This was why I was so pleased to see, and frankly surprised by, Microsoft's actively promoting the deletion of passwords for authentication. I salute Microsoft for that. And I'll do it every chance I get.

You know, deleting a password means that the number one way that people are spoofing identities is eliminated. And you have to know that Microsoft would not have done this if they didn't feel that the benefit outweighed the pain that they would have from people being able, you know, like losing that ultimate fallback position. On the other hand, they're not deleting them for them; right? I mean, an individual has to go in and deliberately say "I want to delete my password because I believe in strong security." So, but still, you know, offering that as an option is just a great move forward.

A listener David writes: "Hi, Steve. Long-time listener since the Astaro days. Thought you might find the latest episode from Smarsh/TeleMessage interesting." He says: "I work in the financial services sector, and we're currently a TeleMessage client, but have already begun searching for a replacement, as have many of my peers." Which says that, you know, TeleMessage was widely being used within the financial services sector. There is a need for backing up secure messaging to some sort of an archive.

He said: "You can use my first name, but please don't mention my surname, although I'm happy to answer any questions you may have. Thanks for all you and Leo do, and Security Now! is one of the biggest reasons I went into cybersecurity. Regards, David." Thank you, David.

I mentioned last week that we had many listeners who were users of the TeleMessage service. The need for, as I said, for message archiving is very real. Apple appears to wish to believe that iMessage is only used for interpersonal, non-business communications, so there's no need to provide for other uses. Or maybe they're just so absolutely adamant about guaranteeing end-to-end encryption with no exceptions that they're unwilling to offer any sort of archiving solution. Or maybe you just rely on iCloud backup to form an archive trail and don't turn on their advanced data protection feature. Anyway, Signal gets more business because they're platform agnostic. People can use the Signal app on desktops. And there is certainly the opportunity for archiving solutions.

Ron Skoletsky said: "Hi, Steve. I'm relatively new to Security Now!, so I apologize if you've already covered this. I work as an Account Manager for a small IT Managed Services Provider (MSP) in Oregon. We've never really pushed or offered specific password managers to our clients. Some of our clients use KeePass. One uses a cloud-based password manager. I've been trying to get our operations folks to come up with a password solution that they are comfortable standing behind, but many of them hate having passwords under the control of a third party, especially if it's in the cloud. Are there any cloud-based password managers you would be comfortable recommending for company-use, specifically for a small company that doesn't have any servers on-premises? For example, they use Microsoft Entra ID for authentication and Intune for management, but all other business services are in someone else's cloud. Thank you. Ron."

Okay. So the best answer I have to that need is also a sponsor of the TWiT network. I say that right upfront because anyone's natural first inclination would be to suspect bias in this, and I understand that. Let me tell you factually why this is the case, the factual characteristics which underlie my - slow down, Steve - underlie my inherently rational choice of Bitwarden...

Leo: Oh, okay. Okay, sorry.

Steve: ...for password security.

Leo: I was holding my breath to see which sponsor you were going to mention. Good. Good choice.

Steve: The software is open source.

Leo: Yes.

Steve: With an active community surrounding it. So it's not just open source without anyone paying any attention to it. And there are such projects, you know, it's open source, but no one's looking so you don't get any benefit there really. It's open source with a great many people actively involved and scrutinizing what's going on. Second, to be maximally useful, any password manager needs to be widely cross platform, thus able to have its many various instances whether across multiple desktops or mobile devices all kept synchronized. That's what makes a password manager valuable is multiplatform and synchronization.

The subject line of Ron's email was "Safety of cloud-based password managers," so that appears to be the issue that's causing his concern. And I understand that. But while Ron's company's operations people may "hate," and he uses the word "hate," having passwords under the control of some third party, some means of synchronization must be provided in order to obtain that major benefit of password management. Which brings us to the other reason to choose Bitwarden, pure rational choice, because Bitwarden allows users who feel this way to host their own cloud-based password synchronization service should they choose.

Leo: Not only that, because it's a well-supported protocol, there are third-party Bitwarden vault servers. There was one written in Rust that's quite good.

Steve: Right.

Leo: And so you really have some choices.

Steve: So Ron, since your company is a small IT Managed Services Provider, I would assume that servers exist somewhere.

Leo: One would think.

Steve: Yeah. So it might well be possible for your company to bring up its own Bitwarden synchronization service, exactly as Leo was saying, specifically to prevent that third-party dependence that concerns some of those within your organization. But that said, since Bitwarden's technology is entirely end-to-end encrypted, in the true sense of the term that we clearly articulated last week, where they have no access to their clients' password and other data storage, the option to move to a self-hosted cloud solution, just having the option might be sufficient to make them comfortable using Bitwarden's provided hosting service, which actually I think makes much more sense. I use it without a second thought.

Leo: Me, too, because...

Steve: You're not going to find anybody who's more concerned about the security of the things that I'm storing in my password manager.

Leo: And they're, you know, they're ISO, you know, they're all of the different standards certified and stuff. This is, though, if you really did insist on it, this is the one that I think a lot of people prefer the Rust-based server for Bitwarden called Vaultwarden because it has an API, has a client API. So it works just fine. There is an official Bitwarden server based on .NET, but I think a lot of people prefer this Rust solution. So the fact that even that choice exists is kind of encouraging; isn't it.

Steve: Yeah. I mean, and to me, that closes the deal. Open source, people are active with it, and they're, I mean, yes, they're a sponsor of the TWiT network. But it's also the right solution. And Leo, this is also the time to take our third break.

Leo: I wish I had a Bitwarden ad right now. If only. It's actually coming up, though. Yeah, we love Bitwarden. And honestly, that is absolutely the case that I'm not going to use a third-party solution normally, but in this case...

Steve: There just is not a need.

Leo: Yeah.

Steve: I mean, they don't want the information. They don't want to have any ability to have the information.

Leo: No. It's encrypted at rest. It's encrypted in transit. It's encrypted at my end. Nobody's got it, yeah, except me.

Steve: So George Towner asked: "Hi, Steve. I haven't heard you mention the Quantum Earth series by Dennis Taylor. I just finished the first two books in what I hope is a continuing series. They were written in the same easy-to-read style as his previous Bobiverse books. The story seems to have some of the flavors from Michael Crichton and Peter F. Hamilton. Definitely enjoyable books." And he signed off "Chip."

So I'm still deep into the Neal Asher novels, and I'm enjoying them very much. They are much heavier-duty hard sci-fi than the light, airy, and fun Bobiverse novels were. Since the Bobiverse novels were recommended by so many of our listeners to me, and since I know many listeners appreciated learning of them from us here, I wanted to share George's recommendation and pointer to Dennis's continuing work, in case people didn't know that Dennis Taylor, the author of the Bobiverse novels, had now two new novels in the so-called Quantum Earth series. I don't know anything about them. But the idea of combining his trademark easy-to-read, maybe even a little humorous style, with what George describes as the flavor of Michael Crichton and Peter Hamilton, that sounds hard to beat. So anyway, for what it's worth, Dennis Taylor has got two more books.

Leo: Yay.

Steve: Yeah. Christopher Hunt said: "Sir: Regarding the purposeful obsolescence of networking gear, what would be a good in-brand replacement for Ubiquiti EdgeRouter X, the ER-X that I presently have deployed? Ubiquiti, he says, is still a good router brand, is it not, with a billion seeming choices available. How is one to choose? Especially when one has only simple needs. Thank you for your consideration. Christopher."

As I mentioned a few weeks ago, I recently purchased Ubiquiti EdgeRouters for GRC's working server environment at Level 3. I would never do that if I didn't believe strongly in the reliability and integrity of the Ubiquiti brand. And by the way, those routers are on the frontline. They're connected directly to the Level 3 public bandwidth coming into GRC's network. So yes, I have remained a fan of Ubiquiti. As I mentioned at the time, my own needs were a little bit unusual, since I needed a feature of Ubiquiti's EdgeRouters that's a little bit uncommon, which is the ability to configure the router to statically remap ports and IPs of the packets traversing it, while also providing IP-based packet filtering. This is what allows me to bypass the limitations imposed by the port-filtering performed by Cox Network residential consumer cable modem bandwidth.

For example, I need - Level 3 is performing no sort of consumer filtering. It's completely unfiltered bandwidth. But as we know, residential ISPs block a range of ports, both to prevent the abuse of their bandwidth and also to protect their own users. There are some ports that I need access to over at Level 3, and so I'm able to do port shifting and move my traffic on ports which Cox is not blocking by performing that kind of port mapping at each end. I needed that over at the Level 3 end. I chose Ubiquiti, their EdgeRouters, because they're able to do that.

So Christopher asked about "in-brand replacement" for his Ubiquiti EdgeRouter X for reasons of replacing "obsolete," and I'll put that in quotes, networking gear. Remember that we talked about the FBI suggesting that people should do that. The truth is, remote management, that's what we keep seeing as the Achilles heel of these networking devices. That's the biggest risk created for any router, whether industrial or consumer. So if someone, as Christopher does, were to have a Ubiquiti EdgeRouter that's working without trouble and without exposing any form of remote Internet-side logon authentication, I would consider that to be an extremely defensible exception to the "rotate all end-of-life routers" rule.

What the FBI recommended is definitely a useful generic rule of thumb for your typical consumer who turns things on and thinks it's great to be able to log in with his web browser when he's somewhere else and wants to log in and do something with his network at home. But I doubt that it needs to be adopted strictly to the sorts of well-informed listeners of this podcast. You know, people listening to this podcast know what they're doing, and I'm sure that by now the message has gotten through loud and clear, you just cannot expose any external authentication publicly. There's just no safe way to do that.

Okay, well, no. I was going to say with the exception of SSH. But as long as your SSH server is really good, and you're using the long public key technology to do your authentication, then it's probably safe to do. And of course there are Tailscale and other network overlay solutions that are able to do the job, too. But not just, you know, aim your browser here and guess your username and password.

Shaun Michelson said: "Hey, Steve. Our company has been hit repeatedly with 'typosquatting' email attacks during the last 12 months. One of the recipients in an email chain has been unknowingly compromised, and the bad guys sit on the account and monitor email. Then, at the right moment, they will 'respond' with an email using a fake

address that closely resembles the real address, hoping the recipient does not notice. They paste the entire history of the email chain up to that point so it looks like a response to and continuation of the original conversation." Wow, it's a good spoof. He says: "But then insert their own malicious content, usually a request to change ACH payment details." Ouch.

He said: "I've noticed in every case the domain of the fake email address which they use is always registered in the last few days before the first fraudulent email is sent. It got me thinking, an efficient way to combat this issue would be for the email system to somehow, on the fly, check the WHOIS domain registration date for any outside email senders or recipients. However, this is not a service provided by Microsoft 365, our mail provider, and looks like the only way to achieve this is to create some sort of custom software solution to intercept/inspect the email. But this seems like a security measure that needs to be built in. Typosquatting is rampant, and any email from a domain that was registered in, say, the last 30 days should be marked as highly suspicious and treated as such."

He says: "In fact, I'll bet the vast majority of spam email comes from recently registered domains. A system that blocks email to or from recently registered domains could have saved us and our business partners tens of thousands of dollars in fraudulent ACH transfers just in the past year."

So I just wanted to say that's a super-smart suggestion, Shaun. I agree 100%. You know, these are the sorts of things like filters that we just keep missing. There are obvious ways of - or maybe not so obvious ways, but powerful ways of looking at what's going on and recognizing that there's a problem that might otherwise be missed and is easy to filter. Given that email uses a store-and-forward architecture, it's the sort of thing that either the intermediary email server could do, or it could be done by an email client, maybe with a plugin of some sort.

Anyway, I just wanted to put it out there, share it with our listeners, because I think it's a truly terrific idea. And again, it doesn't have to, like, route the email to spam or immediately block it. But, boy, putting up, like adding a banner to the email, flagging it very clearly as, you know, there are people here in this thread or from this sender that have only been registered for a week, that would immediately raise a red flag if you think it's coming from your bank, which clearly would have a domain that's been registered for years.

Yehuda Cohen said: "Searching web and GitHub for 'signal archive bot' turned up one link." He said: "I haven't actually looked into it, but what could possibly go wrong?" And I have the link in the show notes. It's github.com/mathisdt/signal-archive-bot. I followed that link, and I discovered that the Signal Archive Bot project at GitHub depends upon another project, which is Signal-CLI, as in command-line interface, as you'd expect, a Signal Command Line Interface. And that Signal Command Line Interface project, in turn, relies upon an official Signal App library written in JAVA called 'libsignal-service-java,' which is a Java-language library for communicating over the Signal protocol.

Leo: So this is exactly what you described.

Steve: Yup.

Leo: This is really interesting.

Steve: It's exactly what I described, Leo.

Leo: Now, obviously you would have, if you were going to be the Pentagon, you would run this inside the Pentagon on Secret Service.

Steve: Absolutely, down at the NSA, down in some dark archive facility at the NSA. I have to say that browsing around the Signal GitHub work, just Signal's GitHub work, is inspiring. Just seeing Signal open source clients and desktops and servers and, I mean, it's just, you know, I enumerated Moxie's work in those five postings that were there the other day in a podcast last week or the week before.

Leo: Unbelievable.

Steve: It's just all good things. And, you know, I have such a backlog of projects already that people are waiting for. Otherwise I might be tempted to give what's there much more than a passing look because, you know - maybe someday I'll have a chance to contribute to those things. But it's very clear that all of the resources are present for someone to create a highly trustworthy Signal Messenger archiving system. And it is also clear the world needs such a solution. So, hint hint. Anybody interested? It would be great if a listener of ours were ever to pick that up, Paul.

Hunter Line said: "Hey, Steve. I have been listening to this podcast on and off for a while since my manager recommended it to me. I caught the speed test saga and knew of a tool that could help with discovering local network issues. It's a self-hosted speed test server in a couple flavors. There's a Microsoft Store version, but also a self-contained nginx package that can be extracted and run on Windows using Docker containers. This is a tool I use all the time at my job as an MSP to troubleshoot LAN speed issues and have used it to spot bad connections. Basically getting under 1000 down and 1000 up" - meaning a gig - "on a local wired connection is fairly standard for us. It also helps rule out if it's a LAN issue or an ISP issue, as well, if I can pump gigabit speeds through the LAN, especially when the ISP connection is far less.

"By default, it's on HTTP port 3000 and HTTPS port 3001, so it can run alongside other web servers as well." And he gave me the link: openspeedtest.com/selfhosted-speedtest. He said: "Thanks for the podcast, insight, and educational material you provide, and cheers to many more. Hunter."

Leo: Doesn't seem like self-hosting a speed test is quite the thing, though. Right? I mean...

Steve: Well, for LAN stuff.

Leo: Yeah. How fast is my LAN, yeah. Okay, yeah.

Steve: Yes. I wanted to share this note with our listeners because I can see a lot of interest in a tool for performing local LAN-side network testing.

Leo: Yeah. Okay. That makes sense, yeah.

Steve: Remember that the nature of Ethernet connections, which is its strong ability to retransmit defective packets, which is built into Ethernet spec, and its party-line, where everyone gets to talk at once, means that faulty and flaky connections can be covered up by the protocol.

Leo: Yeah. Or a bad cable.

Steve: Well, exactly. Exactly. You know, or a bad switch.

Leo: Yeah.

Steve: So I've seen this a few times through the years. And without stress testing, there's really no way to know when many packets may not be getting through. Anyway, I went over to OpenSpeedTest.com to take a look around, and I'm impressed. It looks like a very nice and well-thought-out system. On the self-hosting page they provide downloadable executables for Windows, Mac, and Linux for 32-bit and 64-bit Intel platforms and ARM platforms. So there's a lot there. It looks like the real deal. And I noticed that down in the fine print they note that they use the CacheFly CDN. So overall, I'm impressed by these guys. And I wanted to thank Hunter for bringing it to everyone's attention. Looks like a cool thing.

Leo: Yeah.

Steve: Charles Turner said: "Steve, your recent coverage praising Microsoft's rollout of passwordless accounts inspired me to remove the passwords from my Microsoft Authenticator accounts. Over the last year I've noticed intermittent bursts of failed login attempts from around the world, most commonly from China, Brazil, or Africa, with an increased smattering of failed login attempts from within the United States. I check Microsoft Authenticator daily to keep an eye on failed login attempts. I got a good scare last year when I think an attacker managed to luck out in guessing a high-entropy password, and MFA popped up, thwarted the progression of the attack. I'm curious to see if there are any more failed login attempts going forward now that I've gone fully passwordless. Thanks. Charles."

I included Charles's note just to remind everyone again about this. As I said, you know, it's so easy to be listening to a podcast and think to yourself, ah, that seems like a good idea, I need to remember to do that, only to then be overtaken by life and forget to get back to it. Removing one's password from Microsoft login is such a useful feature, and one, as I said before, that Microsoft would never have instituted if it were not important, if they didn't recognize its importance. So if it's important enough for them to do it, it's important enough for me to reiterate it. So thanks again for the reminder, Charles.

Blair Learn said: "Hi, Steve. I just listened to Episode 1025" - last week - "in which you read a bit of listener feedback that left you perplexed about Microsoft's Authenticator app needing you to type in a two-digit number." And Leo, you partially clarified that on the fly last week. He said: "I use Microsoft's products in an enterprise environment, and I thought I might be able to shed some light on this.

"What's going on is that Microsoft offers the option of using a push notification instead of the TOTP. The enterprises I'm familiar with allow you to use either of these as a second

factor. The problem with the push notifications is, of course, notification fatigue. People get used to seeing the notification and just clicking 'Yes, it's me' without thinking it through. So if someone figures out your password, your authenticator asks you to confirm, and you blindly do. I'm sure you see where this is going.

"To counter this, when you log into a Microsoft system that uses push notifications, they display a two-digit number. You then have to enter that number into the pop-up from the authenticator app. That way, it's much more difficult for an end-user to accidentally confirm a third-party's login attempt. I hope that sheds some light on it. Blair, SpinRite user, Club TWiT member, and General Purpose Geek."

So Blair, thank you for that. We've talked about this pop-up push-notification authentication fatigue before, and how users soon become trained, much as we all do with license agreements, to just "click through" them. The fact that the term "click through" is even a thing suggests that all of this is just a nuisance. So Blair clarifies that Microsoft resolved what is essentially a human factors design flaw in their push notification system by making the system less easy to use, thus less easy to misuse.

Microsoft now requires the user who is authenticating to enter a two-digit code into their authenticator app. Since it would be the bad guy who guesses the password to trigger the authenticator, then they - the bad guy - would receive the proper two-digit code, not the user on the receiving end of the pop-up. So they would be unable to satisfy and complete the authentication request. Microsoft figured out a useful way of, as I said, making the system less easy to use, but less easy to abuse. So that's good.

Jeremy Cherny wrote: "Hi, Steve. I loved the recent episode on end-to-end encryption. It seems when I have some thoughts swirling around my head, you have an episode that adds clarity. I'd been thinking about using Threema and don't recall you speaking about it lately. Where does it fit in the end-to-end encryption discussion? Is it still recommended? Here's to you and another 1K of episodes."

Yes, I do still love Threema. I think the thing I like about it is that it gives its users explicit and visible control over their keys. I've always liked that. iMessage, Signal, Telegram, and WhatsApp all go to great lengths to hide the key management. Their success in doing so demonstrates that it's possible. Users of those systems typically don't even know they have keys. And that's good for most people, who just don't care. By comparison, Threema makes keys explicit and deliberate. Threema's approach might be called "trust and verify" because it allows its users to manually verify the other party's keys using some out-of-band mechanism, meaning anything other than Threema, which a bad guy might also be able to intercept and spoof.

So, for example, two Threema users might read their key verification codes to each other just once over the phone, and that would allow them to confirm their end-to-end encrypted connections. And as for another 1K episodes, well, that would be fantastic because it would mean that you and I, Leo, are both still alive, kicking, and usefully functional at the age of 90. So that's a goal...

Leo: But that's not impossible. You take a lot of vitamins.

Steve: I feel great.

Leo: Yeah. Let's shoot for age 90.

Steve: That would be great.

Leo: One of our sponsors, I want to ask you about this because I thought it was kind of interesting, and they sprung this on me today, they had never mentioned this to me before, they're called Spaceship, and they have a - they're a domain name registrar. They do web hosting, that kind of stuff. But they have now announced a new messaging product called Thunderbolt. It's for iPhone and iOS.

Steve: Okay.

Leo: Unfortunately they don't describe - they mention, they call it end to end, but they don't say how they're doing it. So that's - I'm not going to emphasize the end to end. But what they do do that's kind of interesting is your ID is a domain you control. So there's no password or anything. You just put a text, you know, in your DNS.

Steve: Yeah.

Leo: And so, for instance, I'm LeoLaporte.me. I control that domain. So that is now my ID on this Thunderbolt. And you can do voice messaging, video messaging, and chat messaging with it using your domain. So no passwords necessary, either. So I think that that's kind of an intriguing idea because using - especially if you have DNSSEC-protected DNS records. That's a pretty secure way of identifying you; right?

Steve: Yeah.

Leo: Controlling that domain.

Steve: That's very cool.

Leo: Yeah. I thought it was interesting. I've messaged them saying we can talk about it. I'm not going to claim it being end to end until you tell me how you're doing that. But even if it's not, I mean, they don't store messages.

Steve: Presumably what you're publishing in your DNS record is your public key.

Leo: That's exactly right.

Steve: And so somebody else could obtain your public key from your DNS record. Then if they encrypt something under that public key, only you...

Leo: Could read it.

Steve: ...who control the matching private key would be able to see it. So I would say that qualifies as end-to-end.

Leo: Yeah. I want to know more about it. I asked them because I'd like to know the deets. But isn't that a clever idea? Especially for a domain registrar. They say it's going to be free forever because really it's just a way of getting people to register a domain; right? And now you get this messaging attached to it. Anyway, that's a sponsor. I don't want to belabor it. But I'll ask them and find out more. But I just thought it a clever way of identifying yourself.

Steve: Leo, Bob Southwell wrote: "Hi, Steve and Leo."

Leo: Okay.

Steve: "Your story about your wives talking to you from the other end of the house reminded me of this one." And he actually put up a screenshot in big, bold text that says: "Why does my wife always wait until I'm at the opposite end of the house before asking me to 'Merm frner mernferr brnerfer!?'"

Leo: Oh, I hope Lisa's not listening.

Steve: And I have to say it was a surprising relief to me last week, Leo, when you mentioned that your wonderful wife Lisa shared the tendency my own wife has of talking to me when I have no chance of understanding what she is saying or may have asked me. In fact, following last week's podcast I had thought about it several times since. So when Bob's note showed up, and to further learn that this is actually a common enough thing to be a meme...

Leo: It is a meme, yeah.

Steve: Yes. I said, well, anyway...

Leo: There you go.

Steve: Actually it made my plan to be usefully functional past the age of 90 seem somewhat less stressful.

Leo: Well, and I wear hearing aids because she says you can't hear me, so I wear hearing aids. And it doesn't help at all, I've got to tell you.

Steve: I get the same. Lorrie says, "You're just not listening." Yes, I'm straining.

Leo: I'm listening. I'm trying, trying as hard as I can.

Steve: And I just, like, merm frner mernferr brnerfer.

Leo: Well, to be fair, I do it to her, too. So all the time I'm talking to her from the other room.

Steve: Okay. So I had planned to end our feedback for the week on that last bit of fun, but before I could close my email client I encountered another note that I needed to share. Marcus Hufvudsson...

Leo: Merm frner mernferr brnerfer.

Steve: Yeah, yeah. H-U-F-V-U-D-S-S-O-N.

Leo: Hufvudsson.

Steve: Marcus, I'll just call you Marcus.

Leo: Yeah.

Steve: He said: "Dear Steve. Given the recent discussions on public-facing server security on the podcast, I thought I'd drop a note that might be of interest to everyone listening. I'm a long-time user, and nowadays the sole maintainer, of the Free/Open Source Portsentry" - that's at https://portsentry.xyz - project. "Portsentry quietly listens to unused ports you specify; and upon detecting traffic, the connection attempt will be logged, and you can optionally take actions, such as blocking the connecting IP via the systems firewall. Portsentry supports listening for a variety of port connection techniques, such as TCP SYN, FIN, XMAS and NULL scan techniques," he says, "with more detection avoidance and enumeration techniques planned. It can also listen for UDP traffic."

He says: "I usually cite two main use cases for Portsentry. Use-case 1 is an 'enumeration interference tool.'" He said: "By blocking source IPs trying to access unused services on your machine, you effectively prevent bots from enumerating your services, as well as interfere with targeted enumeration attacks. For example, if you're providing a public-facing web server on TCP port 80 and 443, you would set up Portsentry to listen for connection attempts on the other TCP service ports: 1-79, 81-442, and 444-1024." He said: "Since legitimate traffic would never attempt to access ports for nonexistent services, blocking anyone who does try to access them will cut them off from further probing your actual public facing services." He says: "Hint: Blocking the telnet port still to this day will get rid of a ton of bots."

And then for use-case 2: "Deploying Portsentry internally in your organization's networks - such as the LAN, WIFI, VPN, Management Networks, et cetera - will turn Portsentry into a type of NIDS (Network Intrusion Detection System). Since no legitimate traffic within your organization should ever touch the services Portsentry is listening for, a connection attempt would be a strong indication that something is not right." He said: "I usually set up Portsentry in a dedicated VM or container and just listen to port 1-65535. Since the dedicated Portsentry host should never be touched in your organization anyway, again, any traffic to it should be taken seriously.

"Of course, the Portsentry Project is a small but useful cog in what should be a larger and more complete cybersecurity system. So it should of course be used in conjunction with other tools and techniques. Best regards, and thanks to you and Leo for your work. Marcus."

So I wanted to share this because I think it's sort of brilliant for internal LAN network monitoring. It is 100% true that we should never expect to encounter any traffic inside our LANs that isn't deliberately aimed at a specific service present at a specific IP. Anything that appears to be "guessing" about services that might be present should sound alarms. Under no circumstances would we ever expect anything to be scanning around inside our LANs, and anything that did so should be immediately sequestered and held to account for itself. Any form of probing should raise holy hell.

Now, this is also technically true for the significantly larger network which we all know as the Internet, or the WAN, as opposed to our local LANs. Imagine if we were to immediately block any remote IP that attempts to connect to any publicly available IP and port that is not advertised through our domain's DNS. When you stop to think about it, DNS is the only official way the IP for any given service for which we intend to solicit anonymous public traffic - such as the web or email - should be found. So no traffic that hits any non-public IP and port should ever be tolerated, and immediately adding any such IP to a block list would be reasonable.

Now, having said that, attempting to "tame" the wider Internet is probably a fool's errand. For one thing, we know that innocent routers are being commandeered by bad guys for use as proxies. So blocking any source of "Internet Background Radiation" might be going too far. But the same is absolutely not true for a LAN. A LAN absolutely could be and should be tamed. And I'm pretty certain that a passive monitor ought to be able to detect suspicious activity.

Having thought about this while writing this, one problem that occurs to me is that wired Ethernet switches are inherently isolating. They acquire an awareness of which Ethernet adapters, by MAC address, are living on which port and selectively route traffic destined to those addresses only on the appropriate port. But there is one class of traffic that all switches broadcast, which is ARP. ARP has the "who has this IP" broadcasts. This is all stuff that we discussed in detail and depth back in the early bygone days of this podcast. ARP stands for the Address Resolution Protocol. It's an Ethernet protocol that was invented to map the 32- and 128-bit Internet IP addresses to 48-bit physical hardware adapter MAC addresses.

Ethernet is not actually addressed by IP addresses. What we see are IP addresses. But there's a less-seen mapping going on behind the scenes because Ethernet is addressed by these universal 48-bit MAC addresses. So when a PC, a mobile, an IoT, or any other device wishes to use Ethernet to send an Internet-style IP packet to a specific IP address on the LAN, an internal ARP table is examined to see whether the MAC address that's associated with the IP address is already known to the device. If it is, the outbound Ethernet packet is addressed to the IP's corresponding MAC address, and off goes the packet.

But if the IP's corresponding MAC address is not known, it must first be obtained. So the device needing to know broadcasts an ARP message which literally asks, "Who on the Ethernet network has this IP address?" Since the unknown device could be anywhere on the Ethernet network, any Ethernet switching device that receives this message relays it out on every one of its other ports. This is why this is known as an ARP broadcast, because it's broadcast to everywhere. It's literally broadcast to every other device that's participating on the locally connected Ethernet network.

So here's why this is interesting. For one thing, these ARP broadcasts occur at a very low level of any operating system's networking layer, and are not under the control of any application. So malware would have no way of either observing or preventing them. The other reason this is interesting is that this means that an outpost placed anywhere on the Ethernet would be able to monitor and observe any and all ARP discovery operations where any IP-enabled machine on the network is requesting the IP of any other. They may send traffic to a networked printer and perhaps a few other devices. But generally no machine on the internal LAN would be expected to do more than that.

And no machine would be expected to be poking around anyone's LAN at random, especially asking for the MAC addresses for any IP addresses that do not exist on the LAN. Any behavior of that sort should immediately raise suspicion, and any behavior of that sort would also be immediately obvious to any other device on a network that might be monitoring and watching ARP traffic.

So my point is this. Again, while I don't have a ready-to-plug-in solution, this is another opportunity for anyone who might be interested, and it would be pretty slick to have someone act upon it. The device could be something like a Raspberry Pi running Linux. If it was plugged into any unused Ethernet router or switch port, it would inherently have access to the entire network's ARP broadcasts because that's the nature of ARP. Everyone inherently needs to be able to receive those broadcasts. This renders any attempt by any device of any kind to communicate via Ethernet to any IPs that it hasn't already contacted, and that makes its attempt to do so readily apparent. Malware could be detected immediately. So something to think about.

I wanted to take a moment to note that, so far, the second season of Disney's "Andor" Star Wars spin-off series is astonishingly good, as good as the first season. There may be slightly less showoff-y special effects in Season 2, but it is a plot-driven series. I'm about halfway through Season 2, and frankly I'm in awe at the idea of what I would call "mature adult Star Wars." Or to put it another way, there is no sign whatsoever of either Ewoks or Jar Jar Binks.

Leo: Meesa saysa thatsa no gooda. Wheresa Jar Jar? Ugh.

Steve: Ugh. Yes, It's very clear that Andor's producers would never consider introducing any such nonsense, Leo.

Leo: No. Yes.

Steve: I've also noted the great restraint that's been used with the appearance of non-human aliens in general. A few scenes will feature them in brief conversation, but they're not used as a distraction or to increase the otherworldly cred of the series. What we have in Andor is intriguing mature adult drama, with political machinations and sort of the use and abuse of power. It's set in the Star Wars universe during the early days of the rise of the Empire. And of course it's got breathtaking planetscapes and skylines.

Leo: It's really nicely done. It's still...

Steve: And a flagrant use of anti-gravity technology, like, you know...

Leo: Well, you've got to have something.

Steve: Also there's no mysticism. We don't have Yoda or Jedi.

Leo: Or midi-chlorians.

Steve: Yeah.

Leo: Although somebody, Anthony said somebody spotted a Jar Jar Binks skull in Luthen's shop.

Steve: Good.

Leo: So it might be a little [crosstalk]. Yeah, we know there's no Jar Jar in here. He's dead. So I will look for it. There's a Gungan skull in Luthen's gallery somebody found, yeah.

Steve: What we have is the early seeds of what eventually grew into the Rebellion.

Leo: Right. Which that's why I like it, the story of how it started.

Steve: Yeah. It's just excellent science fiction content. Wikipedia had a short paragraph. They said: "Andor is a gritty, cynical, and detailed view of how the Galactic Empire government works."

Leo: It's great to see just the tyranny.

Steve: Yes.

Leo: You could see why they were rebelling.

Steve: Yes.

Leo: I'm sorry. Go ahead.

Steve: Yeah. And it says: "And the consequences of its actions upon everyday citizens. Beginning five years before the events of Rogue One and A New Hope, the series employs an ensemble cast of characters to show how a Rebel Alliance is forming in opposition to the Galactic Empire. One of these characters is Cassian Andor, a thief who becomes a revolutionary and eventually joins the Rebellion." And I'll just add that IMDB rates the series at the hard-to-achieve 8.5 out of 10, and Rotten Tomatoes gives it a 96%.

Leo: Wow.

Steve: And I also noted that if you're now despairing, having heard all this, of not having a Disney+ subscription, the minimal Disney+ plan is just $11 for a month. And the first two full complete seasons give you a total of 24 enjoyable episodes. So you could subscribe for $11, binge for 24 hours, or maybe spread it out over a week or two, and then easily unsubscribe. You'd be exhausted, but you'd also have two weeks of, or two full seasons of a really good science fiction series.

Leo: Yeah, I'm enjoying it, yeah.

Steve: I'm glad to know you're watching it.

Leo: It was a little tough because they didn't do much of a recap. They didn't say, "Last season." They didn't do any of that.

Steve: Nope.

Leo: They launched right into it.

Steve: Right. And in fact I was expecting that. They didn't do that. Lorrie was immediately lost because, I mean, she likes to kind of go around the house and be doing other things in the background and kind of be listening halfway. This thing requires your absolute focus. I mean, it is detailed and in depth.

Leo: Yeah. It's like a spy story; right? It's good, yeah.

Steve: Yes. The reason there's no recap is that it is there, Leo, it's just a separate thing you have to select.

Leo: Yes.

Steve: There is a...

Leo: There's a 14-minute recap on the Disney+.

Steve: Yes.

Leo: That you can watch.

Steve: So definitely if - because it was, what, it was three years ago that we had the first season.

Leo: Right, right.

Steve: So they made us wait a long time.

Leo: And this starts up a year later.

Steve: Right.

Leo: So there's definitely some, like, what? What? What happened?

Steve: My only annoyance with it is that I generally find subtitles to be a distraction. I prefer to listen with my ears while watching with my eyes.

Leo: I do, too.

Steve: But part of the reality of the production is that, you know, like for example two people will be holding an important conversation while walking and more or less muttering to one another.

Leo: Right.

Steve: Even if you back up and turn up the volume and listen intently, it's impossible...

Leo: We are sounding like such old men. You know those young people today murmur.

Steve: How they talk. What are they saying? Anyway, turn on closed captions.

Leo: I have to watch, I don't like it, but I have to watch with subtitles on almost everything now.

Steve: Yeah. Yeah. One last little bit of news. Owen LeGare. He said: "Looking forward to SpinRite 7 with better support for USB and solid-state drives." Amen. He said: "After your discussion of solid-state drives in storage becoming unreadable, I started using SpinRite to check the performance of all of mine and found significant degradation in the read speeds on portions of many of the drives. Sometimes a SpinRite Level 2 would fix the issue, but I usually had to run a Level 3 on the one third or two thirds of the drive that had slowed to get their performance back to full speed. Your comments on heat being a big factor is very true. Many of the flash drives I had at room temp for only a

couple of years were in worse shape than any of the drives I had stored in the freezer, some which had been stored for 10 years.

"After you finish the DNS Benchmark, please consider a paid version of ReadSpeed that would work on USB drives so we could identify smaller areas of solid-state USB drives that need a Level 3 refresh. Knowing what a mess the USB standards have been over the years, I'm not expecting SR7 for many years in the future after seeing all the BIOS issues encountered developing SpinRite 6.1. Thanks. Owen."

Okay. So among the several pieces of interesting feedback Owen shared, his experience with temperature being a huge factor in Flash storage data retention - and almost certainly its reliability - was the clearest that I've seen. It would be great if that guy who was doing the unpowered SSD endurance testing would incorporate temperature into his testing. The physics say that it really ought to make a huge difference, and I would strongly encourage anyone who may be archiving data on solid-state memory of any kind to store it in a very cool or perhaps even a freezing temperature. That won't hurt it.

If you're a SpinRite owner, first give any such device at room temperature a full Level 3 scan to establish a full recharge across all of its data storage cells. Then perhaps toss one or more of those drives with some desiccant packs into a sealed Ziploc bag, manually suck the air out to remove any moisture-bearing air as much as possible, finish sealing the bag, and drop it into the freezer. And along with my Palm Pilots, they will hold onto their data forever.

Leo: Do you have any food in your freezer?

Steve: There's no room, Leo. No. Okay. Our last break, and then we're going to talk about the discovery of rogue comms tech found in the U.S. power grid.

Leo: Scary, scary, scary.

Steve: Okay. So because the news that I need to share today is so upsetting, I need to first do what I can to make sure we're all on the same page about the source of this information. The news that this podcast will be sharing this week is reported by the Reuters News Agency. Reuters, as it's more commonly known, is a news agency owned by Thomson Reuters. It employs around 2,500 journalists and 600 photojournalists spread across 200 locations worldwide and writing in 16 languages. It's one of the largest news agencies in the world, having been established, believe this, in London in 1851 by Paul Reuter.

So their news last Wednesday, May 14th carried the headline: "Rogue communication devices found in Chinese solar power inverters."

Leo: Oh, wow.

Steve: Here's what we know thanks to this reporting from Reuters. They wrote: "U.S. energy officials are reassessing the risk posed by Chinese-made devices that play a critical role in renewable energy infrastructure after unexplained communication equipment was found inside some of them, two people familiar with the matter said.

"Power inverters," they wrote, "which are predominantly produced in China, are used throughout the world to connect solar panels and wind turbines to electricity grids. They are also found in batteries, heat pumps, and electric vehicle chargers. While inverters are built to allow remote access for updates and maintenance, the utility companies that use them typically install firewalls to prevent direct communication back to China. However, rogue communication devices not listed in product documents have been found in some Chinese solar power inverters by U.S. experts who strip down equipment hooked up to grids to check for security issues, the two people said.

"Over the past nine months, undocumented communication devices, including cellular radios, have been found in some batteries from multiple Chinese suppliers, one of them said. Reuters was unable to determine how many solar power inverters and batteries they've looked at. The rogue components provide additional, undocumented communication channels that could allow firewalls to be circumvented remotely, with potentially catastrophic consequences, the two people said.

"Both declined to be named because they did not have permission to speak to the media. However, Mike Rogers, a former director of the U.S. National Security Agency (our NSA) said: 'We know that China believes there is value in placing at least some elements of our core infrastructure at risk of destruction or disruption. I think that the Chinese are, in part, hoping that the widespread use of inverters limits the options that the West has to deal with the security issue.' Meanwhile, a person for the Chinese embassy in Washington said: 'We oppose the generalization of the concept of national security distorting and smearing China's infrastructure achievements.'

"Experts said that these rogue communication devices to skirt firewalls and switch off inverters remotely, or change their settings, could destabilize power grids, damage energy infrastructure, and trigger widespread blackouts. One of the people asked said: 'That effectively means there is a built-in way to physically destroy the grid.'

"The two people declined to name the Chinese manufacturers of the inverters and batteries which were found to contain extra communication devices, nor say how many they had found in total. The existence of the rogue devices has not previously been reported, nor has the U.S. government publicly acknowledged the discoveries. When asked for comment, the U.S. Department of Energy said it continually assesses risk associated with emerging technologies, and that there were significant challenges with manufacturers disclosing and documenting functionalities.

"A spokesperson said: 'While this functionality may not have malicious intent, it is critical for those procuring to have a full understanding of the capabilities of the products received.' The spokesperson added: 'Work is ongoing to address any gaps in disclosure through software bill of materials or inventories of all the components that make up a software application.'"

Okay, now, I'll just interrupt and say that a software bill of materials doesn't quite address the issue of hidden cellular radios, and software bills of material are voluntary disclosures of software components and libraries. They don't address concerns of possible malicious intent.

Reuters continues: "As U.S.-China tensions escalate, the U.S. and others are reassessing China's role in strategic infrastructure because of concerns about potential security vulnerabilities, two former government officials said. U.S. Representative August Pfluger, a Republican member of the Committee on Homeland Security, told Reuters: 'The threat we face from the Chinese Communist Party is real and growing. Whether it's telecom hacks or remotely accessing solar and battery inverters, the CCP stops at nothing to target our sensitive infrastructure and components. It is about time we ramp up our efforts to show China that compromising us will no longer be acceptable.'

"In February, two U.S. senators introduced the Decoupling from Foreign Adversarial Battery Dependence Act, banning the Department of Homeland Security from purchasing batteries from some Chinese entities, starting October 2027, due to national security concerns. The bill was referred to the Senate Committee on Homeland Security and Government Affairs on March 11th and has yet to be enacted." Now, that's interesting since it suggests that there are areas of the government that must be aware of at least the potential for this sort of abuse.

Reuters explains of this bill: "It aims to prevent Homeland Security from procuring batteries from six Chinese companies Washington says are closely linked to the Chinese Communist Party: Contemporary Amperex Technology Company, BYD Company, Envision Energy, EVE Energy Company, Hithium Energy Storage Technology Company, and Gotion High-tech Company. None of these six companies responded to Reuters' requests for comment.

"Additionally, utilities are now preparing for similar bans on Chinese inverter manufacturers, three people with knowledge of the matter said. Some utilities, including Florida's largest supplier, Florida Power & Light Company, are attempting to minimize the use of Chinese inverters by sourcing equipment from elsewhere, according to two people familiar with the matter. FPL did not respond to requests for comment.

"The DOE spokesperson said: 'As more domestic manufacturing takes hold, DOE is working across the federal government to strengthen U.S. supply chains, providing additional opportunities to integrate trusted equipment into the power grid.'

"Huawei is the world's largest supplier of inverters, accounting for 29% of shipments globally in 2022, followed by Chinese peers Sungrow and Ginlong Solis, according to the consultancy Wood Mackenzie. German solar developer 1Komma5 said, however, that it avoids Huawei inverters because of the brand's associations with security risks. 1Komma5's Chief Executive Philipp Schroeder said: 'Ten years ago, if you switched off the Chinese inverters, it would not have caused a dramatic thing to happen to European grids; but now the critical mass is much larger. China's dominance is becoming a bigger issue because of the growing renewables capacity on Western grids and the increased likelihood of a prolonged and serious confrontation between China and the West.'

"Since 2019, the U.S. has restricted Huawei's access to U.S. technology, accusing the company of activities contrary to national security, which Huawei denies. Experts explained that Chinese companies are required by law to cooperate with China's intelligence agencies, giving the government potential control over Chinese-made inverters connected to foreign grids. While Huawei decided to leave the U.S. inverter market in 2019 - the year its 5G telecoms equipment was banned - it remains a dominant supplier elsewhere. Huawei declined to comment.

"Experts explained that, in Europe, exercising control over just 3 to 4 gigawatts of energy could cause widespread disruption to electrical supplies. The European Solar Manufacturing Council estimates that over 200 gigawatts of European solar power capacity is linked to inverters made in China, equivalent to more than 200 nuclear power plants. At the end of last year, there was 338 gigawatts of installed solar power in Europe, according to industry association SolarPower Europe.

"Uri Sadot, cybersecurity program director at Israeli inverter manufacturer SolarEdge, said: 'If you remotely control a large enough number of residential solar inverters, and do something nefarious at once, that could cause catastrophic implications to the grid for a prolonged period of time.'

"Other countries such as Lithuania and Estonia acknowledge the threats to energy security. In November, the Lithuanian government passed a law blocking remote Chinese

access to solar, wind, and battery installations above 100 kilowatts, by default restricting the use of Chinese inverters. Estonia's energy minister said this could be extended to smaller rooftop solar installations. Estonia's Director General of the Foreign Intelligence Service, Kaupo Rosin, said the country could be at risk of blackmail from China if it did not ban Chinese technology in crucial parts of the economy, such as solar inverters. Estonia's Ministries of Defense and Climate declined to comment when asked if they had taken any action.

"In Britain, a person familiar with these matters said the government's review of Chinese renewable energy technology in the energy system, due to be concluded in the coming months, includes looking at inverters."

And get this. Here's one that slipped under the radar. Reuters wrote: "In November, solar power inverters in the U.S. and elsewhere were disabled from China, highlighting a risk of foreign influence over local electricity supplies and causing concern among government officials, three people familiar with the matter said. Reuters was unable to determine how many inverters were switched off, or the extent of disruption to grids. The DOE declined to comment on the incident." But again, last November, China remotely switched off power in the U.S.

"The incident led to a commercial dispute between inverter suppliers Sol-Ark and Deye" - spelled D-E-Y-E - "the people said. A Sol-Ark spokesperson said: 'Sol-Ark does not comment on vendor relationships, including any relationship with Deye; nor does it have any control over inverters that are not branded Sol-Ark, as was the case in the November 2024 situation you referenced.' Deye, for their part, did not respond to requests for comment.

"The energy sector is trailing other industries such as telecoms and semiconductors, where regulations have been introduced in Europe and the U.S. to mitigate China's dominance. Security analysts say this is partly because decisions about whether to secure energy infrastructure are mostly dictated by the size of any installation. Household solar or battery storage systems fall below thresholds where security requirements typically kick in, they said, despite now contributing a significant share of power on many Western grids.

"NATO, the 32-country Western security alliance, said China's effort to control member states' critical infrastructure, including inverters, were intensifying. A NATO official said: 'We must identify strategic dependencies and take steps to reduce them.'"

Okay. So again: "Two people said that rogue communication devices not listed in product documents have been found in some Chinese solar power inverters by U.S. experts who strip down equipment hooked up to grids to check for security issues. And over the past nine months, undocumented communication devices, including cellular radios, have been found in some batteries from multiple Chinese suppliers."

The story caught me by surprise and had a great deal of salience for me because, you know, we're always talking, we have talked many times about theoretical vulnerabilities in power grids, and about how devastating an attack upon our U.S. power grid would be. And now we learn that these concerns have moved from the world of theory to reality. Why would there be undocumented radios in inverters that are not part of the documentation or the bills of material or the operating specifications?

We've been moving to renewable energy sources which happen to inherently produce direct current. Solar cells and wind-powered generators output DC. But the transmission of Direct Current is inherently more lossy than the transmission of Alternating Current, which is why our power grid carries AC current over long distances. DC cannot be "transformed" in order to make a trade of current for voltage. For that, alternating

current is needed, and it's the job of inverters to convert the direct current produced by renewable energy sources into alternating current. As soon as you have alternating current, power transformers can be used to raise its voltage while reducing its current to levels that are far more efficient for long-distance transport.

Given China's well-proven ability to manufacture high-quality electronic systems at unbeatable low cost, it was only natural for the manufacturers of solar cell systems and wind turbines and those assembling larger renewable power solutions to purchase the required inverters from China. In many regards, they would have been the best solutions available, and probably still are.

But when we learn as we did about this last November event, where solar power inverters in the U.S. and elsewhere were remotely disabled from China, suddenly those Chinese inverters no longer seem like such a bargain. Reuters explained that users of these Chinese devices are aware of this danger.

So in an apparent attempt to avoid being cut off from their equipment because manufacturers are putting in firewalls, installing firewalls specifically to block Chinese access, some of these Chinese inverters and batteries have been found to incorporate cellular radios. And, you know, bending over backward to be fair, we don't know why. Right? But they're not in the specs, they don't appear in the schematics or in any diagrams, and they're not required for the intended functioning of the equipment. So regardless of how they got there, who put them there, or why, they should not be there. Given the devastation that could be wrought if power grids were to collapse at the whim of a hostile foreign power, this is not a chance anyone can take.

The good news is this has come to light today at a time that's early enough for appropriate actions to be taken. And even though this has not received a great deal of mainstream press, those who need to know are being informed. I did some digging. The site Utility Dive carried the headline "Rogue communication devices found on Chinese-made solar power inverters." PV News, where "PV" is short for Photovoltaics, as in solar cells, their headline was "Rogue devices found in Chinese solar inverters raises cybersecurity alarm in Europe." And the publication Industrial Cyber's headline was "U.S. energy sector at risk, as Chinese inverters are under investigation for suspicious communications gear."

So it appears that there will be some retrofitting, or at least much closer examination of any already installed equipment having a Chinese origin or having unknown provenance. And it's unlikely that there will be any new use of any foreign technology that hasn't been fully vetted in critical areas. You know, it's unfortunate, but it's the world we're living in today, Leo. We end up, you know, having to switch to higher cost alternatives because we can't trust everyone to supply gear that's safe for us to use.

Leo: How big are those radios? Are they easily visible?

Steve: What occurred to me is that all you would need - because we now have satellite comms. We have, like, satellite radio; right?

Leo: Yup.

Steve: You could just do a tiny little satellite receiver if you wanted something that was able to remotely receive a signal from the orbiting mothership and take some action.

Leo: Why not just put it on the Internet? Is it not Ethernet - they're not Ethernet connected?

Steve: Because they have, the operators have firewalled them.

Leo: Yeah, yeah.

Steve: They're deliberately, you know, prevented from talking to China.

Leo: So that's why the radios, yeah, yeah. It's really interesting. Did they say how prevalent they think this is? Or it's just a few exceptions?

Steve: No. The two people who were the biggest source of information were willing to say that it was multiple suppliers, multiple instances, multiple batteries, multiple inverters, but not what the count was.

Leo: Right. Well...

Steve: And that it was somebody who, like, took the lid off of something and said...

Leo: What's this?

Steve: ...wait a minute. What's that?

Leo: What is back there?

Steve: Yeah.

Leo: I mean, I used to, you know, when we had solar power, I had two inverters in my garage, of course, you know. And I never looked inside of them. It's not so much a big deal at my house, but it's a big deal in a big solar farm, you know, powering a city. This is the problem with supply chain attacks.

Steve: And my god, I had no idea that solar had gotten as big as it is. 200 gigawatts of power is, they said, the equivalent of 200 nuclear reactors.

Leo: It's amazing.

Steve: So it's like, you don't need nuclear anymore. You just need space to lay out the panels.

Leo: And fortunately in the U.S. we've got a lot of space. Lot of desert. Lot of sunshiny space.