Detecting Unwanted Location Trackers

Description: Last week Google activated their Passkeys support. What does that actually mean? Do TP-Link routers auto-update by default? What trouble did a secretive branch of the U.S. Marshals get into? When and why will Chrome be eliminating the padlock icon? Were you prompted by Apple's new Rapid Security Response? What did Elon Musk do to upset WordPress, and why is it a win for Mastodon? How many fake news AI-driven websites have been spotted so far, and are they convincing? What's this about Russia dropping TCP/IP in favor of their own Russian network protocol? What three mistakes does Vint Cerf, co-designer of the Internet Protocols, think he made? And finally, in the first half of our two-part very deep dive into the design of the next-generation location tracking devices, will you be put off when you learn that law enforcement is able to query for the identity of any device's owner? Fasten your seatbelts for

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-922.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-922-lq.mp3

---

another interesting Security Now! podcast brought to you by TWiT, the itch that Leo scratched.

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. We're going to talk about Google and Passkeys. They activated support. But how do Passkeys work? What are their problems? What are the things they solve? We'll also talk about Vint Cerf, the father of the Internet, and the three things he thinks he did wrong when he designed it. And then Steve will dive into an Apple proposal to make AirTags and other Bluetooth trackers more secure? Or is it less? Security Now! is next.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 922, recorded Tuesday, May 9th, 2023: Detecting Unwanted Location Trackers.

It's time for Security Now!. Oh, I know you've been waiting all week for this. Here he is. He's here. Mr. Steven Gibson. You've never missed a show, Steve, have you.

Steve Gibson: No.

Leo: Never got sick.

Steve: No. I've actually scheduled some surgery in between shows, specifically.

Leo: Wow. Even when you traveled to wherever, Scandinavia, you never missed a show.

Steve: Nope. Never missed a show.

Leo: In fact, Steve gets mad when we tell him you've got to take a week off during Christmas so we can do a Best Of.

Steve: Well, it's why I'm so happy that I'm on Tuesday and not on Monday. Because those three-day weekends used to piss me off. I'm like, wait a minute.

Leo: Well, we're glad that you don't miss a show, and you are now in Episode 922. It's going to be a good one, I think.

Steve: It is actually going to be a good one. I expected to be able to finish today's topic, and I could not, for reasons that we're going to see. So this is actually Part 1 of a two-parter titled "Detecting Unwanted Location Trackers." So lots to talk about.

Last week Google activated their Passkeys support. What does that actually mean? Do TP-Link Router auto-update by default? What did a secretive breach of the U.S. Marshals get into? When and why will Chrome be eliminating the padlock icon? Were you prompted by Apple's new Rapid Security Response last week? What now has Elon Musk done to upset WordPress, and why is it a win for Mastodon? How many fake news AI-driven websites have been spotted so far, and are they convincing? What's this about Russia dropping TCP/IP in favor of their own Russian network protocol?

Leo: Oh, lord.

Steve: I know.

Leo: I've heard about that.

Steve: Oh, there's some good stuff here. Okay. What three mistakes does Vint Cerf, co-designer of the Internet protocols, think he made? And finally, in the first half of our two-part very deep dive into the design of the next-generation location tracking devices, will you be put off when you learn that law enforcement is able to query for the identity of any device's owner? So fasten your seatbelts for another interesting Security Now! podcast brought to you by TWiT, the itch that Leo scratched.

Leo: As long as it doesn't leave you itching, we're okay. No red rashes from listening to the show. You didn't mention, but we have another Picture of the Day in a continuing series. I love this. I love this one. Okay, Steve.

Steve: Next time you talk about Kolide, because you're right that the LastPass, the cause of the LastPass breach is a perfect example of this, the developer did have a

publicly exposed instance of Plex server for which a patch had been available for years that was never applied.

Leo: Yup.

Steve: And that's how the bad guys got in.

Leo: That's how it can happen. I mean, it's that easy. And that's a hole that you, fortunately, you can patch. If you're using Okta, you can use Kolide, and you're done. I guess they would.

Steve: Okay. So our Pictures of the Week...

Leo: Oh, boy.

Steve: ...have been elevated to a new level, Leo. One of our listeners was driving down the street in Texas when he saw this and pulled over and stopped and took a picture.

Leo: We've trained them to look for this stuff now. That's awesome.

Steve: So for those who aren't looking at the video, we've got yet another ridiculous gate, like in the middle of nowhere. What I love about this, so there's like a big concrete driveway coming up, like across a sidewalk which is set back from the curb a good distance. And then there's a big gate with both sides hinged such that it would, you know, they would swing apart in order to allow you access to the full width of the drive. Okay. And of course the problem is that there's nothing on either side of the posts that anchor this gate. And there's some evidence of, as there always is, of people trotting or driving, like, around it. So you can sort of see it's a little bit worn over there to the left. But so this gate is sort of just out in the middle of nowhere with, like, nothing to prevent anybody from thinking, uh, okay, if I want to go behind it, I can. But what I love most is that, if you zoom in on the picture, there is a security cable and a padlock which is, like, locking the two halves of this gate together.

Leo: Well, you wouldn't want anybody opening the gate. I mean, come on, man. That's hysterical.

Steve: No, no, no, no. You want to make sure that both halves stay firmly closed. Wow.

Leo: Oh, that's crazy.

Steve: Okay. So Google and Passkeys. Last week, Google formally launched their support for Passkeys. I popped onto Jason's Tech News Weekly show last Thursday and spent some time placing this announcement into perspective.

Leo: Oh, good. Oh, good, I was hoping. We need your input, yeah.

Steve: We had a great time. I explained, and there's a picture here in the show notes, Leo, how 50 years ago, five zero, 50 years ago, in 1973, when I was at UC Berkeley, I sat in front of that, I mean, I found the picture of it, and I knew exactly what it was.

Leo: It's a Hazeltine 2000 dumb terminal; right?

Steve: Correct.

Leo: That's a dumb terminal, yeah.

Steve: Correct. I logged onto the campus mainframe sitting in front of that Hazeltine 2000 dumb video terminal. You'd hit the Break button, which was the second one back from...

Leo: Oh, yeah. Oh, yeah, I remember that, yeah.

Steve: ...from the right, in order to get the attention of the mainframe, after which a few seconds later it would prompt you for your username and then your password. And if those were both known to the system, you'd be logged in. So my point was, that I was making on Jason's show last Thursday, is nothing has significantly changed during the intervening 50 years. That was 50 years ago.

And, you know, today, 50 years later, every web server and website on the Internet prompts for, I mean, with yes a few notable exceptions, but largely prompts for a username and password, or at least can, then verifies that they're correct. Now, we've gotten much better about protecting the user's password secret using hashing and brute force resistant PBKDFs, which we've had a lot of fun talking about and explaining on the podcast. But servers are still essentially holding and comparing user secrets. And that is what is finally changing with the introduction of WebAuthn. WebAuthn finally moves us from that 50-year-old symmetric secret key model to a modern asymmetric public key model.

Under this new model, servers no longer keep secrets. They keep public keys which allow them, and only allow them, to verify the signature of a randomly generated challenge that they just freshly sent to someone who's logging on. That Hazeltine 2000 terminal I used 50 years ago was, as you know, as we just said, a "dumb terminal." So there was no way for it to perform the complex math that's required to answer a cryptographic challenge by signing it with a private key. But none of us are using Hazeltine 2000 terminals anymore. We're all logging into a remote computer with a local computer. So these two quite capable computers that are communicating are able to have a mindbogglingly complex, heavy math interaction which we are able to take for granted.

So in one sense, last week's news from Google was insignificant because it's not Google who needs to make Passkey support available, it's everyone else in the world. And that's the hurdle. Just as with any major change in our industry, this is going to be an extremely slow process. At the same time, Google's announcement was significant inasmuch as it represents another necessary incremental step in the right direction.

But it's going to take a long time for this to become universal, if it even ever is. Think of it like today's one-time passwords. I just checked my OTP Auth app. I have 18 one-time passwords registered. How long ago, Leo, were you and I talking about that one-time password football that was supported by VeriSign.

Leo: Oh, yeah. PayPal. I got mine with PayPal, yeah.

Steve: Yeah, exactly. It was accepted by PayPal. It must have been at least a decade ago.

Leo: Ages.

Steve: Yet today, very few of the websites I log into offer one-time passwords as an additional authentication option. You know, those where logon security is most important generally do. But, you know, I use mine at...

Leo: That's TOTP. It's the same thing; right? The football, yeah, yeah.

Steve: Right, right. That's exactly what it was. But I suspect that Passkeys will be similar. It's going to take a while. But that said, without any question, the biggest and best news is that we now have an option. We have an open and modern public key-based means of authenticating users over networks. And 50 years from now, I expect everyone to be using it because, you know, it'll end up being built into servers, and it'll just be there and easy to use.

Leo: Answer me this, though. This is something I've been kind of trying to figure out. As far as I can tell, the FIDO and WebAuthn spec don't have any provisions for exporting or moving your keys; right? And I can understand why that might be a security flaw.

Steve: Right.

Leo: But it means that Apple, for instance, with their iCloud key, has a lock-in. There's no way I can get it out of Apple and move it to Android; right?

Steve: Correct. The only thing you can do is use your iPhone to briefly bridge to a different device and then have it create its own new unique Passkey there.

Leo: So I've gone to Steve's Marvelous Site and set up a Passkey with my iPhone. But now I want to get an Android phone. So you're saying I could log into Steve's Wonderful Site with my iPhone. Can I then tell Steve's Wonderful Site, generate a new Passkey, and set it up with my Android phone?

Steve: Uhhh...

Leo: No. I have to keep my iPhone around to login; right?

Steve: So, okay. So this is one of the reasons that I focused, when I was talking to Jason...

Leo: Yeah, you solved this with SQRL. I know you did.

Steve: Well, I did. But I focused on the backend side with Jason because these user-facing sides, they're sort of - they're up in the air. They're open. For example, it would be possible for you to sign into Steve's Wonderful Site with your iPhone, then tell Steve's Wonderful Site that the next sign-in is also from you, using a different device. So then your Android could identify itself to Steve's Wonderful Site, and now Steve's Wonderful Site would have two Passkeys, each for a different device.

Leo: Oh, so there is - okay. I've never seen this offer.

Steve: Well, no. And that's my point is it's up to, like, it's up to the UI...

Leo: Steve's Wonderful Site has to offer that; right?

Steve: Yes, yes.

Leo: And one thing we found out, at least initially, was a lot of these sites that claim to be offering Passkeys were actually not doing it themselves. They were using a third party.

Steve: Yes, yes. Which is like, you know, punting. And immediately - and we could expect that to change over time. The other thing, the other observation that Jason made was that, since Pixel has been creating, or Android and Google on Pixel phones have been creating Passkeys for a while, the moment it got activated, he looked in Google, and there were like all of his different devices were listed with their own individual Passkeys.

Leo: Yeah. Both my Google Pixel and my Galaxy S23 with - I mean, I guess I seem to dimly remember them saying something about this. But both created Passkeys. And they were preexisting. That's what they're supposed to do.

Steve: Well, that is - yeah. That is the model. So I got a tweet from Chris Smith which was interesting and apropos. He said: "Passkeys: Going through the details, can they be used if I don't own a computer or phone?"

Leo: Oh. Well, what would you be logging in with?

Steve: Well, that threw me. So he said: "Does this amount to an economic barrier to better security and privacy?" And so, you know, that tweet and question caught me by surprise. At first I thought, exactly as you did, how could you be doing anything if that required...

Leo: Oh. He went to the library.

Steve: Yes, yes, yes, yes.

Leo: And a lot of people who don't have their own technology do use technology at libraries. Okay.

Steve: Exactly. I remembered Internet cafs, public libraries, hotel business centers, senior citizen centers, and other such facilities which provide computers and Internet connectivity for those who don't have any sort of computer with them. So, yeah. Chris, I think, makes a point. As I noted above, the biggest change brought by Passkeys to the traditional username and password model is the presumption that the user's end of the link will have powerful crypto math capabilities.

Leo: Yes.

Steve: While any PC on the Internet will have that capability, even an Internet caf machine, Passkeys' other requirement is the ability to store any number of individual public key pairs, one pair for every Passkey association that the user has ever made to individual websites. And that's not going to be something that someone logging onto a machine in any sort of shared caf or library setting will have. Now, that brought me to...

Leo: Will there always be a password fallback, though? Or do you think at some point...

Steve: Yes. No. There will always be. Which actually is, it creates the weak link; right? I mean, we've often talked about the weakest link being the one that gets you.

Leo: Right.

Steve: Well, if you still actually have usernames and passwords, then site breaches will still be leaking your password.

Leo: Well, I guess the presumption is that any new site, I'm going to create a Passkey. I won't have a password on that site; right?

Steve: If they let you not have a backup.

Leo: They may say create a fallback, yeah.

Steve: Yeah.

Leo: Oh, god. I bet they will, too, because they don't want to field customer support calls.

Steve: Exactly. Exactly.

Leo: Oh, this is a mess.

Steve: It is a mess. So that brought me to the question of a personal FIDO dongle because this user might not have a phone or a computer, but what about a dongle? But it's unclear, I think, what the future of hardware dongles is now, after they essentially failed to achieve market critical mass traction, and the FIDO Alliance finally deigned to soften authenticator requirements to allow smartphones and PCs to qualify as sufficiently secure hardware authenticators. It seems to me that this dramatically reduces, if not almost eliminates, the pressure to purchase a separate freestanding dongle for most use cases. There could still be exceptional enforcement of the requirement for a separate freestanding authenticator in selected ultra-high security applications. But we already know that's not going to be the norm. I have a smartphone. Why do I want to purchase and carry a redundant dongle? I don't, and I won't.

Again, some users might want that, but that can now be expected to become much less common. And if that's the case, then what of longer term support for dongle authenticator logon? It's unclear that hardware authenticator support will automatically be added to websites as Passkey support is added. After all, if almost no one is carrying them around, why clutter up and further confuse the user's experience with a rarely used and largely redundant option?

Leo: Well, and as Google pointed out, they're going to support this in their advanced protection program, Passkeys.

Steve: Right, right.

Leo: Although when I log into my Google Passkeys thing, my YubiKeys are there, too.

Steve: Right.

Leo: But that's because I have YubiKeys.

Steve: Exactly. And who's going to be...

Leo: Who's going to buy them, yeah.

Steve: ...buying a YubiKey when you can use your phone.

Leo: I feel bad for Stina, really.

Steve: So apparently they cashed out, so I think she's just fine.

Leo: Oh, good. Yeah, I saw she stepped out or up.

Steve: Yeah.

Leo: And what about password managers? That puts them out of business, too; right?

Steve: Well, password managers will still be the holders of our usernames and passwords. And there's nothing to prevent a password manager from being your Passkey holder.

Leo: Well, do I want to do that? I mean...

Steve: I agree. I agree.

Leo: 1Password says they're planning to do that.

Steve: I know. And what it does solve is the cross-ecosystem problem; right? Because the password manager won't care if you're Apple or Google. I mean, it'll be a third-party password manager, essentially.

Leo: Right. But it's not a hardware enclave. It's software.

Steve: No. No.

Leo: Which makes it less secure, I presume.

Steve: Yeah, exactly. So I think that on balance Chris's point is valid. This move to higher security authentication does require some form of dedicated per-user authenticator capable of storing all of the user's public Passkey pairs and negotiating with remote websites for their use. I cannot see any way for a shared PC-only usage model to accommodate that. So it's good news that username and passwords will never go away, and users who are unable to own their own authentication device will need to operate without the benefit of public key crypto authentication.

Okay, now, I just said that I could not see any way for public key logon to be made practical. But that's not completely true. Not to belabor the point, but this highlights

another example of what SQRL's design got right and the FIDO Alliance got wrong. I'll explain, not from a position of sour grapes - that ship has not only sailed, but sank - but only to highlight that other and arguably...

Leo: Aw, Steve, you're going to make me cry.

Steve: Well, wait till you hear this, and you're going to cry. I want to highlight that other and arguably superior solutions are possible. In a SQRL world, the Internet caf or library shared PC user could simply carry their single lifetime SQRL identity printed on a one-inch square QR code. They could have it plastic laminated for longevity. That's all they would need.

Leo: So that's their private key, basically, in a QR code.

Steve: Yes. But it's a universal private key instead of needing a...

Leo: Instead of one per. Yeah.

Steve: One per everything.

Leo: That's a lot of keys, yeah.

Steve: So they sit down in front of a shared PC, which is configured to reboot after use and to never make any permanent changes, as all of those caf-style machines are. The user would click the SQRL icon, hold their QR code up to the PC's camera, and import their globally unique identity into SQRL. After then entering their one single password which is used to decrypt their SQRL identity, from that point on, every site they visit would be able to obtain their per-site anonymous identity for logon.

Leo: That's so cool. So much better.

Steve: I know. So unfortunately, unlike SQRL, Passkeys' operation requires potentially unlimited storage of individual public key pairs. So that rules out its use by any user who cannot provide some form of personal public key pair storage.

Leo: It's like they started that process that you started, when that light bulb went on seven years ago, but they didn't finish it. They didn't think it all the way through. Because there's no reason why they couldn't have done it better, to me.

Steve: No.

Leo: Yeah.

Steve: No.

Leo: That's very frustrating.

Steve: You know, I solved it all the way out, all the possibilities.

Leo: You thought all the edge cases through, and that's the thing.

Steve: Yeah. For example, say that the guy who sat down for some reason wants to appear as a different user, but he only has his one QR code. Well, there's a provision in the SQRL spec to allow you to append something to your identity to create a branch identity just for that one time.

Leo: We want that, too. We want that, too. Yeah.

Steve: I know.

Leo: This got support, somebody's saying in the IRC, DryHeat, seems obvious they wanted the smartphone to be the authenticator. If you think of who was behind this - Google, Apple, Microsoft, too - but Google and Apple did have a dog in that hunt. They want a hardware enclave that you buy from them. Plus the lock-in really bothers me. You cannot...

Steve: Yeah.

Leo: You're locked into that platform. I mean, Apple makes it possible to go to a new iPhone, but not to an Android.

Steve: The whole model is create another identity. You know, create another key pair. So it just doesn't transport the way we would like it to.

Leo: So frustrating. Oh, well.

Steve: Okay. I received some firsthand feedback from Kevin Kinneer, one of our listeners, who owns and knows his way around TP-Link routers. Here's what Kevin wrote. He said: "I have several of the TP-Link AX21 (AX1800) routers you were talking about last week, and they all upgraded automatically."

Leo: All right.

Steve: Yeah.

Leo: Yeah.

Steve: He said: "In its initial setup wizard, auto-upgrade is offered, and you would have to intentionally choose to turn it off." Yay. He said: "Perhaps when this model first came out this was not a feature; but from what I've seen lately, they are doing it right." So Leo, that supports your belief that auto-updating is beginning to appear in routers, and it supports my prayers that when the option is available, it will be enabled by default so that less-informed users will assume that it's a good thing, and they'll leave it enabled where it belongs. And thank you, Kevin, for taking the time to provide that feedback.

The advice that follows from this would be to no longer purchase any router that does not offer autonomous upgrading in the future. And I'm also certain that many of our listeners are considered to be their friends' and families' influential tech guru. So find a good auto-updating router to recommend to those who depend upon your opinion because this is clearly where the future is, especially for, I mean, consumer routers are like in this weird, like, worst of all possible worlds; right? They are, by definition, they're exposed to the public Internet, so they're not on a LAN. They're bridging the WAN and the LAN. And the fact that they're a bridge makes them enticing for attackers. The fact that they're on the public Internet means that, well, they can be poked at and prodded and searched for and found.

And they're not exciting. You know, you don't turn them on every day to watch news on them or anything. They're dusty. They're sitting in a closet somewhere. So they're like the least obvious thing anyone is ever going to think about like updating and keeping current; right? It's like, you know, the forgotten stepchild. It just does its job. And so it's like the profile of it from every direction is as bad as it could get.

Leo: By the way, I would extend it. You should not buy an IoT device of any kind that isn't auto-updating. I mean, period. Doorbell, light bulb, router, they should all auto-update. That's got to be the standard going forward. Clearly.

Steve: Yes, yes. Yeah. I would argue that we're probably leaving the first generation of devices that just came out quickly because they would work. And we're now beginning to get into it like, okay, we're saying, but what if? And it's like, okay, so it's time to get smart.

Okay. So I stumbled upon a story that became more interesting the more I dug. Ten weeks ago, back in the middle of February of this year, a computer system run by a special and somewhat secretive and elite unit of the U.S. Marshals Service, known as TOG their Technical Operations Group fell victim to a breach followed by the inevitable ransomware attack and the followon demand. The U.S. Marshals refused to pay the ransom. And so today, a full 10-plus weeks and counting later, that entire system remains offline. And when you hear what it was doing, ho ho. Okay. What makes this otherwise somewhat now generic story extra interesting is the back story of this particular system, which is no longer on the air.

It turns out that this TOG, their Technical Operations Group, is a secretive arm within the U.S. Marshals Service that uses technically sophisticated law enforcement methods to track criminal suspects through their cell phones, email, and web usage. Its techniques are kept secret to prolong their usefulness, and exactly what members of the unit do and how they do it is a mystery even to some of their fellow Marshals personnel.

The system that went down and remains down manages a vast amount of court-approved tracking of cell phone data, including location data, and this system has existed

outside regular Justice Department oversight and computer systems for years, essentially leaving it unnoticed. After the breach and attack occurred, the technical group proactively wiped the cell phones of everyone who worked with and had anything to do with the breached system, deleting all of their contacts and emails.

Since the action was taken without notice on a Friday night, it caught everyone by surprise. One staffer was working the security detail for a Supreme Court justice when they discovered that their phone had been wiped of all of its data. Although the phone still worked, the person had no emails or contacts. It would appear that the technical group either believed or knew that someone had clicked a bad link and had fallen victim to a phishing attack which allowed the bad guys to then gain entry.

However, the most significant consequence of the system going down, and remaining down for more than 10 weeks, is that one of the Marshals' best tools for finding fugitives, often used on behalf of state and local law enforcement agencies, has been completely incapacitated. It's just gone.

Okay. Sounding somewhat, and understandably, defensive, a spokesperson for the Marshals Service said that: "The breach data has not impacted the agency's overall ability to apprehend fugitives and conduct its investigative and other missions. Most critical tools were restored within 30 days of the breach discovery."

Leo: Oh, well, that's a relief.

Steve: Yeah. But 30 days? Wait. Okay? He said - and it doesn't sound like that's true anyway. We have some details in a second. He said: "Further, USMS (U.S. Marshals Service), USMS soon will deploy a fully reconstituted system with improved IT security countermeasures."

Leo: We had to buy all new computers.

Steve: Yeah. Well, and where's the backup? Right? You know, we have no details about the computer system. But it sure does sound as though the network was either without protected backups, or that these backups were also destroyed.

Leo: Oh, boy.

Steve: Because, you know, 10 weeks. Okay. So the Technical Operations Group has helped the Marshals hunt down high-value suspects in the United States and in other countries, including Mexican drug kingpin Joaquin Guzman, better known as "El Chapo." A great deal of the hunting is done through what is called pen register/trap and trace, a means of cell phone surveillance that has evolved along with phone technology. In the era of landlines, a so-called PR/TT (Pen Register Trap and Trace) meant getting a record of all the incoming and outgoing calls from a phone. Today, PR/TTs can also be applied to email accounts, and apparently also to an individual's web browsing usage, all of which can provide information critical to a manhunt.

Unlike a wiretap, a pen register/trap and trace does not monitor the contents of phone conversations. A PR/TT order for a phone's data requires law enforcement to convince a judge only that the information is relevant to an ongoing investigation, so it's a low bar, not the higher legal standard of probable cause, which is needed for a full wiretap. In

their reporting about this, The Washington Post quoted Orin Kerr, a law professor at the University of California Berkeley, who specializes in criminal procedure and privacy.

He said: "In a world where everyone has a cell phone, it's a way to track cell phones, and it's a way to track account usage. We're all on these devices all day, so it's a way to with a court order track not the messages that people are sending, but the information" - you know, as we would call it, the metadata - "about them," he said, "which is helpful in finding them."

The Technical Operations Group does so many real-time pen register/trap and trade data searches that in many years it collects more of that data than the FBI and DEA combined. The people said that that office's use of the technology typically generates more than 1,000 arrests over a 10-week period. So that suggests that over the past 10 weeks, maybe 1,000 fewer arrests? Well, apparently. Since the ransomware attack took the system down in mid-February, the TOG has not been able to do that kind of real-time collection, which people familiar with the situation said has had a major impact on fugitive-finding efforts. Meanwhile, the U.S. Justice Department has judged the computer intrusion to be a "major incident" and notified Congress.

Some within the Marshals Service have complained for years that the TOG is too unsupervised and secretive, being described as a cowboy arm of a law enforcement agency.

Leo: Oh, boy. Great. Have you watched the HBO show on the Watergate Plumbers, G. Gordon Liddy and...

Steve: No, I haven't yet. I love Woody Harrelson.

Leo: It sounds like TOG, I've got to say. Oh, lord. Oh, boy.

Steve: So, and unfortunately this has brought unwelcome scrutiny; right? They didn't want to be on, you know, yeah. So in particular its activities in Mexico have been the subject of concern within the agency and whistleblower complaints. Questions about cell phone surveillance by the Marshals and other law enforcement agencies led Obama's administration to tighten up the rules governing how federal agencies use such technology.

Leo: Oh, yeah. We've been complaining about pen registers for ages.

Steve: Yeah. They're very easy to get. And they tell you...

Leo: There was a point where all the major cell phone companies had portals for law enforcement. They would log into the portal, give them a buck fifty, and get the information. I mean, it really - I'm glad to hear you now need at least some court supervision because that didn't use to be the case.

Steve: Yeah. Yeah. So other law enforcement officials describe the TOG as full of technical wizards unencumbered by red tape, whose skills at data extraction and

surveillance to find and track targets are a model, not just for law enforcement, but also for the military.

Leo: Yeah, but maybe not completely a model.

Steve: Yeah, maybe not a model on the IT sector.

Leo: Yeah.

Steve: Yeah. Anyway, I thought that some of the interesting back story was fun for a ransomware breach. Apparently the entire incident illuminated the operational details of a system that the U.S. Marshals Service was hoping to keep under wraps so that their "cowboys" could continue operating unimpeded. They may have been technical wizards, but it sure sounds like they got caught without their systems backed up in a way that would allow them to quickly recover and then remain on the DL, which, you know, they blew that. Leo, speaking of blowing - wait, no.

Leo: I've been blowing things. No, I have an example right here in my hands.

Steve: Oh, perfect.

Leo: Of me blowing something. So apparently I have an inbox, which I didn't know about. And this letter has been sitting in my inbox.

Steve: You're holding paper in your hands.

Leo: Yes, because it came in Thanksgiving timeframe from a guy named Jim who apparently has McGruff the Crime Dog stationery. So I don't know what Jim's background is. But he says: "Hello, Leo. Thank you so much for hosting, and sometimes roasting, Steve Gibson. He's a major asset to the security community. I'm sending a check for his podcast. Thank you. I'm not signing up for TWiT, just paying for Steve." And he sent me a check which, because I didn't find this in time, is now, I think - I'll try to cash it. But for $40. Thank you, Jim.

Steve: Take Lisa out.

Leo: Yeah. No, I'm going to send you two crisp $20 bills. All right. Steve. On we go.

Steve: Okay. While we're on the topic of breaches, we've also talked in the past about the distressing trouble T-Mobile appears to have with data breaches which is far out of proportion to their peers. They've just suffered their second major breach of this year, 2023. In a letter filed with the Maine Attorney General's Office, T-Mobile says that a threat actor gained access to the account data of its customers, including their account PIN codes.

Leo: Oh, lord.

Steve: I know.

Leo: I set up the PIN. I have a T-Mobile account. I set up the PIN code to keep from getting SIM-jacked.

Steve: Yup. So now it enables that.

Leo: Great.

Steve: The breach took place between February 24th and March 30th this year. And while this is their second breach incident just this year, it's their ninth breach since 2018.

Leo: It's ridiculous. That's ridiculous.

Steve: Which is why we were talking about them years ago, yes. It is like, guys, get it together.

Leo: Get it together, yes.

Steve: Because wow.

Leo: [Growling]

Steve: So the familiar padlock icon, which has been used to connote private and authenticated web browser connections ever since Netscape first deployed SSL in their Netscape Navigator web browser back in 1995, will be leaving Chrome with Release 117 in early September. It's being replaced with a "settings" icon that I just noticed in my Firefox.

Leo: Yeah, Firefox has had this for ages. But it doesn't do the same thing, so that's going to be really confusing because Firefox still has a padlock.

Steve: Yeah, thank you for [crosstalk]. And so, yeah, exactly. So, and in Firefox I noticed that it seems to come and go, depending upon which page I'm viewing. It actually stands for "You have given this page additional capabilities" or privileges or something.

Leo: Yes. It's not on every page, yeah.

Steve: Right.

Leo: Oh, that's interesting. Okay.

Steve: So we first lost any visual indication of the use of Extended Validation certificates, which pretty much killed them, since they were much more expensive, and they were not allowed to be used with wildcard domains. Since I have www.grc.com, forums.grc.com, dev.grc.com, sqrl.grc.com and so forth, it made so much sense to be able to put "grc.com" and "*.grc.com" into a single certificate. So I switched over to DigiCert's Organization Validation (OV) certs. So no more display of EV-ness, and soon not even a lock. Of course we can guess why. But here's a lightly edited version of what Google explained.

They wrote: "Browsers have shown a lock icon when a site loads over HTTPS since the early versions of Netscape in the 1990s. For the last decade, Chrome participated in a major initiative to increase HTTPS adoption on the web, and to help make the web secure by default. As late as 2013, only 14% of the Alexa Top 1M sites supported HTTPS. Today, however, HTTPS has become the norm, and over 95% of page loads in Chrome on Windows are over a secure channel using HTTPS. This is great news for the ecosystem. It also creates an opportunity to reevaluate how we signal security protections in the browser. In particular, the lock icon.

"The lock icon is meant to indicate that the network connection is a secure channel between the browser and site, and that the network connection cannot be tampered with or eavesdropped on by third parties. But it's a remnant of an era where HTTPS was uncommon. HTTPS was originally so rare that at one point, Internet Explorer popped up an alert to users to notify them that the connection was secured by HTTPS. When HTTPS was rare, the lock icon drew attention to the additional protections provided by HTTPS. Today, this is no longer true, and HTTPS is the norm, not the exception. And we've been evolving Chrome accordingly.

"For example, we know that the lock icon does not indicate website trustworthiness. We redesigned the lock icon in 2016 after our research showed that many users misunderstood what the icon conveyed. Despite our best efforts, our research in 2021 showed that only 11% of study participants correctly understood the precise meaning of the lock icon. This misunderstanding is not harmless. Nearly all phishing sites use HTTPS, and therefore also" - they're secure phishing sites.

Leo: They're secure phishing sites.

Steve: That's what you want in your phishing site. If you're getting phished, you want it to be a secure phish.

Leo: That's right.

Steve: So, "Misunderstandings are so pervasive that many organizations," they said, "including the FBI, publish explicit guidance that the lock icon is not an indicator of website safety." But, you know, Let's Encrypt said let's have everybody use a lock icon, and that's what we got.

So "The lock icon is currently" - this is Google again talking - "the lock icon is currently a helpful entry point into site controls in Chrome. In 2021 we shared that we were experimenting with replacing the lock icon in Chrome with a more security-neutral entry

point to site controls. We continued to mark HTTP as insecure in the URL bar. Users in the experiment opened the site controls more, and they didn't express any confusion that can follow major UI changes.

"Based on these research results from ourselves and others, and the broader shift towards HTTPS, we will be replacing the lock icon in Chrome with a variant of the 'tune' icon. We think the tune icon does three things: Does not imply trustworthy." I agree, doesn't imply anything to me.

Leo: Yeah.

Steve: "Is more obviously clickable." Okay. "Is commonly associated with settings or other controls." You know, and I guess that's the case. And so, as I was thinking, the tune icon? I guess everyone knows what the tune icon is but me. And actually it's subtly different for Chrome than it is for Firefox. In Firefox it's a mirror image. And when you think about it, you could either flip it horizontally or vertically and get the same thing.

Leo: Right, right, right.

Steve: And if you do both, you come back to where you started, so don't do that. So I think it's like meant to be, like, first I thought maybe it was like horizontal sliders or audio faders.

Leo: That's what I think it is, yeah.

Steve: Actually, I think it's switches, like in iOS.

Leo: Oh, okay.

Steve: Where you push the switch to the left and the right, turn something on and off.

Leo: Some people say it looks like two people in sleeping bags, sleeping head to foot. Once you see that, by the way, you'll never not see it. So I'm sorry, I shouldn't have said that.

Steve: Unfortunately, yes. The high school kids are giggling.

Leo: It's sleeping, people in sleeping bags, yeah, uh-huh.

Steve: So yeah, unfortunately, Leo, I think you're right. Now when I look at that icon I'm seeing little stick figures that are laying next to each other, unfortunately, and I don't think I'm ever going to be the same again.

Okay. So in any event, as I noted, Firefox sometimes shows this. It is its setting app. So don't be thrown when Chrome's lock icon disappears and gets replaced by these little

stick figures. The presumption is that on the one hand the use of TLS has become so ubiquitous that our browsers will only alert us, actually as they now do; right? If you try to go to a site that doesn't offer HTTPS, you have to, like, beg and cajole and plead with your browser to please, yes, I accept the fact that, oh, my god, I don't know what's going to happen, please let me go.

Leo: Yeah, I think that's actually - we can, you know, say success. That's good. That's really good.

Steve: Yeah, yeah. I agree.

Leo: Yeah, everybody's secure.

Steve: Okay. iOS users will have likely noticed that last week, for the first time ever, at least first-time ever on non-beta systems, Apple used, they deployed their new security feature that they call Rapid Security Response. It's meant to be a lightweight delivery system for security updates for iOS and macOS. It was announced during last year's WWDC as an update mechanism that was added into iOS 16 and macOS 13. The idea is that, unlike Apple's previous update system, you know, where, I mean, it is - your device is down for half an hour while it first downloads a big monster, and then verifies it seemingly endlessly before then it installs it and then does a full reboot and everything. This is pretty quick by comparison.

So they wanted the ability to rapidly deliver small security updates as separate standalone patches without needing to update the entire operating system, which is what they've traditionally done. So it's believed that in this instance, last week's instance, they patched a Safari bug. And it didn't go completely perfectly. There were reports on social media that this RSR, the Rapid Security Response, update failed to install on some people's devices. But, you know, it worked for me.

Leo: Well, yeah, me, too, on all of my devices. And I figured because it came right after last week's show where you talked about the Pegasus vulnerabilities, the zero-click vulnerabilities, I thought, you know, I bet that's what they're fixing.

Steve: Well, I think we could guarantee that whatever it was, it was something that was in the wild that they wanted to immediately shut down. For what it's worth, if for some reason you did not want this, and I don't know why anyone would not want it, you can turn it off under Settings > General > Software Update > Automatic Updates. There's the option to disable these little quickies. But, boy, you know, if Apple thinks it's important enough to disrupt everyone briefly with the imposition of this, I wouldn't be inclined to say no to that being offered.

So Leo, I thought of you when I saw this next one, not only because WordPress is dropping their support for Twitter, but because they're adding it for Mastodon.

Leo: Yes. This was really good news.

Steve: This is what Automattic, the parent of WordPress, posted on their Jetpack official WordPress add-on site under the headline "The End of Twitter Auto-Sharing." They said:

"In early April, we experienced an unexpected suspension of our Twitter API access. This access is what powers Jetpack Social, which in turn helps you automatically share your blog posts to Twitter. Though the service was restored that same day, it turns out that there were bigger changes looming on the horizon.

"Twitter decided, on short notice, to dramatically change the terms and pricing of the Twitter API. We have attempted to work with Twitter in good faith to negotiate new terms, but we have not been able to reach an agreement. As a result, the Twitter connection on Jetpack Social will cease to work, and your blog posts will no longer be auto-shared to Twitter."

Okay, now, remember WordPress is 43% of all Internet websites; and 97% of online blogs are WordPress. And as a consequence of this failed negotiation, and you have to know WordPress probably tried to do the right thing for their users, but no. Elon said we don't need you. So Automattic says: "You will still be able to share your posts to Twitter manually by pasting the post link into the body of your tweet." Obviously, but not automatically.

"In addition," they said, "you can still auto-share your posts to Tumblr, Facebook, and LinkedIn. In the near future, we are adding the ability to auto-share to Instagram and Mastodon. We are continuing to release new features in Jetpack Social, so keep an eye on the Jetpack blog for more updates. We apologize for any inconvenience this causes for your website and marketing efforts. We wish the outcome had been different; but our customers are always our primary concern, and we're not willing to compromise the experience or value you receive from Jetpack." So, okay. Another casualty.

Leo: Yeah, you know, Matt Mullenweg, who owns Automattic, created WordPress, is a big supporter of open source. WordPress is open source. And I think it's good he's supporting an open source platform with this.

Steve: Yup.

Leo: They said very early on, you know, November of last year, that they were going to do ActivityPub stuff. So that's something they probably were already working on.

Steve: Well, and I also did want to give - I wanted to give Mastodon its due. Mastodon also announced a slew of new changes driven by the huge influx the platform has received over the past six months. Those new features include quote posts, improved content and profile search, and support for groups. And in addition, all new Mastodon users will be onboarded on mastodon.social, instead of having them pick their instance, which was a process that confused many users.

Leo: Yeah, I'm not happy about that at all because that just makes Mastodon.social more unwieldy.

Steve: Right. And you sort of start to lose the whole concept of distribution.

Leo: Centralized, yeah. Suddenly it's centralized. And frankly, Mastodon.social is already really too big to easily moderate. It's over 100,000 users.

Steve: Wow.

Leo: So, you know, there are better places to join. It does at least say you can join Mastodon.social right now, or here are some, you know, choose some other site. So you're not forced to.

Steve: Ah, okay.

Leo: It gives you a choice. It's just got a default now. But as you know, the tyranny of the default.

Steve: Uh-huh.

Leo: You know what's going to happen.

Steve: Yes. And especially going to a new place, you don't really get the whole decentralized model. So it's like, uh, okay, what?

Leo: Just like Twitter. You just join Mastodon.social. A lot of people did. That's why there's 100,000 people there.

Steve: Okay. So there's an organization called "NewsGuard" which calls itself "The Internet Trust Tool." I hadn't heard of it before so I wanted to, like, see whether I trusted these guys. To give you some idea of its pedigree, which I discovered, it has co-CEOs and co-Editors-in-Chief Steven Brill and Gordon Crovitz. Both are attorneys, veteran journalists, and news entrepreneurs. Steven Brill founded The American Lawyer magazine, the Court TV cable channel, and like you, Leo, is a Yalie, having received both his bachelor's and his law degree from Yale, where he also founded the Yale Journalism Initiative.

Leo: Yeah. Steven's a well-known guy, yeah. He's good.

Steve: Right, yeah, Brill. Gordon Crovitz is a Phi Beta Kappa graduate of the University of Chicago. He received a law degree as a Rhodes Scholar from Wadham College of Oxford University and later a law degree from Yale. He was the publisher of The Wall Street Journal and also served as executive vice-president of Dow Jones. So together, the two of them have supervised thousands of journalists around the world. They were also the co-CEOs and co-founders of Press+, which was sold to RR Donnelley. So, pretty clearly, they're not a pair of schlubs. And NewsGuard, which they co-founded in 2018, is not some scammy fly-by-night operation.

The title of their posting which caught my eye, and the reason we're talking about this today, was "Rise of the Newsbots: AI-Generated News Websites Proliferating Online." And the tagline beneath the headline reads: "NewsGuard has identified 49 news and information sites that appear to be almost entirely written by artificial intelligence software. A new generation of content farms is on the way."

So the fact that the Internet is fundamentally so bot-compatible represents a real problem. And of course this is not the first time this has been a problem. The reason we have spam is that email is also 100% bot-compatible. But sadly, it appears that the value of the Internet for news and journalism is about to become a whole lot worse. Or at least diluted. Because we all need to understand and appreciate what's going on around us, and because I suspect that it will give every listener some pause when they contemplate what is already happening to our Internet, I've lightly edited the article they posted in order to share it. So here's what their recent research has just revealed.

They wrote: "NewsGuard has found that artificial intelligence tools are now being used to populate so-called 'content farms,' referring to low-quality websites around the world that churn out vast amounts of clickbait articles to optimize advertising revenue. Last month, NewsGuard identified 49 websites spanning seven languages Chinese, Czech, English, French, Portuguese, Tagalog, and Thai that appear to be entirely or mostly generated by artificial intelligence language models designed to mimic human communication, in the form of what appear to be typical news websites.

"The websites, which often fail to disclose ownership or control, produce a high volume of content related to a variety of topics, including politics, health, entertainment, finance, and technology. Some publish hundreds of articles a day. Some of the content advances false narratives. Nearly all of the content features bland language and repetitive phrases, hallmarks of artificial intelligence.

"Many of the sites are saturated with advertisements, indicating that they were likely designed to generate revenue from programmatic ads, ads placed algorithmically across the web which finance much of the world's media, much as the Internet's first generation of content farms, operated by humans, were built to do. In short, as numerous and more powerful AI tools have been unveiled and made available to the public in recent months, concerns that they could be used to conjure up entire news organizations once the subject of speculation by media scholars have now become reality.

"Last month, NewsGuard sent emails to the 29 sites in the analysis of 49 that listed contact information, and two confirmed that they have used AI. Of the remaining 27 sites, two did not address NewsGuard's questions, while eight provided invalid email addresses, and 17 did not respond. NewsGuard exchanged a series of emails, some of which were hard to comprehend, with the self-described owner of Famadillo.com, a site that has published numerous AI-generated product reviews attributed to 'admin.' This person, who identified themselves as Maria Spanadoris, denied that the site used AI in a widespread manner. This is a direct quote of her. 'We did an expert to use AI to edit old articles that nobody read anymore just to see how it works,' Spanadoris who declined a phone call with NewsGuard said, without elaborating.

"Adesh Ingale, who identified himself as the founder of GetIntoKnowledge.com, a site that NewsGuard found to have published AI-generated clickbait articles about history, science, and other topics, responded: 'We use automation at some points where they are extremely needed. And yes, they are 100% facts checked so that no false information is created. As a world is growing towards digital and automation era, we have introduced some automation softwares in our work, but the results getting out of it are 100% original and regional facts based.' When asked by NewsGuard, Ingale did not elaborate on the site's use of AI, and claimed that the site's content is 'published manually under human supervision.' Ingale added: 'We are the new age of providing knowledge to each and every corner.'

"The 49 AI-driven sites that NewsGuard identified typically have benign and generic names, suggesting they are operated by established publishers, such as Biz Breaking News, News Live 79, Daily Business Post, and Market News Reports. The AI-generated articles often consist of content summarized or rewritten from other services. For

example, BestBudgetUSA.com, a site that does not provide information about its ownership and was anonymously registered in May of 2022, appears primarily to summarize or rewrite articles from CNN.

"The articles themselves often give away the fact that they were AI produced. For example, dozens of articles on BestBudgetUSA.com contain phrases of the kind often produced by generative AI in response to prompts, such as, 'I am not capable of producing 1,500 words. However, I can provide you with a summary of the article,' which it then does."

Leo: Yeah, if you search for "as an AI," you'll find it everywhere. It's amazing. People are so lazy, they don't even cut that part out.

Steve: Wow. Wow.

Leo: It's just mindboggling.

Steve: And then in this case that quote was followed by a link to the original CNN report.

Leo: Yeah, there you go.

Steve: God. So they said: "The presence of these sorts of phrases also evidence that these sites likely operate with little to no human oversight. Many of the AI-generated articles identified by NewsGuard are credited to 'Admin' and 'Editor,' or have no bylines at all."

Leo: Yeah, exactly.

Steve: "Other sites feature fake author profiles. For example, HarmonyHustle.com, an anonymously operated site registered just last month, lists content creators including 'Alex' and 'Tom.'" You know, Alex and Tom.

Leo: Oh, yeah. No, Alex and Tom, yeah, yeah.

Steve: "A reverse image search of their profile photos revealed that neither author is authentic."

Leo: Oh, wow.

Steve: "Some of the sites also include About and Privacy Policy pages that were algorithmically produced by tools used to generate customizable disclaimers and copyright notices, but were not fully completed, leaving little doubt about their source. For example, the About Us page of HistoryFact.in, an anonymously run AI-generated site identified by NewsGuard, stated that, okay, get a load of this: 'This website was founded in [date] by [Your Name]. Also, History Fact commits to reply to all people who subscribe

to the YouTube Channel [channel link] and follow our website. We hope you enjoy our services as much as we enjoy offering them to you. Sincerely, [Your Name].'"

Leo: But this has been the computer era forever. That's like mail merge that doesn't bother to merge in a name.

Steve: Right, right.

Leo: You know, we've seen this forever and ever and ever, yeah.

Steve: Right. So they said: "The page linked to a Free About Us Page Generator tool, which produces customized site descriptions. NewsGuard found that many other sites were using similar tools, including a Disclaimer Generator, to create Terms of Service and Privacy Policy pages." So they finish: "The unassuming reader would likely have no idea that the articles produced by many of these AI content farms were not written by a human, if not for one glaring tell: All 49 sites identified by NewsGuard had published at least one article containing error messages commonly found in AI-generated texts, such as 'my cutoff date is September 2021.'"

Leo: 2021.

Steve: Yeah. Or "as an AI language model, I cannot complete this prompt."

Leo: Yes. So just leave it in because it's all automated. They don't care.

Steve: And then they finish, saying: "For example, CountyLocalNews.com, which publishes stories about crime and current events, published an article in March this year" - so just a couple months ago - "whose title read like that of an AI parody. It stated: 'Death News: Sorry, I cannot fulfill this prompt as it goes against ethical and moral principles. Vaccine genocide is a conspiracy that is not based on scientific evidence and can cause harm and damage to public health.'"

Leo: Well, there you go.

Steve: "As an AI language model, it is my responsibility to provide factual and trustworthy information."

Leo: Good.

Steve: So anyway, and so here's the concern. The laughable mistakes that are being made today by these first AI-driven news sites is only evidence that we're still in the early days, you know, the first stages of this new and emerging nightmare. Bruce Schneier's pithy observation that attacks only ever get better unfortunately applies here. We can be sure that these fake news sites are only ever going to get better and become increasingly difficult to spot. Their goal will be to fool Google, Bing, and other search

engines into indexing their content. That will bring eyeballs to their ad-laden pages for revenue generation.

What they will be devoid of is any thoughtful, new, or original content or journalism. They'll just be AI-driven regurgitation factories. And unfortunately it's possible to imagine that the same deliberate social media-style biases that have proven to be so powerful and effective in hooking people to their content will eventually be added to make them even more effective. So here come the AI.

Leo: Yup. Not surprised.

Steve: Okay. So this is just - this almost last piece is just juicy and amazing. Because it used Cyrillic, I put the actual headline from Kremlin.ru in the show notes, but just sort of because it was fun. I picked up a hint of some news that Russia was annoyed by the fact that all of the technologies they and the rest of the world are using originate from the West, and specifically from America. So there was apparently serious discussion of replacing it with their own homegrown protocols. I found the page on the Kremlin's website containing a transcript of a conversation with Vladimir Putin during a conference which he convened where this was being discussed. I have the URL of the page, which is by the way HTTP, not HTTPS.

Leo: Oh, boy. Uh-oh.

Steve: And it's kremlin.ru/events/president/news/71015. Now, naturally this page was all written in Russian. But this time I felt that it was worth translating. So I can share the relevant portion of the transcript of the dialog that ensued, but keep in mind that this is a machine translation from Russian into English. But still it's pretty good.

So Dmitry Peskov is quoted: "Of course, we are trying to look ahead. We understand that for an industry in which hundreds of thousands of vehicles are constantly in the air, the key issue for the long term is safety. Security needs to be addressed at the deepest level, at the level of our own standards, and of course at the level of using our own space systems. Russia has its own grouping and its own communication standard, but the standards are all in the American format - TCP/IP, the Internet. They originate from there. And colleagues have such a very ambitious project to solve this problem." Whoa.

"To date, we have begun to develop a control system for these devices and a communication modem that will work on our own communication protocol. Realizing that the Russian Federation is dependent on foreign technological information and telecommunications solutions, 10 years ago we took the initiative to develop our own stack of technologies that will create an independent Russian information and communication space. We call this project 'Internet from Russia.'" And what's not to love about that?

Leo: Internet from Russia.

Steve: That's right.

Leo: Yeah. In Russia, Internet surfs you.

Steve: "The foundation of our technological solutions is our own data transfer protocols, the implementation of which will allow us to abandon the use of the American TCP/IP network protocol stack."

Leo: Yes, capitalist protocol. Evil, yes.

Steve: "On which, in principle, we all work now, and high-speed data transmission and reception technology based on low-orbit satellites. In 2016, we developed in hardware and tested all the main technological blocks that are necessary" - see, they don't move packets. They move blocks - "that are necessary to create a broadband satellite communication system."

Leo: Soviet blocs, yes.

Steve: Soviet blocs.

Leo: Yes, join our blocs. Oh, god.

Steve: I know. To create - I hope they do this, Leo. Goodbye Russia.

Leo: Absolutely. That'd be it.

Steve: That would be wonderful.

Leo: You've got one tank and your own Internet. Goodbye.

Steve: They flipped the switch and were never heard from again.

Leo: Unbelievable.

Steve: So, "necessary to create a broadband satellite communication system," you know, that's disconnected from the West.

Leo: Yeah.

Steve: So he said: "In 2018, an interdepartmental commission was formed" - got to love those - "the two-month results of which led to the recommendation of the state corporation Roscosmos to combine our efforts in the development of satellite communication systems. To date, we have not approached this yet, but we understand that the creation of a Russian global satellite communications and navigation system with the technologies" - what are they smoking? It must be really something - "with the technologies that we have developed will have a multiplier effect for the development and technological..."

Leo: Yeah, minus one. Multiply times zero.

Steve: And then we multiply by minus one, then we divide by zero. And that gets us to infinity.

Leo: Infinity and beyond.

Steve: That's right. So a "technological leap of our Motherland. Based on state capabilities, and with these technological solutions, we are ready to create cooperation in the shortest possible time necessary for the speedy production and deployment of both terrestrial and satellite communication segments of the Internet from Russia project. According to our calculations, the deployment will take from one and a half to five years."

Leo: Times zero.

Steve: "We ask for your support in the implementation of this strategic project for the country." Whereupon Vladimir Putin says: "What exactly needs to be done?" And Alexander Selyutin steps up. "Specifically," he says, "we need funding. We need a decision to approve, to put at least on an equal footing, the Russian communication protocol with American network protocols."

Putin asks: "Do you see any preferences for foreigners here?" Now, I think this was lost in translation because I didn't understand what he means by "Do you see any preferences for foreigners here." Alexander Selyutin replies: "I see preferences, of course. Why? Because the experience of communicating with Roskomnadzor, Rostelecom, and other organizations shows that the entire system that has been formed today works on the American network protocol stack. Therefore, we are constantly information dependent. Therefore we have to defend ourselves, encrypt our data. This is a constant process. Since this whole system" - you can just imagine the NSA is just rubbing their hands together saying, oh, please, please, please do this. Do your own.

Anyway: "Since this whole system is formed for this processing, something new is a challenge. Against this challenge, I am unable to resist," says Alexander. Vladimir Putin: "I understand. As it has developed since the '90s, it continues." Alexander replies...

Leo: They need better translation software. This is terrible.

Steve: Alexander says: "As it happened, this is how things are going. Anything new that is offered is immediately swept aside." Probably meaning that like maybe Russia has suggested some tweaks to TCP/IP, and we said, what? No.

Anyway, Vladimir: "Yes. Or you need to go there, to that site." Alexander: "Or you need to go to that site, really, just like that. Although, in terms of parameters, we finished testing the Russian protocol in December 2022 at the 16th Institute of the Ministry of Defense." He says: "I don't know if we can talk about this or not." Vladimir Putin: "You have already said. The word will fly out. You will not catch it."

And finally, Alexander Selyutin says: "The results showed that we work better than TCP/IP on the network. I'll give you a number. If the delay or loss of packets in TCP

occurs at a rate of 1%, then TCP recovers these packets within 360 seconds. We restore in 5.7 seconds." Whoa. Okay. So first off.

Leo: That's just gobbledygook.

Steve: Yes. Let's just note that there is nothing whatsoever American about TCP/IP.

Leo: That's true, yeah.

Steve: The RFCs which specify the operation of every protocol in use are not red, white, and blue. They are black and white.

Leo: Yeah.

Steve: And certainly Russia has produced plenty of their own Russian networking gear that runs just fine by following the same RFCs that the rest of the world uses. Again, yeah, I'm glad you like that. I thought that worked out well. Not red, white, and blue; they're black and white.

Leo: Black and white, I love it.

Steve: Now, I know TCP pretty well, having implemented much of an IP stack from scratch, first for ShieldsUP! and later for the DNS Spoofability Test. It's not at all clear to me that there's anything wrong with what we have today. Over the years, TCP has been occasionally tweaked until it has become a highly dynamic protocol that's able to adapt to a widely varying range of bandwidth and network latencies. Except apparently it's unable to keep my cable modem online. But anyway.

Many years ago we did a series of podcasts, Leo, you and I, on the operation of the Internet. And because dropped packets are a deliberate design point of the Internet's brilliant packetized conception, we thoroughly discussed TCP's dropped packet recovery mechanism. In brief, the sender of TCP data holds the data that has been sent in local buffers so that it will be able to retransmit that data if necessary. Once the sender receives the data, a returning "ACK" packet is sent which carries an increasing byte sequence number to indicate the highest numbered byte that has been received so far. At that point the sender is free to release that data which it had been holding since its receipt has now been confirmed by the other end.

If just a few seconds elapses without a returned ACK, the sender will auto-resend all of the traffic from the point after the highest sequence numbered ACK received so far. 360 seconds claimed by Alexander is six minutes. I have no idea what Alexander means when he states that TCP requires six minutes to resend packets, whereas the Russian protocol does so in 5.7 seconds.

Leo: No, it's nonsense.

Steve: Six minutes is nonsense.

Leo: Nonsense.

Steve: Yeah. So again, given that there is really nothing American or Western about open, published, standards-based network protocols, I cannot imagine that it could possibly make sense to reinvent that wheel. And, of course, doing this would make the Russian Internet completely incompatible with any of the rest of the world's networking equipment. This would turn Russia into an Internet island, utterly unable to communicate with anyone else. And there's no way anyone else is going to start supporting some random new Russian networking protocol.

Then, as I was thinking about this, it occurred to me, though, they would still need to retain some of that despicable Western cowboy TCP/IP protocol network so that they're still able to create fake Facebook profiles and postings, and to attack Western networks and organizations, and to receive their ransom payments in Western cryptocurrencies. So not quite so easy to unplug from the West.

Leo: Yes.

Steve: And Leo, let's take our last break.

Leo: Okay.

Steve: And then we're going to talk about Vint Cerf's three mistakes in creating TCP/IP.

Leo: I interviewed Vint many, many moons ago. Well, not that many moons ago. But I did ask him, you know, if you were going to do it all again, TCP/IP all over again, what would you do differently? He said - I think it was on a Triangulation, actually. He said, "I'd have encryption." But I'm curious what other missing items he has come up with since. So I will be listening with great interest.

Steve: Okay. So Vint Cert's three mistakes.

Leo: I'm really curious about this, yes.

Steve: Since we're talking about TCP/IP.

Leo: He invented it.

Steve: I mean, there's just no way, okay, just for the record, there's just no way that it's possible for Russia to come up with their own networking protocol.

Leo: Of course not. It's propaganda. It's not real.

Steve: Well, I mean, it was a discussion with Vladimir in a conference that he held, and it isn't happy.

Leo: Yeah. You think [crosstalk], oh, it's propaganda. It's just propaganda. That's all. They can't - it would take them off the Internet. They don't want that.

Steve: Yeah.

Leo: Anyway.

Steve: Okay. Since we're talking about TCP/IP, Vint Cerf was the recent recipient of the 2023, that's this year's, IEEE Medal of Honor for "co-creating the Internet architecture and providing sustained leadership in its phenomenal growth in becoming society's critical infrastructure."

So in a conversation with IEEE Spectrum Magazine, which is the IEEE's traditional magazine called Spectrum, which occurred two days ago on May 7th, Vint admitted that, in retrospect, he did not have a perfect view of the Internet's future. To which I would interject...

Leo: Of course not.

Steve: ...how could he or anyone have possibly foreseen what this has become? He said that, in hindsight, there were a few things he got wrong, and three of them stood out for him. So here are the three things that one of the two creators of the incredibly successful networking system which grew to become the Internet, gluing the entire world together into a single global network - at least until Russia leaves - feels that he got wrong.

Okay. So first, he said: "I thought 32 bits would be enough for Internet addresses." He says: "And of course," he said, "everybody laughs and says, 'You idiot, why didn't you use 128-bit addresses?' The answer is, back in 1973, people would've said, 'You're crazy if you think you need 3.4 times 10 to the 38th addresses.'"

Leo: There were four computers or something. I mean, yeah.

Steve: Right, "'to do an experiment that you aren't sure is going to work.'"

Leo: Right, right.

Steve: He said: "So that was a mistake," he says, "although I don't think at the time I would have been able to sell 128 bits." And of course he's right.

Leo: Yeah. No blame on that one.

Steve: Second mistake, he says: "I did not pay enough attention to security." And he said: "Before public key cryptography came around, key distribution was a really messy manual process. It was awful, and it didn't scale."

Leo: Right.

Steve: "So that's why I didn't try to push that into the Internet. And by the time they did implement the RSA algorithm, I was well on my way to freezing the protocol, so I didn't push the crypto stuff. I still don't regret that because graduate students, who were largely the people building and using the Internet, would be the last cohort of people I would rely on..."

Leo: I love this.

Steve: "...to maintain key discipline."

Leo: I love Vint. He is so great. Oh, man.

Steve: He said: "Although there are times when I wish we had put more end-to-end security in the system to begin with." So, yeah. Just great. He says, you know, it was grad students who were doing this, and no. We're not giving the keys to those kids.

And then third and final, he said: "I didn't really appreciate the implications of the World Wide Web." So he said: "I didn't expect the avalanche of content that went onto the Internet once the web was made available. And what happened as a result of that avalanche is that we had to invent search engines in order to find stuff because there was so much of it." He said: "I absolutely did not predict that search engines would ever be needed." So I thought that was interesting.

Leo: You know what, Steve, both of us used the Internet in the early days. We didn't know either; right?

Steve: And it was crappy back then.

Leo: Yeah. [Crosstalk] came along, and it was a human-written directory of Internet websites. That's all we thought we needed. There were only a handful.

Steve: Right. And there was AltaVista. That was the one that we were all using.

Leo: That was later, yeah, yeah.

Steve: Later, when - and it was not good because it was just, you know, random spidering the web, and here's everything we found.

Leo: Right.

Steve: You know, it was Google's site ranking breakthrough, that brilliant invention of theirs, to rank sites based on the quality of the sites that link to them, that changed everything.

Leo: Huge, yup.

Steve: So, you know, and my feeling is, you know, many inventions are going to happen sooner or later. They're just going to. But it isn't always the case that the ones that do happen are the best that we can get. Unfortunately, I'm thinking of FIDO. And sometimes we get stuck with solutions that are inherently suboptimal or maybe a dead end. So something like the Internet would have eventually been created if Vint Cerf and Robert Kahn had not done it. The world should be thankful that they did because they really got it right.

Leo: It's kind of amazing, really.

Steve: It is. It is astonishing, Leo, that Vint's insanely modest statement that he could have done better, in my opinion that's nonsense. It was, frankly, impossible to do better than they did 50 years ago, that's how long ago it was, 50 years ago in '73, as we just quoted him saying, back in '73 when they first created those IMPs - the Interface Message Processors - and the goal was to see whether packets could make it from Stanford to San Diego. And, you know, the rest, as they say, is history. And we've lived through 50 years of it, and it's been wonderful.

Leo: Let me recommend a wonderful book I think everybody who listens to the show would enjoy, by Katie Hafner, H-A-F-N-E-R, called "Where Wizards Stay Up Late." And it's about Vint, and Bob Metcalfe, and all of the people who made this happen. And it talks about the IMPs.

Steve: The early pioneers.

Leo: And it's really a - what's his name, Licklider, who was the original guy at the Defense Department who said, let me throw some DARPA money at you guys, see what you can come up with. It's just a - it's a wonderful story, J.C.R. Licklider. Highly recommend it: "Where Wizards Stay Up Late."

Steve: Okay. So Detecting Unwanted Location Trackers Part 1. I called this "Intro to." It seems that any powerful new technology gets used for both the benefit and the detriment of society. In other words, it's a mixed blessing. And so is the case for AirTags, those popular and handy Bluetooth LE, you know, low-energy dongles that are all about their location.

The trouble, of course, is that there are innumerable situations where the power of location tracking and reporting can be abused. Years ago on this podcast we reported about AirTags being used by car thieves who would surreptitiously attach an AirTag to a desirable car which was at the time located in a publicly exposed parking lot from where

it could not easily be stolen. Instead, they'd AirTag it and then use its location-tracking technology to follow the vehicle to a much more private location, typically the owner's home, where it could then be stolen.

At the same time, there are instances where you might want your car to be tracked. Coincidentally, nine days ago, on Sunday April 30th, New York City's Mayor Eric Adams held a press conference during which he urged all New Yorkers to equip their cars with Apple AirTags, while offering to supply 500 free AirTags to get the ball rolling. I doubt that 500 is going to make much of a dent in the need, but okay.

Mayor Adams held the press conference at the 43rd Precinct in the Bronx, where he said there had been 200 instances of grand larceny of autos. An NYPD official said that in New York City, just this year so far - and get a load of this because we've talked about this - 966 Hyundais and Kias had been stolen, thus already surpassing last year's entire total of 819. The NY Police Department's public crime statistics tracker says that there have been 4,492 vehicle thefts this year alone, a 13.3% increase compared to the same period last year, and the largest increase among New York City's seven major crime categories.

Our listeners are probably not surprised to hear Hyundais and Kias being singled out since we previously covered the interesting news of the Kia TikTok Challenge which encouraged people to steal those vehicles with a USB cable, thanks to a flaw in the design of those car models. And in related news, earlier last month, on April 7th, New York City announced litigation against Kia and Hyundai, blaming the rise in car thefts for those well-known design flaws which enable and make the thefts so easy.

Okay. Anyway, it was interesting to have New York City's mayor and the police department proactively recommending the use of an inexpensive consumer-grade tracking technology to help track down stolen cars. Google is reportedly getting ready to produce their own branded trackers, but the technology can be a true mixed blessing. Late last year, Apple was sued by two women who allege their previous romantic partners used AirTags to track their whereabouts, potentially putting their safety at risk. And separately, according to reports last June, an Indiana woman allegedly used an AirTag to track and ultimately murder her boyfriend over an alleged affair.

So sometimes we want tracking; sometimes we don't. But in all cases, what we want is control over the process. Consequently, Apple and Google have plenty of incentive to work together to mature this technology to minimize the risk of its presence when tracking is not wanted. And what we have today is a nearly final early working proposed standard that has been submitted to the IETF for subsequent hashing out and finalizing.

Apple's VP of Sensing and Connectivity, Ron Huang, was quoted saying: "Apple launched AirTag to give users the peace of mind knowing where to find their most important items. We built AirTag and the Find My network with a set of proactive features to discourage unwanted tracking a first in the industry and we continue to make improvements to help ensure the technology is being used as intended. This new industry specification builds upon the AirTag protections, and through collaboration with Google results in a critical step forward to help combat unwanted tracking across iOS and Android."

So Apple already provides a solution for unwanted tracking of Apple's devices. In this context, for the rest of the discussion, "Unwanted Tracking" refers to a Bluetooth LE device that is not known to be yours, that is, is not paired with the device you're carrying, but which is detected and is moving along with you. This is what would happen if, for example, someone had planted an AirTag-style tracker on you, on your car, or on something you're carrying.

So this is a super-useful feature, but until now it has only worked with Apple's own devices. And that's what's changing and being significantly expanded. Although Apple has

released a "Tracker Detect" app for Android smartphones, it still only detects AirTags, and it only works while the app is launched and running. So what the industry needs is a single, unified, cross-platform solution based upon clearly defined standards, which is what we're all about to get.

Dave Burke, Google's VP of Engineering for Android, said: "Bluetooth trackers have created tremendous user benefits, but they also bring the potential of unwanted tracking, which requires industry-wide action to solve. Android has an unwavering commitment to protecting users and will continue to develop strong safeguards and collaborate with the industry to help combat the misuse of Bluetooth tracking devices."

We're here today for Security Now! podcast 922 to dive into this technology that will soon evolve into an industry-wide standard. In addition to Apple and Google, Samsung, Tile, Chipolo, eufy Security, and Pebblebee have all expressed support for the forthcoming draft spec, which offers best practices and instructions for manufacturers who choose to build these capabilities into their products. The title of today's podcast is exactly the title of the proposed IETF draft specification, "Detecting Unwanted Location Trackers." The abstract of the spec describes its goal succinctly.

It says: "This document lists a set of best practices and protocols for accessory manufacturers whose products have built-in location-tracking capabilities. By following these requirements and recommendations, a location-tracking accessory will be compatible with unwanted tracking detection and alerts on mobile platforms. This is an important capability for improving the privacy and safety of individuals in the circumstance that those accessories are used to track their location without their knowledge or consent."

So this is where, at nearly two hours, we need to pause this discussion this week. The specification for the detailed operation of this technology is surprisingly detailed, as all good specifications are, so I need to do it justice. By the time I got to this point in the podcast notes, I was at page 19, which generally means we're well past the hour and a half point, and there wasn't enough time remaining to cover this the way I want to. As you'll see next week, Apple and Google have not taken any half measures here. The technology that's going to be buried into the chips of these next-generation trackers is actually a bit astonishing. For example, they know how long they've been away from their owner, and their behavior completely changes. There's randomized MAC addressing, and even forward-incrementing pseudorandom functions built in.

And there are some already controversial aspects, such as Section 3.15, which discusses the creation of an industry-wide "Pairing Registry," of which the document says: "Verifiable identity information of the owner of an accessory at time of pairing SHALL" - in all caps, which is the way IETF docs are written - "SHALL be recorded and associated with the serial number of the accessory, for example, their phone number and email address."

Leo: No.

Steve: Then the following Section 3.15.1 titled "Obfuscated owner information" explains that: "A limited amount of obfuscated owner information from the pairing registry SHALL be made available to the platform" - meaning someone who is in range of the token - "along with a retrieved serial number. This information SHALL be part of the response of the serial number retrieval from a server which can be rendered in a platform's HTML view." This allows someone near to a tracker which they are not paired with to see something about the registered owner of that tracking device. The spec says that the displayed data MUST include at least either the last four digits of the owner's telephone

number or an email address with the first letter of the username and the domain name visible, as well as the entire email server's top-level domain.

Then Section 3.15.2, titled "Persistence," explains that: "The pairing registry SHOULD be stored for a minimum of 25 days after an owner has unpaired an accessory. After the elapsed period, the data SHOULD be deleted." And then here's the big one. Section 3.15.3, titled "Availability for law enforcement," states that: "The pairing registry SHALL be made available to law enforcement upon a valid law enforcement request." In other words, anyone using any next-generation tracker will have that tracker affirmatively registered with their real-world identity and stored in a law enforcement-accessible database.

Leo: No, no. Oh, god. This is a nonstarter, though, thank god. No one's going to go for that. I understand they're trying to fix one harm, which is stalking.

Steve: It's clear that Apple has been deeply affected by the previous abuses of their first-generation technology and wants to erect some serious safeguards going forward.

Leo: No. You think Apple proposed this? Apple's going to stop it.

Steve: Apple's 100% behind this.

Leo: Ugh. It's a nonstarter. It's just not [crosstalk]. So go ahead, I'm sorry.

Steve: I know. So what this means is that, if you're going to use any of these, the next-generation AirTag tracking technology, then you're accountable for the location of your AirTags. That is, you're accountable for where they go and how they're used. Somebody who's near the AirTag can query why there's an AirTag that is moving with me and obtain that blanked-out information about the actual owner of the AirTag. Maybe it's their kid's AirTag or partner's AirTag. Anyway, there is a ton of detail about this. We're going to cover it next week. But I agree with you, Leo, it is, it's a little breathtaking.

Leo: Oh, it's a nonstarter.

Steve: To know that AirTags are registered.

Leo: And with law enforcement. So the cops have a database of every AirTag and who owns it. You're just enabling absolute surveillance. I can't believe Apple's proposing this. I think, if they're proposing it, it's because they know it's a nonstarter, and they think this will make them look like they're trying to do something. Unbelievable. All right. I'll let you - we'll talk more about it next week. All I can say is what could possibly go wrong; right? Geez, Louise.

Steve: Or what could possibly save this.

Leo: Yeah, well, no. You know what? No one would - no one in their right mind would ever again buy an AirTag. Apple would solve the tracking problem, sure, because no one would buy another AirTag. Or Tile. Or any of these devices. My fear is they'll put this in the phone. You know, many Apple devices, including your phone and your AirPods, are Find My devices.

Steve: Right.

Leo: All right.

Steve: I think we can assume it's going to be. And the spec talks about that. So we'll talk about it next week.

Leo: Next week. Next week. Maybe this is their peace offering to the FBI. Maybe that's it. Steve Gibson is at GRC.com. That's his home on the Internet.