SECURITY NOW!

Transcript of Episode #898

Wi-Peep

Description: This week we note that Firefox moved to v107 and that Google recently reached a nearly $400 million user-tracking settlement. Red Hat has started cryptographically signing its ZIP distributions, the FBI purchased the nefarious Pegasus spyware, and Greece paid 7 million euros for the similar Predator spyware. Passkeys have a directory listing sites where they can be used, the OMB has decreed a quantum decryption deadline, and 33 U.S. state attorneys general have asked the FTC to get serious about online privacy regulation. We have some engaging listener feedback, and SpinRite is finally a day or two away from starting its final testing. And we're going to wrap up by examining some chilling research which allows the physical location in space of every WiFi device within range to be accurately determined by someone walking past or flying a tiny drone.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-898.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-898-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here with lots to talk about. Red Hat cryptographically signing its ZIPs. How do you do such a thing? We'll talk about the FBI. Apparently they tried to use Pegasus. How legal is that? And then we're going to talk about Wi-Peep, a new way to map WiFi access points or, more threateningly, to track people using WiFi devices. All that and more coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 898, recorded Tuesday, November 22nd, 2022: Wi-Peep.

It's time for Security Now!, the show where we cover your security, your privacy, your online exploits, your offline deploits, with this guy right here, Mr. Steven Gibson. Hi, Steve.

Steve Gibson: Leo, great to be with you again for...

Leo: As always, good to see you.

Steve: ...what is this, the pre-Thanksgiving episode.

Leo: It is.

Steve: Yeah.

Leo: And we're almost in the 900s, which is a scary place to be.

Steve: Actually, it was interesting because when I got Elaine's transcript last week she said this was 897. And she reminded me, she said, okay, that means that we are 102 episodes from 999, and there are 51 episodes per year because we skip one for the holidays.

Leo: She's paying attention.

Steve: Which means exactly, precisely two more years of Security Now!.

Leo: Wow. I'll put that in my calendar.

Steve: Oh, honey, I'm not going to let you forget.

Leo: By then you might say, oh, I'd like to keep doing this.

Steve: You know, Leo, I may have the hang of it by then. So it'd be like not that big a deal.

Leo: Well, as somebody who just quit the radio show after 19 years of doing that, I can kind of understand. After a while you get to a point where it's like, you know, I've done everything I'm going to do.

Steve: And imagine, now, I would say you could sleep in on Saturday except that the show didn't start till 11:00. So if this really changes your sleeping habits, then we have a different problem.

Leo: I get to do stuff on Saturday, which is, I mean, I've worked weekends for 19 years. That's a long time.

Steve: Well, yes. In fact, what was happening was you were only working six hours, two days, Saturday and Sunday for three hours. And then you and I were meeting once a month up in Toronto. You were spending four days up there to record 15 Call for Help shows.

Leo: I'm getting PTSD just hearing about it. That's just crazy.

Steve: And you had three unfilled weekday weeks. And so you said, you know...

Leo: Let's do some shows.

Steve: I've got a little time on my hands here.

Leo: What was I thinking?

Steve: Well, aren't you glad now that you have a podcast network?

Leo: I am. I've been telling people this is the first time I've not been working for anybody in my whole working life. I'm working for myself for the first time ever, something you know a lot about.

Steve: Well, except now you have a wife. So, you know.

Leo: Well, as Patrick Norton once told me, because I said, "I don't want to work for the man," he said, "Leo, there's always a man." And in this case the man is a woman, but still. No. We're partners. But it is kind of interesting that I've been an employee, a W-2 employee since I was 16 years old. So that is a big change. I think the podcast thing might work out, that's all I'm saying.

Steve: It might turn out to be something.

Leo: I don't need to keep this job anymore, yeah.

Steve: It might. So we're going to note this week many things. We've got a new version of Firefox. Google recently reached a nearly $400 million user-tracking settlement. We've got some interesting legislative things to talk about during this next couple hours. Red Hat has started cryptographically signing its ZIP distributions. It's like, what? You can sign a ZIP? Well, not really. The FBI purchased, turns out, the nefarious Pegasus spyware, just to kind of see what it's about. Uh-huh. Greece paid 7 million euros for a similar spyware called Predator. Passkeys has a directory listing the sites where they can be used, so that will be exciting. The OMB, the U.S. Office of Management and Budget, has decreed a quantum decryption deadline.

Leo: Oh.

Steve: Oh, yes. And of course we're all going to pay attention to that. Also 33, speaking of paying attention to the FTC, 33 U.S. state attorneys general have asked the FTC to get serious, my friends, about online privacy regulation. We'll see how that turns out. We've got some engaging listener feedback. And SpinRite is finally a day or two away from its final testing to begin.

Leo: What?

Steve: Yeah, it's done. I'll explain. I'll explain. I have a couple drives here. There's three drives which are weird, and you should not write to these drives. So I'll explain about that. And then we're going to wrap up by examining some chilling research which allows the physical location in 3D space of every WiFi device within its range, like within a multistory building or whatever facility, to be accurately located within a meter or so by someone simply walking past or flying a tiny drone, for about 20 bucks. So that's the Wi-Peep thing. So we're going to talk about all this. And we have a Picture of the Week that had you almost falling off your chair.

Leo: It was pretty funny.

Steve: Yeah, it's a pretty good one.

Leo: It was cute. I liked it.

Steve: So I think a good podcast for our listeners.

Leo: All coming up on this fine 898th edition of Security Now!. That's kind of amazing, isn't it.

Steve: 898!

Leo: Agh. Well, it's funny because the last Tech Guy show is December 18th, and it's going to be I think Episode 1955. I'm one shy of my birth year. And I thought, if I can just do - or actually it'll be 1954, and then the Best Of will be 1955. If I could just do one more. That's okay. That's okay. I'm ready, my friend, for the Picture of the Week.

Steve: Okay. So for those who are not video-connected here, as always I have to explain this. We have a flatbed transport vehicle, a flatbed trailer sort of thing. And it looks like there's on the left is a sort of a rust-colored red container where someone probably said, hey, we need you to pick up some dirt, so bring a container, and we've got some dirt for you. Well, apparently the container that they brought was too small.

Leo: Yeah, couldn't fit.

Steve: Because it's about one third of the back of this trailer, this flatbed. And the rest of it has been piled up with the overflow dirt that didn't fit in the container. Now, in a sane world they would throw a tarp over this whole thing, right, and lock the tarp down. But maybe they didn't have a tarp? Anyway, some apparent rocket scientist here decided, well, you know, I've got to do something, right, because I've just got this exposed dirt, the big pile on the back of this trailer. So they used what they had. They threw about a two-inch diameter belt across the top of the pile of dirt which is about, I don't know, it covers maybe 4% of the pile, the rest of it exposed to the air. Now, you can sort of see also on one side, Leo, facing us, there is like, looks like the strap was somewhere else initially because you can sort of see that some of the dirt was flattened on the side there.

Leo: Oh, yeah. It's moved a little bit, yeah.

Steve: So it looks like, well, it looks like maybe the strap was originally anchored on the slot in the trailer one notch further forward.

Leo: I think you've spent too much time looking at this.

Steve: So there was, yes, there was - are you telling me there's a reason SpinRite took three years to get...

Leo: No, I think Logan 5 in our chatroom may have come up with something. It's not to prevent slippage. It's to prevent theft.

Steve: Oh. It's brilliant. You wouldn't want someone to steal the dirt.

Leo: Steal your dirt. Don't, don't steal my dirt, man.

Steve: So this is like that pole that we saw that had the bike lock around it.

Leo: Right, yeah.

Steve: Where it didn't - it indicated an intention without actually providing any enforcement.

Leo: I love it. Oh, my.

Steve: Anyway.

Leo: Oh, my. Oh, my.

Steve: Once again, we seem to be drifting here a little bit off of the security-related topics recently.

Leo: But it's secure dirt. No, no, it's secure.

Steve: Well, part of the goal of podcasts is to have some fun. And so we're providing some entertainment.

Leo: We do that, yes.

Steve: Okay. With Firefox v107, which was released last Tuesday, a week ago, nothing was earth shattering. There were no critical security fixes, but there were a very large and welcome collection of high-severity things fixed. No zero days that were noted. There were also a couple moderate severity repairs. So, you know, it appeared to primarily be released just to fix those things, since there were not otherwise even a large number of new features. A couple little developer things, you know, they're continuing to push the standards which Firefox supports forward because, you know, the web people can't keep their hands off of, like, ooh, how about if we added the ability to read your mind? That would be good. It's like, well, okay, we don't have that technology. Or no, but let's develop an API for that so that when we do, web pages will be, I mean, this is what's going on. So a little bit more of that is happening. Nothing else to see.

It was interesting to me to see that Google recently settled something that we discussed four years ago. This was a suit brought against Google by 40 states attorneys general. They settled for $391.5 million. Where that number came from only the attorneys know. As I said, we talked about this four years ago, back in 2018, when these offices of those 40 states attorneys general sued Google, alleging that Google had been lying and misleading users into thinking that they had disabled location tracking in their account settings. The lawsuit followed some reporting that was produced by the Associated Press, which found that Google was continuing to track its users even after they had enabled the account privacy setting that claimed to turn off location tracking.

So in that settlement Google agreed to pay this $391.5 million in restitution and also, of course, to change the way it handled location tracking in the future. The first thing we're reminded of is that the wheels of justice, when they don't completely fall off the wagon, do tend to turn slowly, at least in the United States. So it took us four years to get to this point. The other thing we learned is that, thanks to Google's posting about this, their own posting, we've learned what has changed since then.

So their posting last week was titled "Managing your location data." And it brings new meaning to the phrase "putting on a happy face." They wrote: "Location information lets us offer you a more helpful experience when you use our products. From Google Maps driving directions that show you how to avoid traffic, to Google Search surfacing local restaurants and letting you know how busy they are," like all the benefits, right, of Google knowing where you are. They said: "Location information helps connect experiences across Google to what's most relevant and useful." And, okay, yeah, that's certainly the case, or can be.

They said: "Over the past few years" - right, while this lawsuit was in the works - "we've introduced more transparency and tools to help you manage your data and minimize the data we collect. That's why we" - and then they have three things - "launched auto-delete controls, a first in the industry, and turned them on by default for all new users, giving you the ability to automatically delete data on a rolling basis and only keep three months', 18 months', or 36 months' worth of data at a time." And if that sounds familiar to our listeners, it's because, yes, we covered this when this was happening. Second thing they did: "Developed easy-to-understand settings like Incognito mode on Google Maps, preventing searches or places you navigate to from being saved to your account." And third: "Introduced more transparency tools, including Your Data in Maps and Search, which lets you quickly access your key location settings right from our core products."

And they said: "These are just some ways that we have worked to provide more choice and transparency. Consistent with these improvements, we settled an investigation with 40 U.S. state attorneys general based on outdated product policies that we changed years ago." Okay, you know, in addition to the $391.5 million, "outdated product policies that we changed years ago. As well as a financial settlement, we will be making updates in the coming months to provide even greater controls and transparency over location data." So things to come.

These updates include three things: "Revamping user information hubs. To help explain how location data improves our services, we're adding additional disclosures to our Activity controls and Data & Privacy pages. We're also creating a single comprehensive information hub that highlights key location settings to help people make informed choices about their data." Okay, so more transparency.

Second thing: "Simplified deletion of location data. We'll provide a new control that allows users to easily turn off their Location History and Web & App Activity settings and delete their past data in one simple flow. We'll also continue deleting Location History data for users who have not recently contributed new location data history to their account." And third: "Updated the account set-up. We'll give users setting up new accounts a more detailed explanation of what Web & App Activity is, what information it includes, and how it helps their Google experience."

So they finished: "Today's settlement is another step along the path of giving more meaningful choices and minimizing data collection while providing more helpful services." So it seems clear that what was going on during these four years, I mean, lots of back-and-forth, was some negotiation about the things that Google was being asked to do proactively in order to make what they were doing, make this tracking behavior which initially got them into such trouble that these 40 attorneys general decided to gang up and say, look, this needs to change.

So stepping back from this a bit, it must be truer than I guess I'm able to understand that the more information an advertiser has about someone, the more revenue is generated by showing that person advertisements. I mean, as our listeners know, I've always been somewhat skeptical about that, I mean, that it can mean that much. But it seems to me that advertisers would not be trying so hard if it didn't really make them more money, since they also know that no one wants to be profiled and tracked across the Internet. So they wouldn't be risking our wrath to the degree they are if it really wasn't valuable to them. So anyway, we've got, in a minute, we'll be talking about a different issue with some more attorneys general and the FTC.

I caught wind of a mention that Red Hat had started cryptographically signing its deployment ZIP files. Okay, now, that made me curious since I'd never heard of ZIP files being cryptographically signed. We're always talking about executables being signed. And we know that web assertions of their identity are signed. But that was new for me for ZIPs. And with all the problems that we've been seeing with supply chain poisoning, obtaining verifiable assurance of an archive's unmodified authenticity, that would be great. So a cryptographic signature could do that.

And cryptographic signing makes way more sense than the old-school practice of publishing the hashes of files on the same site where the files are being hosted for download. Doing that never made any sense to me since, if a bad guy was able to compromise a web server to alter the files being downloaded from that site, what would keep them from also updating the hashes shown at the same site as proof of a file's authenticity? Talk about a false sense of security. So anyway, this is a lot better than that.

So I looked into what was going on, and I found a posting by Red Hat titled "Cryptographic signatures for zip distributions." I've paraphrased what they posted to remove a lot of their over-simplified descriptions for our audience. So they wrote: "Our build system, Brew, produces our RPM and zip distributions and automatically hashes the archives it makes. The hashes are used to validate that the files have not changed before they're uploaded to our CDN and made available to customers. We've taken advantage of this aspect of our build process and extended it by combining all of the hashes for a particular release and packaging them into an SHA256SUM file, so S-H-A-2-5-6-S-U-M file.

"This file is in a standard format that lists the hash and the corresponding filename of the particular file." "Artifact" is the term they use. "It is commonly used across the industry to provide integrity to binary files. However, it's not limited to that. The SHA256SUM command on Red Hat Enterprise Linux, other Linux distributions, and macOS natively support this file format."

They said: "Once our software production team has completed their verification procedures, they sign off on the release from both a process and technical perspective. The SHA256SUM file they created is signed by our latest release key, which produces a .asc file. This file is an ASCII-Armor formatted detached signature file that proves the integrity and provenance of the SHA256SUM file and, transitively, the zip file artifacts enumerated within that file. The GPG command on Red Hat Enterprise Linux, other Linux distributions, and macOS supports the file format natively.

"Due to the potential damage that a lost or stolen private key could cause, we have taken additional steps to add assurance to the signatures we produce. The primary technology behind this is our signing server. To sign these files we use a high-strength, 4,096-bit private key, and our public keys are available on our website and the MIT (Massachusetts Institute of Technology) public key server."

Okay. So that's what they posted. Red Hat's mention of a detached signature simply means that the signature itself resides in a separate file. The signature is just an SHA-256 hash of the file it's signing, which is then encrypted under Red Hat's super-secret, and in this case very long, 4,096-bit private key, which they're careful not to let loose. Just like my GRC code-signing keys, it probably resides in an HSM, a hardware security module, where it literally cannot be extracted. It can only be used. So there's no reason for that signature file not to stand alone, that is, again, it's just - so there's this composite file which contains the hashes and the files that they were hashed from. That's just a listing, a textual listing, an ASCII file. That file is then SHA-256 hashed. That's the file whose integrity you want to verify.

That SHA-256 hash is then signed with their signing server, which is to say that the SHA-256 hash is encrypted with the private key. So that creates an encrypted blob which is the signature. And it's a freestanding file. So somebody who then wants to verify that uses Red Hat's private key, which is available from several sources so you don't have to worry about that being screwed with, in order to decrypt the blob. That will bring - decrypting that blob restores the SHA-256 hash which you can then use to verify that the file of the hashes that you've got matches and has not been tampered with.

So this is a welcome move as a deterrent to the abuses that we are now seeing and talking about more and more of today's supply chain. And it's probably where the broader open source community will need to go. The glitch here, the glitch to doing that, is that Red Hat Enterprise Linux Corporation, you know, Red Hat Corporation, has no problem maintaining a signing server and buying a certificate that asserts their identify. But the open source world has always had a problem with the need to pay for certificates.

As we know, Let's Encrypt solved this problem by making TLS certificates free for web servers. But the challenge here is not the same. Let's Encrypt offers no guarantees about the identity of a site. It provides domain validation certificates where the only requirement is for the certificate to match the server's domain name. Specifically, it does not offer, that is, Let's Encrypt does not offer OV (Organization Validation) certificates. In order to issue OV certificates, any certificate authority must by universal agreement perform some significant reconnaissance to positively verify the identity of the entity requesting the certificate so that the OV-ness means something. And what's more, of course, many open source projects are just some guy working alone without any organization to be validated.

So maybe the solution will be, for example, to come up with a secure means for submitting repositories to GitHub for its signing with its signature, then using some much stronger means for asserting the identity of the individual requesting the signing service. For example, that process might require much more rigorous multi-factor authentication. Something, again, you're really wanting to put it out of the reach of bad guys to get in there and screw this up so that it means something. So it's a problem that needs to solved. But one way or another, we need a solution to this current supply-chain pollution problem. And, you know, the application of a bit of crypto might be a place to start. So hats off to Red Hat for doing a little pioneering here in that way.

Okay. Now, the FBI purchased Pegasus, you know, that's the NSO Group's infamous smartphone spyware platform. They said it was for "research and development purposes."

Leo: Yeah, what are they developing, I wonder?

Steve: Uh-huh, yeah. Last week The New York Times ran a story with the headline "Internal Documents Show How Close the FBI Came to Deploying Spyware." Now, I have a little bit different take on this, but we'll get to that in a second. The New York Times reported that, last December, FBI director Christopher Wray (W-R-A-Y) told Congress, this was closed-door testimony, that the Bureau purchased - "bureau" as in, you know, Federal Bureau of Investigation - the Bureau purchased the infamous Pegasus phone hacking tool for "research and development purposes." Well, it turns out that FOIA, the U.S. Freedom of Information Act, can be quite handy for figuring out things that really happened.

Here's how the Times explained what they found. They wrote: "During a closed-door session with lawmakers last December, Christopher A. Wray" - spelled W-R-A-Y - "the director of the FBI, was asked whether the Bureau had ever purchased and used Pegasus" - so, like, directly asked - "the hacking tool," writes the Times, "that penetrates mobile phones and extracts their contents. Mr. Wray acknowledged that the FBI had bought a license for Pegasus, but only for research and development, 'to be able to figure out how bad guys could use it, for example,' he told Senator Ron Wyden, according to a transcript of the hearing that was recently declassified.

"But dozens of internal FBI documents and court records tell a different story," writes the Times. "The documents, produced in response to a Freedom of Information Act lawsuit brought by The New York Times against the Bureau, show that FBI officials made a push in late 2020 and the first half of 2021 to deploy the hacking tools made by the Israeli spyware firm NSO in its own criminal investigations." That is, in the FBI's own criminal investigations. "The officials developed advanced plans to brief the Bureau's leadership, and drew up guidelines for federal prosecutors about how the FBI's use of hacking tools would need to be disclosed during criminal proceedings." Like, okay, how did you get this information? Uh, well, it came to us. Uh-huh.

So the Times writes: "It's unclear how the Bureau was contemplating using Pegasus, and whether it was considering hacking the phones of American citizens, foreigners, or both. In January, the Times revealed that FBI officials had also tested the NSO tool Phantom, a version of Pegasus capable of hacking phones with U.S. numbers. The FBI eventually decided not to deploy Pegasus in criminal investigations in July of 2021, amid a flurry of stories about how the hacking tool had been abused by governments across the globe. But the documents offer a glimpse at how the U.S. government, over two presidential administrations, wrestled with the promise and peril of a powerful cyberweapon. And despite the FBI decision not to use Pegasus, court documents indicate the Bureau remains interested in potentially using spyware in future investigations."

And of course the Times reporting brings up the question of Christopher Wray's apparently misleading testimony in front of Congress. Senator Ron Wyden is not happy about that. In a statement from his office, it read: "It is totally unacceptable for the FBI director to provide misleading testimony about the Bureau's acquisition of powerful hacking tools, and then wait months to give the full story to Congress and the American people." So the Times revealed in January that the FBI had purchased Pegasus in 2018 and, over the next two years, tested the spyware at a secret facility in New Jersey. Since the Bureau first purchased the tool, it has paid approximately $5 million to the NSO Group.

Now, it seems to me that the issue with Pegasus is less about its use than its potential for misuse and abuse. The worry is that, once they have it, repressive governments would be unable to resist the temptation of using it to spy on political rivals - we'll see an example of that here in a moment - and, of course, dissidents and other non-criminal actors. And of course Pegasus doesn't respect geopolitical boundaries. So anyone who has it can aim it at anyone else, anywhere. But in the United States we have a system for obtaining court orders for searching and for making legal, within bounds, what would otherwise be illegal reconnaissance.

So as long as the FBI would only be using Pegasus within our constitutional protections, I think that it would be a useful tool to empower their criminal investigations. And yes, they would be required to tell a judge that this is what we want to do. This is how we're going to do it. And we have probable cause and all the other requirements of getting a court order to pursue things like a wiretap and so forth. So it seems to me, yes, it is problematical because it can be abused. But if we're going to have systems that are otherwise not prone to be subject to court order search, then maybe this is the way it happens.

Leo: Yeah, I mean, we allow wiretaps under court order.

Steve: Exactly. Exactly.

Leo: Is Pegasus somehow too dangerous to be used?

Steve: I think the concern is just - it's control. All the reports we have suggest that it is a zero-click tool which it is possible to target at an individual smartphone, and it goes in against all of the attempts by Apple and Google, you know, iOS and Android to keep it out. There are enough ways in that it gets in, and then it's able to provide the entity that deployed it with information the equivalent of someone unlocking their phone and also being eavesdropped on. It's able to, you know, it is a surveillance tool.

Leo: I guess the question always is, is it - first of all, it's going to be very expensive. It's a million-dollar surveillance tool; right. It's very, very expensive.

Steve: Yeah, multi.

Leo: Multi. Because it can't be used too often or it loses its usefulness. Because as soon as the companies find it, they'll defend against it. So these zero-days are very, very expensive, especially if it's no-click.

Steve: And it also might very well be that it is tightly tethered.

Leo: Well, that's actually, and this may have been the problem, as I understand it, the NSO Group is responsible for the hack. You don't just give the FBI Pegasus and say, "Have fun with it, guys."

Steve: Yes.

Leo: It doesn't work that way; right?

Steve: Right, right.

Leo: So that's another problem is that some international company, Israeli company, would then be privy to what you're doing.

Steve: Right.

Leo: Yeah, that might be a bigger problem.

Steve: And other entities may not care, but that may be something that we can't get over. And in fact, maybe that was the beginning, you know. In testimony like this there's typically some piece of truth. So probably, yeah, the FBI said maybe we need to be empowered with this tool.

Leo: Right.

Steve: Because we're unable to get in any other way. So let's buy a copy, and let's learn how it works. Let's have the Pegasus experience so that we can decide if this is something that we can sell to the greater government.

Leo: It's my understanding that what you're buying really is the NSO Group...

Steve: Access.

Leo: They trigger it on - let's say they want to spy on my phone. The NSO Group gets into my phone, triggers it, and then hands control over to the FBI.

Steve: Right.

Leo: So I'm sure that's illegal in the U.S. because that's an Israeli company, not even a governmental entity, but a business that the FBI says, okay, we want to hack Leo's phone, here's his phone, hack it for us. It can't be legal. And you're right, it literally was research. They just wanted to know, well, how, you know, let's understand it a little bit better. But I can't imagine the NSO giving the keys to the kingdom to the FBI either. That's why they do it that way; right?

Steve: No. Yes. And in fact I have another related story that sort of speaks to that. Greece, the Athens government, Greece, bought a related program, Predator, for 7 million euros.

Leo: Wow.

Steve: A recent report in the Greek press claimed that Greece's government paid 7 million euros to Intellexa, I-N-T-E-L-L-E-X-A, Intellexa, for access to the Predator surveillance and spyware platform, and an additional $150,000 euros for the ability to rotate 10 new targets per month. So that says, yes, they were not given carte blanche. They had to - it is tightly tethered under Intellexa's control. So this little bit of accounting news follows the massive scandal of the Greek government having been caught using the spyware to go after not only rival political parties, but also journalists and prosecutors investigating government corruption.

So this is the double-edged sword is that it seems to be impossible for governments that purchase this to behave themselves. Again, I would hope that, if it were made possible for the FBI to acquire this technology, it would be done aboveboard. It would be done within the constitutional protections of the government. I'm sure there are those who, you know, Edward Snowden, who don't believe it could be possible. But we do have - we've set up a situation where the technology that our private citizens and corporations are using is not subject to court orders. And thus the tension that we're currently under.

So anyway, again, as I said, it seems to me the problem is less about the tool than how it's used. It is technology. It already exists, and it's going to exist. So it makes more sense to me to properly regulate and control its use than to attempt to deny it completely, which, you know, just forces its use underground.

Leo: And maybe it's old-fashioned, but I also feel like we're the United States. We should be better than those other guys. You know?

Steve: I agree.

Leo: You know, we should have higher standards. Just because other countries use these tools doesn't mean we have to.

Steve: Yeah. I agree.

Leo: All right, Steve. On we go.

Steve: The password manager 1Password has added support for passkeys to its offering. And in a nice promotion of passkeys, they've created a community-supported online directory listing online services currently supporting passkey authentication. I've been waiting for this because I want to play with passkeys. I've got iOS 16.1.1, I think, now. And it's supposed to support passkeys, but I've never tried it. So now we can.

So this directory is at passkeys.directory. I didn't know "directory" was a TLD. Really, they've just gotten out of control, Leo. Is there a .leo? There probably is. Anyway, so, again, passkeys.directory takes you to this listing. It currently has 43 companies listed with their URLs, although some are flagged as MFA so, you know, multifactor authentication, so I suspect that they might not be pure passkeys login. They may be passkeys plus another factor, which would be annoying. So anyway, some notable names on the list, which do appear to be pure passkeys authentication without the MFA tag, include a 1Password Passkeys Demo page. Of all things, Leo, Best Buy.

Leo: Whoo, about time.

Steve: Yeah, it supports passkeys. Carnival Cruises.

Leo: Good.

Steve: eBay.

Leo: Good.

Steve: Kayak, you know, the travel site. Microsoft.com. Again, Nescafe? Like what?

Leo: Sure. Why not?

Steve: Nvidia, PayPal, and Robin Hood. So anyway, I just discovered this as I was putting the podcast together, so I have not made any time to experiment with and explore. But I am an avid buyer on eBay.

Leo: Oh.

Steve: Often buying, like, old hard drives that I need to make sure that SpinRite works with. Or in fact I'll be talking about SpinRite in a few minutes here because I actually did just buy four drives from eBay which were specific drives that I needed. So anyway, I ought to be able to give logging into eBay on passkeys a try.

Leo: I think I'm seeing it. Let me log in, and I'll show you. I'm going to log into - I'll go to Carnival Cruises. And it says "Create an account." Here, let me show you this. I'm going to make it bigger. And see, that "Login with your phone's Face ID or fingerprint," that's passkeys.

Steve: Ohhh.

Leo: It may not say "passkeys." Right?

Steve: Right. Right, right, right.

Leo: So scan this QR code - all right, let me try it - with your phone's camera. So this is - yeah. That's cool. I'm so glad. This is the first time I've seen it. All right. I scanned it with my camera. I'm logging in. Enter your email. Okay.

Steve: Ooh, and the site knows, Leo. Look what it's doing.

Leo: Oh, it does. It knew I did something.

Steve: Yeah.

Leo: How would it know that?

Steve: Oh, because you're...

Leo: I'm going to a special URL. That QR code is...

Steve: Yes.

Leo: Okay. Connection lost. Something went wrong. Try again. Oh, crud. Hah. Well...

Steve: Well, after all, it is Carnival Cruises, so...

Leo: We're working, working on it. So now what do I do? Now what do I do? Do I take another picture? Let's do it again. No, I can't.

Steve: Do you have an account at Carnival?

Leo: Not Carnival, no.

Steve: Okay. How about Kayak?

Leo: Well, I think the idea is you would have to - oh, should I go somewhere I already have an account? You want to see what that looks like?

Steve: I don't know. I don't know.

Leo: Let me just - I didn't do it quickly enough, probably.

Steve: Let me try. What happens if I log out of eBay? If I'm, like, statically...

Leo: Yeah. So now I'm pressing Continue. Oh, here it is. Do you want to allow Carnival.com to use Face ID? Continue. I'm using Face ID. It worked. And look at this, on the phone it now says - let's see if I could find that - passwordless sign-in enabled.

Steve: Enabled.

Leo: Fast login by own ID. But this is passkeys. Right?

Steve: It's got to be passkeys.

Leo: Yeah. So that's cool. So now it wants, you know, complete your profile, blah blah blah. But now I presume from now on I can just use my phone. I love it. Yay.

Steve: Yes, yes.

Leo: Now I have a Carnival Cruise Line login. Not sure I want that. Oh, actually the cruise line that we do go on is owned by them, so I guess that's - one of the cruise lines we like to go on.

Steve: Again, this is the weirdest list. Like Best Buy. Carnival Cruises.

Leo: Don't you think I should trust these people because they're at least on top of it?

Steve: Chase is not there. BofA is not there.

Leo: No, banks is going to be a higher standard.

Steve: You know, Nescafe, but not Starbucks. It's like, okay. I don't know what's going on. But anyway...

Leo: I think it's going to be lower stakes companies, don't you think, initially? A bank, that's going to be problematic, probably.

Steve: Yeah, I guess Microsoft has become lower stakes; you're right.

Leo: Would I like to receive emails? No. Do I accept their terms and conditions? Yes. Have you already booked a cruise? No. Okay. Now I guess the next time I go there - let's go on another computer. This is the first time I've ever used this. It's cool. So now I'm going to see "Login," and it's going to say "Login in with your phone's Face ID or fingerprint." I'm going to click that. Scan. Oh, I have to scan it again. Is that right? Is that what it should be doing?

Steve: Yes, yes, because you don't have - you haven't transferred your passkey into that computer.

Leo: And then it says do you want to login using a saved account? Yes. Logging in. Bingo. "Hi, Leo. When's your cruise?" It works. It's a little onerous. So will I always have to scan my QR code to get in?

Steve: Well, so what you're doing is you're using your phone's passkey in order to authenticate across to a different device.

Leo: Right.

Steve: And this was the problem that I talked about is that SQRL would - there would only be one. But so you need to create another passkey in your laptop. And so there is, there should be a way to - you can't export the passkey, but you can link them. You can create another passkey and then link them so that they're identified as the same.

Leo: Yeah, see, I don't see - I already have a Microsoft account, but I don't see any way to log in, if you haven't set up passkeys, with passkey; right? I'm just going to the Microsoft site. Now, I do have an account, and I could sign in, but I'm going to say could I do this with my passkey. No. But maybe if I go into my account I could set that up.

Steve: What's that little thing down at the bottom?

Leo: Sign-in options, but that's - I already looked at that, and it just gives me GitHub or forgot my username. That's not passkey. I bet you I have to go to the Microsoft account, log in normally, and then say "and I would like to establish passkeys with this account."

Steve: Probably, yeah.

Leo: That would make sense. I'll try it while you're talking.

Steve: Anyway, all of our listeners now, again, passkeys.directory. You can check back there, and maybe eventually some more interesting sites will be available.

Leo: I think it's hysterical that Robin Hood is using it.

Steve: Yeah, yeah.

Leo: I bet FTX would have.

Steve: Go to passkeys.directory and see what it...

Leo: Oh, there's more than just this. Okay, yeah, yeah, yeah.

Steve: Oh, yeah, yeah, because there was a bunch of things that also had MFA tags for some reason. So you could see that the little green dots are just...

Leo: Sign-in.

Steve: Sign-in.

Leo: And here's Cloudflare MFA.

Steve: Yeah. So I didn't know what that meant.

Leo: It probably means I need a password to log in and passkeys. Like it's two-factor. Yeah?

Steve: That's what I'm thinking, too.

Leo: Use DocuSign. I can sign in, or I can - interesting. Well, I have a GitHub account. Let me play with that a little bit and see.

Steve: Ah, okay, cool.

Leo: Yeah.

Steve: In other news - again, passkeys.directory, our listeners. Okay. So from the Having Fun with Bureaucracy department comes an edict from the OMB. The U.S. Office of Management and Budget has ordered federal agencies to scan their systems - oh, yes, scan those puppies carefully.

Leo: Scan them, man.

Steve: Scan them, and provide an inventory of assets containing cryptographic systems that could be cracked by quantum computers in the coming years.

Leo: Just how would you know?

Steve: Well, Leo. Okay. First of all, there is probably not a single computer in the government that doesn't use and depend upon some public key crypto. And none of the currently deployed public key crypto...

Leo: There's no way.

Steve: ...is quantum resistant.

Leo: Yeah.

Steve: So the OMB could have simply said, give us a list of all your computers.

Leo: That's a good point.

Steve: And by the way, stop using them.

Leo: Yeah.

Steve: Okay. So the next point worth noting is just a reminder that no one has come near to building a quantum computer anywhere, so far as anyone knows, that could even begin to think about breaking actual public key crypto. Oh, yes, factoring the number 27, we can do that. It's magic. But the number 35? Uh, we're not quite there yet. Give us another 10 years or so, and we'll be able to factor 35.

Okay. Now, that said, I'm on the record agreeing that there's absolutely no reason not to move us to quantum-safe crypto sooner rather than later. You know, let's not wait till we need it because we know how slow and painful these moves can be. So just as soon as we are absolutely sure that we're not going to be making a big mistake because that's possible. Remember that one of the candidates that had already been chosen, already selected, was recently cracked by conventional computers. So it would be a lot better for us to stay where we are, where we know we can't crack today the algorithms we're using, before moving prematurely to something that we presume some future non-existent mythical quantum computer should also be unable to crack.

So the OMB edict stated that federal agencies had until May 4th, 2023. So like this coming May 4th. I don't know why May 4th, but that's it. And the NSA ordered that all government agencies handling classified information must use quantum-resistant encryption by 2035. Okay. So that's 13 years from now. By then we ought be up to factoring 45. So good to be switching over to quantum computers, you know...

Leo: Any minute now.

Steve: Before we need them, yeah. Okay. So this other piece of attorneys general news that I wanted to share, one of the developing themes of this podcast is the observation that we're still in the Wild West stage of the creation of the Internet. It remains an unregulated or only very loosely regulated medium. And of course globally it's an uncoordinated total disaster. The idea that we've linked our fundamentally insecure networks to those of openly hostile nations should give anyone pause. Yet that's what we've done.

Chinese, Russian, and Iranian cybercriminals under the protection of their nation states, who have no love for the U.S., are able to openly attack the networks of U.S. corporations and its private citizens. And yes, there's reciprocity. You know, the U.S. is able to do the same to them, and presumably that's happening, too, although there seems to be a surprising lack of information about that. You know, but reciprocity doesn't make any of this sane. It's like, you know, mutually assured destruction. So we can only hope that the Internet our grandchildren will use as adults 30 years from now will be much different from the one we've been watching being born through these past 30 years.

I bring this up because various democracies around the world, notably the EU and the U.S., among others, are inching forward cautiously in an attempt to provide their citizens with some legally enforceable rights to privacy and personal information. At the moment, we have clear statutes outlawing overt network intrusion and attack. And when those laws are broken, people lose their freedom for doing so. But nothing yet prevents or regulates the passive collection of as much Internet user data as possible. Google was sued by those 40 states attorneys general, not for tracking, but for tracking after they said they weren't. As long as a company doesn't say that they won't do something, they can do pretty much anything they want.

So how do we get this to change? Here's a hopeful example: Last Thursday a coalition of 33 state attorneys general co-signed a letter formally urging the U.S. Federal Trade Commission, our FTC, to pass legislation which would regulate online data collection practices. Might not happen, but it's a good start. These AGs said they are "concerned about the alarming amount of sensitive consumer data that is amassed, manipulated, and monetized." And they also said that they regularly receive inquiries from consumers in their states about how their own data is being hoarded and abused.

Okay. Since we've still got a bit of time, and I think this is extremely important, I'm going to first share just the introduction in the letter which was submitted to the FTC and signed. It's really pretty, like different colors of ink on the signatures. I don't know how they pulled this off. But it was like, you know, signed by 40 states attorneys general.

So in the beginning of this letter they said: "We, the Attorneys General of Massachusetts" - I'm not going to read them all because they didn't list them all, but they did some - "Massachusetts, Connecticut, Illinois, New Jersey, North Carolina, and Oregon, joined by the respective Attorneys General of the undersigned states, write to the Federal Trade Commission in response to the August 22, 2022 Advanced Notice of Proposed Rulemaking on Commercial Surveillance and Data Security." So this was something that the FTC put out there and asked for comments. So that was an Advanced Notice of Proposed Rulemaking on Commercial Surveillance and Data Security. That all sounds great.

So they said: "As the chief consumer protection officials in most of our respective states, we hope to inform the Commission as it contemplates new trade regulation rules governing commercial surveillance and data security. The State Attorneys General commend the FTC for its comprehensive review of corporate surveillance and data security in preparing the Notice. We, too, are concerned about the alarming amount of sensitive consumer data that is amassed, manipulated, and monetized. Our offices frequently receive outreach from consumers concerned about the privacy and security of their information. Research supports that consumers are worried about commercial surveillance and feel powerless to address it."

Leo: Oh, really.

Steve: Imagine that.

Leo: That is interesting.

Steve: We're just going on the record here. "Many consumers believe that tracking by companies is inevitable, yet often do not even know what is being recorded. These fears intensify when they learn more about the commercial surveillance economy, and in particular consumers fear falling victim to identity theft and data misuse. A majority doubt that their data can be kept secure. Contributing to these concerns is the fact that companies are often collecting more data than they can effectively manage or need to perform their services.

"Our consumer privacy-related enforcement actions and investigations have resulted in settlements" - like Google - "that have provided significant business practice changes to strengthen data security and privacy going forward, but there is still more work to be done. Our submission highlights the heightened sensitivity of certain categories of consumer information, the dilemma of data brokers and how they surveil consumers, and how data minimization can help mitigate concerns surrounding data aggregation."

Okay, then the letter goes on at quite some length detailing five general categories of abuse. Unfortunately, in an effort to be very clear and to drive their points home, that part is too long to share. But I found a separate release about this action from New Mexico's Attorney General Hector Balderas. It addressed each of these five points by reference quite succinctly, so those I want to share because it's good stuff.

So first, so there's five categories. Location data, he said, or his office said: "According to the letter, many consumers are not even aware that their location information is being collected; and, when a consumer wishes to disable location sharing, their options are quite limited. The attorneys general recognize the sensitive nature of this information, which can reveal intimate details of daily life such as where they live and work, their shopping habits, their daily schedule, or whether they visited the doctor or pharmacy. Laws passed in states like California, Connecticut, and Virginia that restrict the use and collection of location data can provide a framework to inform the FTC through the rulemaking process." So this is him saying, or his office saying, for location data things, look at what California, Connecticut, and Virginia have done. Use that, you know, consider using that as a framework.

Biometric data: "The coalition urges the FTC to consider the risks of commercial surveillance practices that use or facilitate the use of facial recognition, fingerprinting, or other biometric technologies. Many consumers provide this information to companies for security purposes or to learn about their ancestry. But consumers are not always made aware of when their data is collected, how it is used, or if it is resold for purposes to which they never meaningfully consented."

Medical data: "The FTC should also consider the risks of practices that use medical data, regardless of whether the data is subject to the Health Insurance Portability and Accountability Act of 1996" - popularly known as HIPAA - "and the Privacy Rule. Medical data not necessarily covered by HIPAA is referred to as 'health adjacent data,' which can be collected by many devices, for instance, smartwatches, heart monitors, sleep monitors, and health or wellness phone applications. The letter also highlights medical information risks through examples such as the storage of health-related Internet searches, or appointment scheduling information being passed to others through online tracker tools." In other words, you get a sense for how comprehensive this letter was that the 40 states attorneys general submitted to the FTC.

Two more to go. Data brokers: "The attorneys general reiterated to the FTC the persistent dangers of data brokers. Data brokers profile consumers by scouring social media profiles, Internet browsing history, purchase history, credit card information, and government records like drivers' licenses, census data, birth certificates, marriage licenses, and voter registration information. Data brokers use this information to create profiles of certain consumers which can be purchased by almost anyone based on susceptibility to certain advertising or likelihood to buy certain products. This scale of aggregation of anonymously gathered information can identify consumers and put consumers at risk of scams, unwanted and persistent advertising, identity theft, and lack of consumer trust in the websites they visit."

And lastly, data minimization: "The attorneys general say that it is vital that the FTC consider data minimization requirements and limitations. With respect to data collection and retention, the letter encourages the FTC to examine the approach taken in California, Colorado, Connecticut, Utah, and Virginia consumer privacy laws which mandate that businesses tie and limit the collection of personal data to what is 'reasonably necessary' in relation to specified purposes. Limiting the collection and retention of data by businesses will improve consumer data security as businesses will have less data to protect and less data potentially available to bad actors."

Okay. So I think, if nothing else, this is a useful start. In the United States, where we exalt capitalism, no one wants to strangle innovation. But we all know that we're a long way from being in danger of that. Much of what is going on today is only able to happen under the cover of darkness, because consumers are blissfully unaware. You know? What did Apple discover when they started requiring their apps to proactively obtain cross-application tracking permission? They found that nearly everyone who was asked, declined. No thanks, and no surprise. So we can expect any improvements to be slow going. As I always say, change is slow. But the pressure is there, and it's not going to go away. At least I think we're moving in the right direction. And 40 states getting behind this, you know, one wonders why it's not 50. Well, who knows? Some presumably buckled to some pressure.

Okay. Some closing-the-loop things that I think are interesting. Vincent Stacey shot me a note that I wanted to share regarding - we were talking about the concern that was raised by a different listener about the ZimaBoard and how, when he changed its credentials, it was only away from the logon of casaos/casaos, it was only for the web portal logon, and all of the other credentials remained the same. He was concerned that the lack of changing of other credentials was unknown to ZimaBoard users, and that they might get themselves in trouble, for example, if they turned this thing into a router.

Anyway, Vincent Stacey tweeted: "Hi, Steve. pfSense installs its own version of Linux and won't have the default users of another distribution." And that's a very good point for anyone who's interested in using a ZimaBoard as a pfSense router, though just for the record it's actually FreeBSD Unix that pfSense runs on top of and brings along with it. But the main reason why a ZimaBoard would not be my first choice as a router is that, unless a network expansion board were to be plugged into its PCIe x4 slot, it only has a pair of LAN NICS built in, and I would expect a router today, certainly one that any of our listeners would be using, to have a few more network interface controllers, a few more NICs for implementing useful multi-network isolation. So I can't see it being really popular as a router. There are some better fanless solutions like that, what is it, the SG-1000, I think, that I've talked about before.

Charles Turner tweeted: "As possible fodder for a Listener Feedback section in a future episode of Security Now! podcast, I have a question arising from the discussion you and Leo had on Tuesday" - okay, he says November 15th, that was last Tuesday - "during Security Now! Episode 897, Memory-Safe Languages." Yup, last podcast. He says: "With the future of Twitter in doubt, what is your prediction on the long-range fate of Mastodon? The cynical part of me gives Twitter a 50/50 chance of either, A, rebounding back to its former glory and beyond; or, B, becoming a $44 billion version next iteration of MySpace and FTX."

Okay. So it's clear to us all that Twitter is currently in turmoil. And I don't have any firsthand sense for just how fragile Twitter's technology is internally. And it seems to me that matters a lot. If the previous regime engineered really solid bulletproof systems, then it ought to be able to withstand Elon's shaking of its foundation. But overall I'm a big believer in inertia and in things generally changing much more slowly than we expect. Now, of course, Elon could trip over the main power cord, and Twitter could go dark until someone plugged it back in. And I suppose I'm interested in what Elon is doing there; you know? He's an interesting character, and somehow he's managed to get other people in the past, at least, to do some truly amazing things. I'll never forget the sight of those twin booster rockets returning to and landing on that floating platform for reuse.

Leo: Oh, yeah. That was something, yeah.

Steve: That was truly astonishing technology.

Leo: Inspiring, yeah.

Steve: And it's Elon's SpaceX Starlink technology, which actually works, that's enabling Ukraine to survive Russia's increasingly aggressive attacks against its infrastructure. Again, thanks, Elon. Mostly, though, my take is that I think Elon is just having fun with his life, as is his right. Right? You know, I hope he's having fun.

Leo: Expensive fun.

Steve: You know?

Leo: But what about our lives?

Steve: He doesn't care.

Leo: He doesn't care, no.

Steve: No, he doesn't.

Leo: He thinks we're simulations, that's why.

Steve: Because it's his life. And he's not a guy who likes to make small waves. Elon's waves are big. And let's not forget that Twitter made him do it. They insisted that he honor his wildly overpriced purchase offer. He didn't want to buy Twitter. They made him buy it. So it seems to me that Twitter is getting what it deserves: the Elon treatment. He's showing them that he can do anything he wants to with it.

Leo: Yeah.

Steve: So all of this made me curious about what he is doing with it. You know, I pick up little bits here and there, but I don't follow news feeds, or even Twitter, because they interrupt my work and my train of thought. So it was with some joy that I stumbled upon a site, which I figured had to exist somewhere. The site's called TwitterIsGoingGreat.com.

Leo: In the spirit of Molly White, yup.

Steve: And yes, of course, it's offering up its share of schadenfreude. So keep in mind that it's naturally going to biased. But it's still a lot of fun. The site hosts a simple timeline of Twitter's Elon-related happenings. So now I can check in from time to time whenever I want to get a sense for what's going on over there. I mention it because I imagine that some of our listeners would also appreciate knowing about this nicely distilled timeline event resource.

Leo: It's hysterical because it's all tweets.

Steve: Yes.

Leo: I guess that's a best source of what's going on at Twitter, I guess, yeah.

Steve: TwitterIsGoingGreat.com.

Leo: I'll show you another one that you should read.

Steve: Okay.

Leo: This is from a Twitter reliability, site reliability engineer. I think former. Matthew Tejo, he's on Substack. And I think you would enjoy this. I barely understood it, but he talks about all of the redundancies, all of the automation. He says: "When I came in, the list of servers was on a spreadsheet. Now, of course, it's a much better system." And he did a really good job. It sounds like he and his team did a really good job of making it run. He was in charge of the cache, the cache team, which is a pretty big deal because everything you're getting is served from cache. None of it's served from the source.

Steve: Leo, I didn't want to interrupt you, but has anyone stopped to think about what it does?

Leo: Oh, it's phenomenal, yeah.

Steve: It is un-be-frigging believable.

Leo: This is just a fraction of it.

Steve: What Twitter actually does.

Leo: Yeah, yeah.

Steve: I can't imagine building this system.

Leo: Oh, yeah.

Steve: It just astonishes me.

Leo: Well, read this. I think you'd enjoy it. And it's just a fraction of what is going on. And but his point is these things are designed to run unattended. We automated everything we could. And so it should, unless something, you know, nobody's going to kick the plug out of the socket. I hope there's more than one plug. But it could have a...

Steve: You piss off Elon, he might...

Leo: Well, that might - yeah.

Steve: He just might pull the plug.

Leo: So, but you wouldn't expect it to all fail all of a sudden. There may be bugs here and there and stuff. And the real problem is there may not be somebody to solve that problem which cascades to another one, et cetera. I've read a number of articles. We had Phil Libin on, who was the founder of Evernote.

Steve: Very, very well rounded. I was very impressed with Sunday's...

Leo: He's a smart guy. And he was saying, you know, give Elon, as you do, give Elon some credit. There was a good article by a former Tesla engineer that says Elon did exactly the same thing in 2018 to Tesla. He was firing people. He was spending the nights there. He was bemoaning they might be bankrupt. This was all in the lead-up to the Type 3, the Model 3 of the Tesla. They said this is kind of how Elon works. Obviously, for some people, not the ideal situation. That's why so many have left Twitter voluntarily, as well as involuntarily. But I've also read articles that say, you know, this is how he - he's reinventing Twitter. You have to get rid of almost everybody and then build a team of people who believe in your vision. He hasn't really communicated that, apparently, but who believe in your vision first.

Steve: And so he's still making it up. He's making it up as he goes along.

Leo: Nobody, you know, I don't - I'm confused. I see stuff that looks crazy. He says we're going to have a committee to approve who comes and goes. And then he just says, no, I'm going to bring them back. And, you know, it's just - it seems chaotic.

Steve: There was one piece there that said he sat down and explained to the core team how advertising should be tweets. And they said, uh, they are, Elon.

Leo: He said native, yeah, they should be native. It is, yeah, that's exactly my problem with the advertising. You know, so he's coming somewhat from ignorance. But you're right, he's also a pretty interesting, apparently...

Steve: Probably sleeping there.

Leo: He says he is.

Steve: He's there, you know, 24 hours a day. And, you know, he'll figure this thing out.

Leo: He's a weirdo. And some of the things he's tweeted, I'm not thrilled about some of the pictures and stuff. But this is from 1:20 a.m. at Twitter he - this is when he, you know, came, had everybody come in Saturday, or Friday night and Saturday morning, to explain how Twitter works. And these are the skeleton crew. There he is sitting with them. But this is his picture of what they drew on the whiteboard. This is not a code review. This is explaining in rudimentary fashion to somebody who doesn't know how this stuff works, how it's working. I get, you know what, we don't know yet. He may - this may be Twitter 2.0 he's inventing. And maybe this is how he works. I would never want to work for him. But people will, and we'll see what happens.

Steve: There was an interesting moment. I was watching a press conference when Biden was off in the East. And it was that awkward press conference where he meant to say "Cambodia," and he said "Colombia" three times. It's like, oh, Joe.

Leo: Oh, Joe.

Steve: But someone in the press pool asked him about Elon. And so understand that our relationship, the government has a relationship, right, with Elon because he's now SpaceX, and we've got all these contracts.

Leo: Right.

Steve: So Biden just locked up. He didn't know what to say because it's like, oh, you know, I don't dare piss off Elon, or we're going to be in real - we're not going to have any missile launches.

Leo: He says we're looking into it, though. You know, it's very complicated because Elon has relationships with not just the U.S. government, but many other governments. Tesla sells a lot of cars and builds them in China. It's a complicated system. And he's kind of a bull in a china shop, but we'll see.

Steve: I just think he's a - he's a character.

Leo: It's fascinating.

Steve: And I think he's having fun with his life.

Leo: Yeah.

Steve: And, you know? And we're all just sort of observers.

Leo: Too bad, though, because Twitter is a valuable resource. It's not a public resource. It's not even a publicly held company anymore. And he's...

Steve: It's incredibly valuable.

Leo: But it's a shame if he crashes it; you know?

Steve: And that's why I think - well, I mean, he's let a bunch of loons back on recently. And, you know, but I don't ever see tweets from loons. I have a very quiet experience with Twitter. I just talk to our listeners.

Leo: Right.

Steve: And they talk to me. And it's just a great little channel. So I don't care who says that vaccines are garbage. Who cares?

Okay. Leslie MacFarland said: "Hi, Steve." Uh-oh. "If Twitter implodes, are you going to Mastodon or somewhere else? Your Security Now! podcast is top-notch security and quality." Well, thank you, Leslie.

So, okay. In order to get the word out to 18 years' worth of SpinRite owners, I will shortly, and I mentioned this before on the podcast, be setting up an old-school email facility. One of the several lists that I'll be maintaining will be for Security Now! listeners who would like to subscribe to the weekly links and the show notes and a description of each week's podcast, which I post to Twitter. And it'll be nice to have more than 280 characters for that. So that will be a possibility. And, you know, as for Mastodon, I have no idea. I'm not...

Leo: Folks, remember, it took me 10 years to get Steve on Twitter.

Steve: Right.

Leo: Patience.

Steve: Thank you, Leo.

Leo: We'll get him.

Steve: I'm really not looking - I'm not looking for more connectivity. We'll see how Twitter goes. As it is, I spend most of my time in GRC's quiet newsgroups.

Leo: There you go. That's the best place.

Steve: Getting actual work done. And now we have GitLab for managing SpinRite bugs and feature requests. And I have GRC's web forums, which will soon be quite active since that's where SpinRite's tech support will be hosted. And a lot of new users are going to be using SpinRite 6.1 and have questions. Or maybe not because it's pretty much the same as it was, it just works a lot better. So anyway, I just don't have any additional bandwidth available for new conversation opportunities. I doubt that Twitter can actually implode. It's, as you said, Leo, it's too big and too important. You know, I doubt that even Elon can or will kill it. You know, and I have an alternative means for communicating my and GRC's events to anyone who cares through good old email.

Leo: And I will extend this offer after Episode 999. You can always use us to tell the world. I would bet a lot of SpinRite users and owners listen to various other things we do. And we have a lot of different channels, including Twitter channels.

Steve: And Leo, we still have two years. Who knows, two years from now, what'll be going on.

Okay. Someone said, where did his name go? I didn't have his name here. Shoot, I think it was Walt. Anyway, he said: "Steve, did you see there's a 'Project Hail Mary' in IMDB?"

He said: "Crossing my fingers." Anyway, indeed there is. A "Project Hail Mary" movie is in the works.

Leo: Well.

Steve: It is currently flagged as in "in development."

Leo: Well, if you had listened to our interview with Andy Weir some months ago when it came out, he had already optioned it.

Steve: Wow.

Leo: And he told me, and I wasn't too thrilled, I don't know how well I hid my discomfort, that Ryan Gosling had signed on for the lead.

Steve: Yes, I saw that, too. I saw that, yeah.

Leo: And I went, oh. Okay, yeah. But yeah, we're going to, you know, Andy was going to be on some months ago, but he'd just had a baby. We'll get him back on. And by the way, Daniel Suarez has a new book. The sequel to his "Delta" book is coming out soon, I think next January.

Steve: We had a lot of fun reading those.

Leo: Those were great.

Steve: Yeah.

Leo: So we'll get him on, too. So, yeah, we'll keep an eye. I'll have Andy on long before a movie gets made. We'll get the latest on that one.

Steve: Okay. So speaking of books we've loved. So many people have written to me, telling me that they're loving the Silver Ships series, that I want to share a tweet I received two days ago from the first person I know, or we know, who has finished the entire 24-book series. I was horrified, as I started to read the tweet, that he might have written something of a spoiler, but that concern was misplaced.

So here's the content of the DM that Bob Grant sent. He wrote: "Wow, wow, wow. Superb ending to the series. There was enough great writing and new intrigue in the first part of this final book in the Silver Ships series to be a great book in and of itself. However, the wrapping up of all the various storylines from the previous 23 books," and he says, "(20 Silver Ships and the related four Pyrean books) at the end was superb. There were joyful and poignant endings to each of the major characters from the books. I have to say that this is the best series I've ever read. Not to take away from Weber's Honorverse" - and of course he's talking about David Weber's Honor Harrington series that was one of the early series that we talked about on this podcast - "or Ryk Brown's Frontiers series," he says, "both of which I've enjoyed. But these 24 books have been a joy to read from beginning to end."

And then he said: "After a little break to catch up on some other reading, I plan to start the new Scott Jucha series called 'Gate Ghosts' whose first book is 'Axis Crossing.'" And as I mentioned to you, Leo, there are six more in that series after these 24. So anyway, obviously Bob has been following along with my previous reading discoveries. He knows of and read David Weber's Honorverse series and Ryk Brown's work in progress, Frontiers Saga series. And for what it's worth I'm in complete agreement with him about this being the best series I've ever read. I'm at the start now of book 19 of those 24 so I have six to go. And having already made this large investment in this series, I'm delighted to learn in advance that it ends wonderfully. So anyway.

One last piece. Simon, he said: "Hi, Steve. Persistence paid off. I was able to disable one-time code 'feature,'" he has in quotes. He's talking about PayPal. He said: "You can call PayPal and ask to 'unconfirm' your phone number. It may impact use of the PayPal app. But as long as you do not confirm phone number, it will not text security codes." So that's very cool.

Leo: Wait a minute. Which is less secure? Having no two-factor or having SMS two-factor?

Steve: Oh. No, no, no. You can still use...

Leo: Oh, you still have Authenticator or a YubiKey.

Steve: Yes.

Leo: Oh, oh, oh.

Steve: You still - yes, yes, yes.

Leo: No, I did that on Twitter, too. You had to have SMS to enable 2FA on Twitter. But once you set up a key or an Authenticator, you could then disable it. So you're saying you did the same on PayPal.

Steve: Yes, although there is no UI for doing it.

Leo: Oh, interesting.

Steve: You need to contact them. You have to contact them and say please unconfirm the phone number. And that makes sense, right, because the phone number can go to somebody else.

Leo: I don't have it anymore, right.

Steve: Right. Anyway, it was Simon who originally noticed and communicated that it was always possible to cause PayPal to send an SMS code for account/password recovery. However, I should note, someone else sent me a note, and I apologize to that person for letting his name slip, but he sent me a note that if users set up their own personal account recovery questions, you know, those like who was your favorite high school teacher and what was the name of your first dog or whatever, if you set those up, they cannot be bypassed.

Leo: Ahhh.

Steve: So that's another solution: Deliberately choose impossible to guess, no matter how well someone knows you, account recovery questions, and assuming that that information is correctly provided, then you will be safe from hijacking because nobody else will know what it was that you set up.

Leo: It's just like three more passwords, basically.

Steve: Yeah, yeah.

Leo: Yeah.

Steve: Okay. Finally, I mentioned last week that I thought SpinRite's new AHCI driver was not working correctly. I was wrong about that. It was working correctly. It was the location in my code where I was taking the hash of SpinRite's results that was causing a false positive detection. So I found and fixed that and made some other final improvements. Then, as planned, I updated GRC's server to get it ready to manage all subsequent downloads of the pre-release testing versions of SpinRite that will be forthcoming. That work is finished, and the server has been restarted and is now standing by to make SpinRite available.

I have one final feature to add which came up about 10 days ago. SpinRite 6.1 has four levels, or degrees, of its operation. The first level never performs any writing to a drive under any circumstances. It's strictly read-only. I'm not sure why, but it always seemed like it ought to offer that, so it always has. The second level is allowed to perform data recovery, so it will selectively rewrite only those regions of the media that are in need of repair. Level three goes further. Since refreshing any drive's data is generally good for it, and that's because latent and evolving soft errors are completely hidden by all modern drives, level three always rewrites the drive's data as it's moving through the drive. And level four goes even further, writing inverted data, reading it back to verify it, then rewriting the original data and reading it back to make sure that it was written correctly.

Okay. I mention this because there are three classes of drives that I refer to as being "write-hostile," and should only be used under SpinRite's first two "read mostly" levels. Those drives are SSDs whose media we know is incrementally fatigued by writing to it, hybrid drives which incorporate an SSD on their front end to serve as a non-volatile cache, and SMR drives where SMR stands for Shingled Magnetic Recording. Shingling, exactly like it sounds, refers to the deliberate overlapping of adjacent tracks in order to push track density to insane levels. If you picture a shingled roof, you cannot change an embedded shingle without pulling up the shingle above it, and then the shingle above that one, and the shingle above that one, and so on. The same is true for SMR drives, which makes writing to them something you want to do as little as possible.

As I mentioned, this issue just came up in SpinRite's newsgroup discussion a couple of weeks ago. Since I want SpinRite to continue doing everything possible for its user, in this case warning them if they are about to perform a level 3 or 4 scan on any drive which should not be written to needlessly, I need to be able to detect that. But I didn't own any hybrid or SMR drives. So I immediately tracked some down on eBay, and those four drives have all arrived. The last two just came in yesterday's mail. So after today's podcast I'll be adding detection of those drive technologies to SpinRite so that it can take responsibility for warning its users if they're about to do something that they probably don't want to do.

And then, with that last bit of technology in place, as far as I know, SpinRite 6.1 will be ready to start its final stage of pre-release testing. And as for that I'm absolutely certain there are things I've missed, things I just can't see because I'm their author. But that's why we test. What I am confident of is that at this point so much testing has already been done, by far the bulk of the work, that there are no showstoppers remaining. It should be a matter of cleaning up debris. So by next week's podcast it will have been under test for - I'm hoping that this is a Thanksgiving present for our testers. So I should have a good calibration on where we stand.

Leo: Nice. Incidentally, "Project Hail Mary" is the Book of the Month for Stacey's book club in January. If you have read it or want to read it, that's a good book to read, and discussion in the club.

Steve: And we all read it and loved it. It was great.

Leo: Oh, it's a great book.

Steve: Yeah.

Leo: And if you can listen to the audiobook, there's some features the audiobook has that the written page cannot that makes it kind of fun, too. Anyway, it's good either way.

Now, whatever it is, I want to know, what is Wi-Peep? Little Wi-Peep.

Steve: Okay. Little Wi-Peep. So imagine a technology that allows someone walking past a multistory building or a drone flyby to accurately locate and pinpoint within that building or any other similar space, closed or open, with a positional accuracy of about a meter, the location of every WiFi device such as security cameras, and locks and switches, and anything else on WiFi. That capability, which jumps off the pages of science fiction movie scripts, is not only here now, but it costs about $20. The two researchers who figured out how to make this WiFi mapping technology real named it Wi-Peep. They presented their research during the recent ACM MobiCom '22 which was held last month, in October, in Sydney, Australia.

Here's how they described what they accomplished. They said: "We present Wi-Peep, a new location-revealing privacy attack on non-cooperative WiFi devices. Wi-Peep exploits loopholes in the 802.11 protocol to elicit responses from WiFi devices on a network that we do not have access to. It then uses a novel time-of-flight measurement scheme to locate these devices. Wi-Peep works without any hardware or software modifications on target devices and without requiring access to the physical space that they're deployed within. Therefore, a pedestrian or a drone that carries a Wi-Peep device can estimate the location of every WiFi device in a building.

"Our Wi-Peep design costs $20 and weighs less than 10 grams. We deploy it on a lightweight drone and show that a drone flying over a house can estimate the location of WiFi devices across multiple floors to meter-level accuracy. Finally, we investigate different mitigation techniques to secure future WiFi devices against such attacks."

Okay. So this has never been done before. The key components here are the non-cooperative nature and the fact that this is being done by a probe which is not on the WiFi network. So they set this up and framed the problem, explaining the problems they encountered and how each such problem was solved.

They said: "We live in an era of WiFi-connected TVs, refrigerators, security cameras, and smart sensors. We carry personal devices like smartwatches, smartphones, tablets, and laptops. Due to the deep penetration of WiFi devices into our lives, location privacy of these devices is an important and challenging objective. Imagine a drone that flies over your home and detects the location of all your WiFi devices. It could infer the location of home occupants, security cameras, and even home intrusion sensors.

"A burglar could use this information to locate valuable items like laptops and identify ideal opportunities when people are either not at home or away from a specific area, for example, everyone is in the basement, by tracking their smartphones or smartwatches. The promise of pervasive connectivity has been to merge our physical and digital worlds, but the leakage of such location information brings arguably the worst aspect of the digital world, pervasive tracking, into the physical world.

"In this paper, we show that there are fundamental aspects of the WiFi IEEE 802.11 protocol that leak such location information to a potential attacker. We demonstrate that it is possible to reveal accurate location of all WiFi devices in an indoor environment, A, non-cooperatively, without any coordination with WiFi devices or the access points; B, instantaneously, without waiting for devices to organically transmit packets; and, C, surreptitiously, without any complex infrastructure deployment in the surrounding. Our goal is to expose the security and privacy vulnerabilities of the 802.11 WiFi protocol by demonstrating a first-of-its-kind non-cooperative localization capability. We hope that our work will inform the design of next-generation protocols."

So they said: "We note that there's been much past work in WiFi-based positioning. However, such past work does not enable non-cooperative, surreptitious localization of WiFi devices. First, most of this work relies on cooperation from end devices, for example, the client needs to switch channels or physically move or share inertial sensor data. Second, state-of-the-art techniques, such as ArrayTrack, rely on antenna arrays with multiple antennas, that are typically bulky and cannot be easily carried by a person or a small drone. Deploying multiple such antenna arrays near a target building makes the attack less practical and easier to detect." And I don't know if they said, but way more expensive, obviously.

"Third, RSSI-based" - and remember that's Received Signal Strength Indicator - "RSSI-based techniques rely on fingerprinting or trained models that require physical access to the target space. Finally, most of these need client devices to continuously transmit WiFi packets or share their received WiFi packets by installing an application, an access we cannot assume for such privacy-revealing mechanisms."

So they say: "We present Wi-Peep, a system that is quick, accurate, and performs non-cooperative localization. It does not require any access to target devices or the network access points. It does not even need the attacker to connect to the same WiFi network. In our attack, the attacker, a lightweight drone or a pedestrian, passes by the house carrying a small WiFi capable device and estimates the location of all WiFi devices in the target environment. We exploit the design of the 802.11 protocol to first generate WiFi traffic from non-cooperative clients, then use a novel time-of-flight based-technique to locate these devices. Wi-Peep solves the following challenges."

Okay. The first challenge, generate WiFi traffic without cooperation. They explain: "We must, A, identify all devices in the network quickly at the start of the attack; and, B, generate WiFi traffic continuously from such devices to perform location estimation. A simple solution to identifying devices is to passively wait for WiFi devices to transmit a packet. This approach is problematic because it requires the attacker to linger around for a long time. Instead, we exploit the 802.11 power-saving mechanism, which is available in all 802.11 standards from 11a and b up through 11ax by injecting a fake beacon imitating the access point that tells all connected WiFi devices to contact the access point to receive buffered packets. This beacon elicits a response from all devices on the target WiFi network.

"Once we've identified all devices, we use targeted packets to each of these devices. To perform time-of-flight measurements on these devices, the attacker requires exchanging packets directly with target devices. Therefore, natural traffic from a target device cannot be used. Recent work has shown that 802.11 devices always respond to packets with an ACK, even when the packets emerge outside the WiFi network and are unencrypted or incorrectly encrypted. We use this flaw to perform time-of-flight measurements to any target device. The challenge in using WiFi is that WiFi devices are in the sleep mode most of the time, and their radios turned off. We have designed a technique that allows an attacker to keep the radio of target devices on during the attack so that they keep sending ACKs."

Okay. So basically what these guys did was to recognize there was a way to, after learning about the beacon in a residence or a corporate facility or wherever, to simulate a broadcast from the beacon which will induce all WiFi devices to respond. When they respond, they're going to get each device's MAC address. That then allows them to individually target those devices selectively and in real time, so basically they get an instant inventory, and then they switch into an active tracking mode where they're spewing out packets, measuring roundtrip time which they call "time-of-flight," in order to determine an instantaneous distance they are away from each of the devices.

And of course as they move, all of those various vectors are changing length. And by changing their path, they're able to infer where the device must be in order for its vector to have changed as it did over time. So then they explain the second problem they had was localization in the face of noisy, what they call SIFS, which is short for Short Interframe Space.

So they explain: "In 802.11, ACKs are sent at a fixed interval after receiving a data packet. This interval is called Short Interframe Space or SIFS as illustrated in" the figure that they have in their notes. They said: "Wi-Peep measures the roundtrip time between a packet transmission and an ACK reception and subtracts the SIFS. This allows Wi-Peep to estimate the time-of-flight and hence the distance between the attacker and the target device. Unfortunately, our experiments reveal that even though the WiFi protocol mandates SIFS to be 10 microseconds, in practice this delay can vary from 8 to 13 microseconds. Such errors can randomize the location estimation process. We build a new algorithm to correct for such variations in time-of-flight estimates."

And finally, dealing with multipath effects. They explain that the time-of-flight measurements are error-prone because multiple copies of a signal arrive at the receiver from multiple paths, you know, reflection of signals within an environment. They said: "The strongest path may not necessarily be the direct path. Since the attacker is far away and obstructed from the target, this problem is further exacerbated. Indeed, our measurements reveal that Wi-Peep's individual time-of-flight measurements are error prone for this reason. To counter this challenge, we take [what they call] the 'wisdom-of-the-crowd' approach. Even though each measurement is noisy, Wi-Peep involves quick packet-ACK sequences at the millisecond level." So they're doing thousands per second.

"Therefore, we can collect hundreds of measurements as the attacker flies by, or walks by, the target. We exploit the spatial diversity of these measurements to get an accurate position estimation of our targets." So, you know, that's a brilliant and completely workable solution. Individual measurements are noisy, but the truth can be found by sorting through thousands of measurements made over time from different positions.

And then they talk about their implementation. They said: "We've implemented our design on an ultra-light DJI mini 2 drone" - you probably have one, Leo - "using off-the-shelf..."

Leo: Well, I have the mini 3, but okay. Is there something finally I can do with it?

Steve: Yeah.

Leo: Yeah. There's a picture in their paper of it. It's kind of cool, yeah.

Steve: Yeah, it's neat. Yeah, it's sort of like stuck on the front of it.

Leo: Yeah, I don't know how well it'd fly with that on there, but I guess it's not too heavy.

Steve: They managed to do it. Anyway, they said: "Using off-the-shelf ESP32 and ESP8266 WiFi modules. Our hardware weighs 10 grams and costs less than $20. It can be deployed on lightweight drones or carried by a person. Our evaluations in a real environment shows that Wi-Peep finds the location of target devices in an 802.11ax WiFi 6 network on three different floors of a house with a median error of 1.2 meters in around two minutes.

"The contributions of this paper are: We present a new way for 802.11 protocol features to perform time-of-flight-based positioning of WiFi devices without having any control over target devices. We find that many devices deviate from the standard time for SIFS which creates a challenge for localization. We design a localization technique that finds a target device without knowing the exact SIFS used by the device. We present a solution for future WiFi chipsets that allows authenticated devices to perform localization, while disabling non-cooperative attacks."

So consider these fact which they then enumerate: "The Wi-Peep attacks work with any WiFi device without instrumentation, in other words, without any application or firmware-level changes. It does not need physical access to the enclosed physical space and does not need to break the encryption of the WiFi network. Once the target MAC address is obtained, the target device doesn't even need to be connected to WiFi. Due to the ease of attack, Wi-Peep has many privacy and security limitations," they write. "We list some example implications below. In these scenarios, we assume that it is common for a person to carry a WiFi-capable device such as a smartphone or a smartwatch. Also note that the type of device - iPhone versus smart sensors - can be identified through various means like the vendor specific information in the MAC address."

Okay. So, and they give us four examples, one impacting security. "An attacker can track the location of security guards inside sensitive buildings, for example, banks, if they carry a smartphone or a smartwatch. And notice that this is real-time. So moving targets are fine. They will get real-time feedback as things move within the area that they're surveilling. A privacy implication, an eavesdropper can fly a drone over a hotel to find the number and types of rooms currently occupied. This can be done by a rival hotel trying to find detailed information of how target business is performing. WiFi devices that belong to a room such as smart TVs can be filtered based on MAC addresses. If other devices such as tablets and laptops are found in a room, it can be considered occupied. And this can be done in the middle of the night when most guests are in their rooms."

Or a privacy/security implication. "If the MAC address of a device that belongs to a person of interest is known, Wi-Peep can track that person individually in a crowd..."

Leo: Oh, that's scary.

Steve: Uh-huh, "...or inside a building like a shopping center or an airport, even when their device is not connected to any WiFi network."

Leo: So this is - so you could tail somebody with one of these in your pocket.

Steve: Yup.

Leo: Oh, that's interesting, yeah.

Steve: Security: "Wi-Peep could be used by burglars to find out the occupancy status of specific parts of a building. For example, the burglar can find out all the people are on the second floor, and the basement is empty. Wi-Peep can also be used for positive use cases." And I like this. "For example, in a hostage situation, the police can fly a drone over the building to find out where the hostages are kept because many hostages might have smart devices on them, and they would be collected together in a dense group and not moving. It might also be possible to track the attackers, as well."

Okay. Anyway, through the balance of their paper, which is lengthy, they proceed to deal with every aspect of their system and present its solution. So my point is, the method to do this today is now in the public domain. So anyone who wants to do it, and has the skill set to replicate their work, can. You know, I could do that. Many of our listeners could do that. And I would not be surprised if we didn't eventually see an off-the-shelf turnkey Wi-Peep mapping system that would allow anyone with only a few dollars to spare to obtain this potentially powerful remote WiFi mapping capability, very much the way script kiddies are using scripts that they were unable to write.

Until now, we've had a general sense that the goings on inside our homes and offices were at least moderately private. The idea that someone standing outside in the middle of the night could first take a complete inventory of all WiFi devices within the area -

non-cooperatively, without connecting to or knowing our network's password - and then determine the approximate location of every one of those devices, whether they are upstairs or downstairs, and generally where, might not be unsettling to some people. But there are likely some situations and installations where having such knowledge in real time could be very valuable to the wrong people.

The authors spend some time near the end of their paper talking about possible future mitigations. And the overall outlook there is bleak. The bad news is that since this is a hardware-level attack which only leverages standard WiFi features which are implemented in the core WiFi silicon, nothing can be done in firmware or software. All WiFi chips today will and do respond to the probe request packets sent during the use of this technology. It will take a future generation of WiFi chips to deliberately break the WiFi specification or the spec to be updated in order to sanction this by not replying within a microsecond or two, but by deliberately randomizing the Short Interframe Space interval so that time-of-flight information cannot readily be determined. Doing that will allow WiFi to work, while still making location impossible.

Leo: That is why Apple randomizes MAC addresses on its iPhones, though. I wonder if that is effective as a countermeasure.

Steve: Actually, that's different than this. This doesn't need MAC addresses.

Leo: But if I were following you around, I would need your MAC address to know it's you. I'm not thinking about the mapping feature.

Steve: Ah, that is true, yes, yes, yes. That is true. MAC addresses, as we know, are fixed when the phone is attached to a network.

Leo: Right.

Steve: They're only randomized when it's not - when it hasn't joined a network. Once it has, then it uses its actual MAC address. But you're right, following you around, the MAC address - I forgot exactly what the algorithm is.

Leo: I think they change it every 15 minutes.

Steve: Okay.

Leo: And I wonder if, you know, since you know it's him for 15 minutes, and then the MAC address changes, there might be some way to say, ah, yeah, he's just changed his MAC. I don't know.

Steve: Not in a crowd.

Leo: Not in a crowd.

Steve: Because you would be, yeah, because you would be getting - you would be - so first of all you would be - you'd be only pinging that, and then suddenly there would be no reply.

Leo: Yeah, yeah, right.

Steve: So it would then go dead, and you'd have to go back...

Leo: You would have lost it by that point, yeah.

Steve: Yeah, you have to go back into broadcast mode in order to get replies from everybody in the neighborhood.

Leo: I'm less concerned about somebody mapping my house.

Steve: I knew you wouldn't be.

Leo: For WiFi access points. But the tracking thing is concerning. I think others, I think there's Android phones that also randomize MAC addresses.

Steve: Well, remember that it's not the WiFi access point that they're locating, it's all your security cameras.

Leo: Right, right. Anything WiFi; right.

Steve: Anything WiFi.

Leo: Yeah, yeah, yeah. Again, less worried about that. But the tracking thing is a real cause for concern. But it must have been a threat anyway. That's why they're randomizing MAC addresses, I would think. There must be other reasons; right?

Steve: Apple did that for privacy.

Leo: For privacy; right.

Steve: Yeah, yeah.

Leo: Yeah. It's interesting. Really clever. The multitask thing is what I find most interesting. The algorithm to get around that. Fascinating, yeah. Wi-Peep. Little Wi-Peep.

Steve: Wi-Peep.

Leo: The paper is in the show notes, if you want to read it. It actually makes pretty good reading. Mr. Steven "Tiberius" Gibson, again we have come to the end. And every time we come to the end of a show, I think we're one step closer to Episode A00. But it's okay, man, it's okay. Because I'll survive. We will survive.

Steve: Yeah, Leo, we've got two years.

Leo: Two years.

Steve: Let's not worry about it yet.

Leo: Two years is a lifetime.

Steve: Lots can happen.

Leo: For a mayfly. Steve is at GRC.com. That's his website. That's, of course, where SpinRite lives, the world's finest mass storage maintenance and recovery utility, GRC.com. It's also where you'll find copies of this show. It's one of the places, but it is the place to find the 16Kb version for the bandwidth-impaired, and the beautifully handcrafted, human-crafted transcripts. He also has a 64Kb audio. You can leave feedback at the website, GRC.com/feedback. Those SpinRite forums are there. Lots of other stuff, including ShieldsUP. There's so much great stuff. GRC.com. We have the show at our site, as well, of course, TWiT.tv/sn. There's a full-time YouTube channel dedicated just to Security Now!, all 898 episodes. And of course you can subscribe, and that way you'll get it automatically whenever there's a new one, like right now.