Security Now! #898 - 11-22-22

Wi-Peep

This week on Security Now!

This week we note that Firefox moved to v107 and that Google recently reached a nearly $400 million dollar user-tracking settlement. Red Hat has started cryptographically signing its ZIP distributions, the FBI purchased the nefarious Pegasus spyware and Greece paid €7 million 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 US 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.

Keep that dirt strapped down!

Security News

Firefox v107 was released last Tuesday.
Nothing earth shattering here. No critical security fixes, but a large and welcome collection of high severity and some moderate severity repairs. This update appeared to be released to fix those things, since there were not otherwise a large number of new features. When I checked, that's the version I was already running. So I imagine that's likely true for everyone else.

Google settles for a cool $391.5 million
Google recently settled a multi-state data privacy lawsuit which had been brought against it by a team of 40 states. We talked about this four years ago, back in 2018, when the offices of those 40 states' Attorneys General sued Google, alleging that Google had been lying and misleading users into thinking they had disabled location tracking in their account settings. The lawsuit followed some reporting 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.

In the settlement Google agreed to pay $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 wheels of justice, when they don't completely fall off the wagon, do tend to turn slowly in the US. It took four years to get here. The other things we learn, thanks to Google's posting about this, is what they have since changed.

In their posting last week titled "Managing your location data", which brings new meaning to the phrase "putting on a happy face", Google 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, location information helps connect experiences across Google to what's most relevant and useful.*
>
> *Over the past few years, we've introduced more transparency and tools to help you manage your data and minimize the data we collect. That's why we:*
>
> - *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 3, 18 or 36 months worth of data at a time.*
> - *Developed easy-to-understand settings like Incognito mode on Google Maps, preventing searches or places you navigate to from being saved to your account.*
> - *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.*

These are just some ways that we have worked to provide more choice and transparency. Consistent with those improvements, we settled an investigation with 40 U.S. state attorneys general based on 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. The updates include:

- 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.

- 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 History data to their account.

- Updated 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.

Today's settlement is another step along the path of giving more meaningful choices and minimizing data collection while providing more helpful services.

It must be truer than I am able to understand, that the more information an advertiser has about someone, the more revenue is generated by showing that person advertisements. As our listeners know, I've always been skeptical of that. But 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. They wouldn't be risking our wrath if it weren't really valuable to them.

Red Hat Signing its ZIP file Packages

I caught wind of a mention that Red Hat had started cryptographically signing its deployment ZIP files. That made me curious since I hadn't ever heard of ZIP files being cryptographically signed. With all problems we've been seeing with supply chain poisoning, obtaining verifiable assurance of an archive's unmodified authenticity would be great. A cryptographic signature would do that.

And cryptographic signing certainly makes a lot more sense that 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 a site, what would keep them from also updating the hashes shown at the same site as proof of the file's authenticity? Talk about a false sense of security! *"Oh, yeah, don't worry! I just verified that the hashes match. It's got to be okay!"* Yeah. Right.

https://www.redhat.com/en/blog/cryptographic-signatures-zip-distributions

Anyway, I looked into what was going on and found a posting by RedHat titled "Cryptographic signatures for zip distributions". Paraphrasing and removing some of the over-simplified descriptions, 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 are uploaded to our CDN and made available to customers. We have taken advantage of this aspect of our build process and extended it by combining all the hashes for a particular release and packaging them into a "SHA256SUM" file.*
>
> *This file is in a standard format that lists the hash and the corresponding filename of the particular artifact. It is commonly used across the industry to provide integrity to binary files. However, it is not limited to that. The `sha256sum` command on RHEL, other Linux distributions and macOS natively supports this file format.*
>
> *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 artifacts enumerated within that file. The "gpg" command on RHEL, 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 4096-bit private key, and our public keys are available on our website and the Massachusetts Institute of Technology (MIT) public key server.*

Red Hat's mention of a detached signature simply means that the signature itself resides in a separate file. The signature is just a SHA256 hash of the file it's signing which is then encrypted under Red Hat's super-secret private key which they are careful not to let loose. Just like my GRC code signing keys, it probably resides in an HSM – a hardware security module. So there's no reason for that signature file not to stand alone so long as the file it's a signature for, is clear.

Overall, this is a welcome move and as a deterrent to the abuses of today's supply chain, it's probably where the open source community will need to go. The glitch here is that 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 here. Let's Encrypt offers no guarantees about the identity of a site. It provides domain validate certs where the only requirement is for the certificate to match the server's domain name. Specifically, it does not offer OV — Organization Validation — certificates. In order to issue OV certs, a certificate authority must perform some reconnaissance to

positively verify the identity of the entity requesting the certificate. And what's more, many open source projects are just some guy working alone without any organization.

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 some much more rigorous multi-factor authentication.

One way or another, we need some solution to the current supply-chain pollution problem. So the application of a bit of crypto might be a place to start. Hat's off to Red Hat.

The FBI purchased Pegasus for "research and development purposes".
Last week the New York Times ran a story with the headline *"Internal Documents Show How Close the FBI Came to Deploying Spyware"* and what spyware would that be you ask? The New York Times reported that last December, FBI director, Christopher Wray, told Congress that the bureau purchased the infamous Pegasus phone hacking tool for "research and development purposes". Uh huh. It turns out that FOIA — the U.S. Freedom of Information Act — can be quite handy for figuring out what 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, the director of the FBI, was asked whether the bureau had ever purchased and used Pegasus, the hacking tool 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. 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. 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.*
>
> *It is 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 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 his office said: *"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."*

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.

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, 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, 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 under our constitutional protections, I think that it would be a useful tool to empower their criminal investigations.

Greece bought Predator for €7 million:
And on the subject of spyware abuse, a recent report in the Greek press claimed that Greece's government paid €7 million to Intellexa for access to the Predator surveillance and spyware platform, and an additional €150,000 for the ability to rotate ten new targets per month. This 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.

Again, it seems to me that the problem is less about the tool than with how it's used. It's just 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 will just force it underground.

A passkeys support directory
1Password has added support for passkeys to its password manager. And in a nice promotion of passkeys, they've created a community supported online directory listing online services currently supporting passkey authentication. It's at: https://passkeys.directory/ The directory currently lists 43 companies and their URLs, though, some are flagged as MFA so I suspect that they might not be pure passkeys logon. Some notable names on the list, which do appear to be pure passkeys authentication without the MFA tag, include:

- 1Password Passkeys Demo — future.1password.com
- Best Buy — bestbuy.com
- Carnival Cruises — carnival.com
- eBay — ebay.com
- KAYAK — kayak.com
- Microsoft — microsoft.com
- Nescafe — nescafe.com
- Nvidia — nvidia.com
- PayPal — paypal.com
- Robin Hood — robinhood.com

Since I just discovered this, I haven't made any time to experiment with it and explore. But since I'm an avid buyer on eBay (often of specific old hard drives) I ought to be able to give it a try.

Quantum decryption deadline

From the "Having fun with bureaucracy department" comes an edict from the OMB. The US's Office of Management and Budget has ordered federal agencies to scan their systems – oh, yes, scan those puppies carefully – and provide an inventory of assets containing cryptographic systems that could be cracked by quantum computers in the coming years.

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 is quantum resistant. So the OMB could have simply said, give us a list of all of your computers.

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 cryptography. Oh... yes, factoring the number 27, we can do that. It's magic. But the number 35? Nahh we're not quite there yet. Give us another ten years or so.

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, just as soon as we're absolutely sure that we're not going to be making a big mistake. Remember that one of the candidates that had already been chosen was recently cracked by conventional computers. So, it would be better that we stay with what we know we can't crack today before moving prematurely to something that we presume some future non-existent mythical quantum computer should also be unable to crack.

The OMB edict stated that Federal agencies had until May 4, 2023. And the NSA ordered that all government agencies handling classified information **must** use quantum-resistant encryption by 2035. Hmmmm. So that's 13 years from now. So we ought be up to factoring 45 by then.

Attorneys General ask the FTC for online privacy regulation:

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 globally it's a total uncoordinated 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 US, are able to openly attack the networks of US corporations and its private citizens. And, yes, there's reciprocity, the US is able to do the same to them — which doesn't make any of this sane. 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 have watched being born through these past 30 years.

I bring this up because various democracies around the world, notably the EU and the US 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. 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 wouldn'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 US Federal Trade Commission, our FTC, to pass legislation which would regulate online data collection practices. These AGs said they are *"concerned about the alarming amount of sensitive consumer data that is amassed, manipulated, and monetized,"* and that they regularly receive inquiries from consumers about how their data is being hoarded and abused.

Since we have a bit of time left today, and since I think this is extremely important, I want to first share the introduction in the letter that was submitted to the FTC:

> *We, the Attorneys General of 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. 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.1 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 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.

The letter then goes into 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'll share:

Location Data
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.

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 (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.

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 driver's licenses, census data, birth certificates, marriage licenses, and voter registration information. Data brokers also 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.

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 the 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.

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 that being a danger. Much of what is going on today is only able to happen under the cover of darkness, because consumers are blissfully unaware. 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 surprise there.

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 we're moving in the right direction.

Closing The Loop

vincent stacey / @vincent_stacey

Hi Steve, pfsense installs its own version of Linux and won't have the default users of another distribution.

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.) The main reason why a ZimaBoard would not be my first choice as a router, is that unless a network expansion board is plugged into its PICe x4 slot, it only has a pair of LAN NICS and I would expect a router to have a few more NICS for implementing useful multi-network isolation.

Charles Turner / @capturner

As possible fodder for the listener feedback section of a future episode the Security Now! podcast, I have a question arising from the discussion you and Leo had on Tuesday (November 15, 2022) during SN #897 (Memory-Safe Languages). 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 as beyond or b) becoming a $44 billion version next iteration of Myspace and FTX.

It's clear that Twitter is currently in turmoil. And I don't have any firsthand sense for just how

fragile Twitter's technology is internally. It seems to me that matters a lot. If the previous regime engineered really solid bullet-proof systems then it ought to be able to withstand Elon's shaking its foundation. But overall I'm a big believer in inertia and in things generally changing much more slowly than we expect. Of course, Elon could trip over the main power cord and Twitter would go dark until someone plugged it back in. And I suppose that I'm interested in what Elon is doing there, since he's an interesting character and somehow he has managed to get other people 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. Truly astonishing technology. And it's Elon's SpaceX StarLink technology – which actually works – that's enabling Ukraine to survive Russia's increasingly aggressive attacks on its infrastructure. Again, thanks Elon.

Mostly, though, I think Elon is just having fun with his life, as is his right. I hope he is. 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.

Anyway, this has all made me curious about what he **is** doing with it. I pick up little bits here and there, but I don't follow news feeds, or even Twitter, because they interrupt my train of thought. So it was with some joy that I stumbled upon a site, which I figured had to exist somewhere, called: *"Twitter Is Going Great (dot) Com"* https://twitterisgoinggreat.com/ And yes, of course, it's offering up its share of schadenfreude, so keep in mind that it's heavily 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: *"Twitter Is Going Great (dot) Com"*.

leslie macfarland / @lemacfarland

> *Hi Steve, if Twitter implodes, are you going to mastodon or somewhere else? Your SN podcast is top notch security and quality.*

In order to get the word out to 18 years worth of SpinRite owners, I'll shortly be setting up an old-school eMailing facility. And one of the several lists I'll be maintaining will be for SN listeners who would like to subscribe to the weekly links to the show notes and a description of each week's topics. It'll be nice to have more than 280 characters for that.

As for Mastodon, I don't know. I'm really 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 getting actual work done. And now we have GitLab for managing SpinRite bugs and feature requests. And I have GRC's forums which will soon be quite active, since that's where SpinRite's tech support will be hosted and happening. So I don't have any additional bandwidth available for new conversation opportunities. I doubt that Twitter can actually implode. It's too big and too important. I doubt that even Elon can kill it. And I will have an alternative means for communicating my and GRC's events to anyone who cares... using the one thing we all have, which is eMail.
Walt Stoneburner @waltomatic

Steve, did you see there's a Project Hail Mary in IMDB? . . . crossing my fingers.
https://www.imdb.com/title/tt12042730/

Wow. It's shown as "in development" but some slots are assigned. That would be amazing!

And speaking of books we've loved. So many people have written to tell me that they're loving the Silver Ships series that I want to share a tweet I received from the first person I know who has finished the entire 24-book series. I was horrified that he might have written something of a spoiler, but that concern was misplaced. Here's the content of the DM that Bob Grant sent:

Bob Grant / @swguru

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 (20 Silver Ships and the related 4 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 or Ryk Brown's Frontiers Saga, both of which I've enjoyed... but these 24 books have been a joy to read from beginning to end.*

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.

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 of book 19 of 24, so I have six to go. And having already made this large investment I'm delighted to learn that the series ends wonderfully.

Simon / @Talk_2Simon

Hi Steve, persistence paid off :-) I was able to disable one time code "feature". You can call PayPal and ask to "un-confirm" phone number. It may impact use of PayPal app, but as long as you do not confirm phone number, it will not text security codes.

That's very cool. 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. Someone else send me a note that if a use were to set their own personal account recovery questions into their account, THEY would not be bypassed. So that's another solution: Deliberately choose impossible to guess account recovery questions and, assuming that this information is correct, you should be much safer from hijacking.

SpinRite

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 pre-release testing versions of SpinRite. That work is finished and the server is standing by to make SpinRite available.

I have a final feature to add which came up about 10 days ago: SpinRite v6.1 has four levels, or degrees, of 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 has always seemed like it ought to offer that. 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: Refreshing any drive's data is generally good for it, because latent and evolving soft errors are completely hidden by all modern drives. So level three always rewrites the drive's data. And level four goes even further, writing inverted the data, reading it back to verify it, then rewriting the original data.

I mention this because there are three types of drives that are truly 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, 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 it and the shingle above it, 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. 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. 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 arrived 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.

And then, with that last little bit of technology in place, as far as I know, SpinRite v6.1 will be ready to start its final stage of pre-release testing.

I'm absolutely certain that there will be 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 no showstoppers remain. This will be a matter of cleaning up debris.

By next week's podcast I'll have a good calibration on where we stand.

Wi-Peep

Imagine a technology that allows someone walking past a multi-story building, or a drone fly-by, to accurately locate and pinpoint within that building or any other space, closed or open, with a positional accuracy of one 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 $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 Sydney, Australia. Here's how they describe what they accomplished:

We present Wi-Peep – a new location-revealing privacy attack on non-cooperative Wi-Fi devices. Wi-Peep exploits loopholes in the 802.11 protocol to elicit responses from Wi-Fi 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 are deployed in. Therefore, a pedestrian or a drone that carries a Wi-Peep device can estimate the location of every Wi-Fi device in a building. Our Wi-Peep design costs $20 and weighs less than 10 g. We deploy it on a lightweight drone and show that a drone flying over a house can estimate the location of Wi-Fi devices across multiple floors to meter-level accuracy. Finally, we investigate different mitigation techniques to secure future Wi-Fi devices against such attacks.

They then set this up and frame the problem, explain the problems they encountered and how each such problem was solved...

We live in an era of Wi-Fi connected TVs, refrigerators, security cameras, and smart sensors. We carry personal devices like smartwatches, smartphones, tables, and laptops. Due to the deep penetration of Wi-Fi devices in 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 of your Wi-Fi 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 (e.g., 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 – to the physical world.

In this paper, we show that there are fundamental aspects of the Wi-Fi (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 Wi-Fi devices in an indoor environment (a) non-cooperatively – without any coordination with Wi-Fi 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 Wi-Fi 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.

We note that there has been much past work in Wi-Fi-based positioning. However, such past work does not enable non-cooperative, surreptitious localization of Wi-Fi devices. First, most of this work relies on cooperation from end devices – e.g., 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. Third, RSSI-based (received signal strength indicator) 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 Wi-Fi packets or share their received Wi-Fi packets by installing an application, an access we cannot assume for such privacy revealing mechanisms.

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 point. It does not even need the attacker to connect to the same Wi-Fi network. In our attack, the attacker (e.g., a light-weight drone or a pedestrian) passes by the house carrying a small Wi-Fi capable device and estimates the location of all Wi-Fi devices in the target environment. We exploit the design of the 802.11 protocol to first generate Wi-Fi traffic from non-cooperative clients, then use a novel time-of-flight based technique to locate these devices. Wi-Peep solves the following challenges:*

1. Generate Wi-Fi traffic without cooperation – *We must (a) identify all devices in the network quickly at the start of the attack, and (b) generate Wi-Fi traffic continuously from such devices to perform location estimation. A simple solution to identifying devices is to passively wait for Wi-Fi 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/b to 11ax) by injecting a fake beacon imitating the access point that tells all connected Wi-Fi devices to contact the access point for buffered packets. This beacon elicits a response from all devices in the target Wi-Fi network. Once we have 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 Wi-Fi network and are unencrypted/ incorrectly encrypted. We use this flaw to perform ToF measurements to any target device. The challenge in using Polite Wi-Fi is that Wi-Fi devices are in the sleep mode most of the time and their radio is 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.*

The next problem they encountered and solved was Localization despite noisy SIFS:

2. Localization despite noisy SIFS – *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 Figure 1. Wi-Peep measures the round trip time between a packet transmission and ACK reception and subtracts the SIFS. This allows Wi-Peep to estimate the time-of-flight and hence the distance between the attacker and a target device. Unfortunately, our experiments reveal that even though the Wi-Fi protocol mandates SIFS to be 10 μs, in practice, this delay can vary from 8 to 13 μs. 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:

> ***3. How to deal with the multipath effect –*** *The ToF measurements are error-prone because multiple copies of a signal arrive at the receiver from multiple paths. 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 the 'wisdom-of-the-crowd' approach. Even though each measurement is noisy, Wi-Peep involves quick packet-ACK sequences at the millisecond level. Therefore, we can collect hundreds of measurements as the attacker flies by (or walks by) the target building. We exploit the spatial diversity of these measurements to get an accurate position estimate of our targets.*

That's brilliant and completely workable. Individual measures are noisy, but the truth can be found by sorting through thousands of measures made over time from different positions.

And then they talk about their implementation:

> *We have implemented our design on an ultra-light DJI mini 2 drone using off-the-shelf ESP32 and ESP8266 Wi-Fi modules. Our hardware weighs 10 grams, and costs less than 20 dollars. 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 Wi-Fi 6 network on 3 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 using 802.11 protocol features to perform time-of-flight based positioning of Wi-Fi 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 factors which they enumerate:

> *The Wi-Peep attack works with any Wi-Fi 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 Wi-Fi network. Once the target MAC address is obtained, the target device doesn't even need to be connected to Wi-Fi. Due to the ease of attack, Wi-Peep has many privacy and security implications. We list some example implications below. In these scenarios, we assume that it is common for a person to carry a Wi-Fi capable device such as a smartphone or a smartwatch. Also, note that the type of a device (e.g. iPhone vs smart sensors) can be identified through various means like the vendor specific information in the MAC address.*
>
> *• [Security] An attacker can track the location of security guards inside sensitive buildings (e.g. banks) if they carry a smartphone or a smartwatch.*

• [Privacy] 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 the 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. This can be done in the middle of the night when most guests are in their rooms.

• [Privacy/Security] If the MAC address of a device that belongs to a person of interest is known, Wi-Peep can track that person in a crowd or inside a building like a shopping center or an airport (even when their device is not connected to any Wi-Fi network).

• [Security] Wi-Peep could be used by burglars to find out the occupancy status of specific parts in a building. For example, the burglar can find out that all people are on the second floor and the basement is empty.

Wi-Peep can also be used for positive use-cases. 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 attacker(s) as well.

Through the balance of their paper they proceed to detail every aspect of their system and solution. 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. 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 obtain this potentially powerful remote WiFi mapping capability.

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 an area – non-cooperatively without connecting 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 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 by not replying within a microsecond or two, but by deliberately randomizing the chip's Short InterFrame Space interval so that time-of-flight information cannot readily be determined. Doing that will allow WiFi to work, while making location impossible https://randompaper1234.tiiny.site/