Security Now! #1052 - 10-18-25

Global Cellphone Tracking

This week on Security Now!

• Apple introduces a new Digital ID inside Wallet. • Checkout.com refuses to pay a ransom demand. • Google announces "Private AI Compute" in the cloud. • Google backpedals on their "devs must register" demand. • Win11 added a Passkeys API which 1Password & Bitwarden support. • Russia tracks SIM card appearances to thwart drone usage. • Google sues Chinese Phishing as a Service platform. • Lots of interesting listener feedback. • Global cellphone tracking is alive, well, malware free and a distressingly common commercial enterprise.

How bad must things become before you decide to stop and refactor the code?

Security News

Apple adds a Digital ID

Last Wednesday I received a notice on my older iPhone 12 which I had upgraded to iOS 26. The notice was announcing something called Apple's "Digital ID" which was now ready for me. The hook for me was that while the announcement was focused upon this new Digital ID's use as a replacement for the RealID which U.S. TSA agents want to see at airports, the announcement also noted the app's use for age verification.

At this time, Apple's new Digital ID is tied to a passport. Fortunately, I happened to have one, which I had originally obtained when I was joining Leo in Toronto for appearances on Call-4-Help, which I later renewed for the OWASP SQRL presentations in Sweden and Ireland.

The process that I went through to establish the Digital ID was fascinating. The Digital ID app first required me to aim the iPhone's camera at the photo page of the passport, whose image it acquired and processed. Then it had me scan the RFID chip that's embedded in the back cover of the password. The app showed me where to position the phone and, sure enough, it made happy sounds and acquired the digital data from the passport. Presumably the chip contains much of the same data as the visual photo page. Finally, the app required me to follow its step-by-step instructions in "selfie mode" to prove that I was alive and looked like my passport photo. I was instructed to position my face in the frame, then to look at the screen. Then it told me to close my eyes until the phone vibrated in my hand. So I did that and after a few moments it was satisfied and vibrated. Next, it told me to smile until the phone vibrated again, which I did. And, finally, I was told that everything had been done and that my identity was being verified and I'd be notified once the verification was complete.

I was poking around in the app when I noticed that the process appeared to finish, and I was invited to store my shiny new passport-based Digital ID in the phone's wallet.

Last Wednesday Apple's announcement was headlined: *"Apple introduces Digital ID, a new way to create and present an ID in Apple Wallet"* with the tag line: *"Digital ID offers a secure and private way for users to create an ID in Apple Wallet using information from their U.S. passport, and present their ID with iPhone or Apple Watch"*

> *Apple today announced the launch of Digital ID, a new way for users to create an ID in Apple Wallet using information from their U.S. passport, and present it with the security and privacy of iPhone or Apple Watch. At launch, Digital ID acceptance will roll out first in beta at TSA checkpoints at more than 250 airports in the U.S. for in-person identity verification during domestic travel, with additional Digital ID acceptance use cases to come in the future.*
>
> *Digital ID gives more people a way to create and present an ID in Apple Wallet even if they do not have a REAL ID-compliant driver's license or state ID. Digital ID is not a replacement for a physical passport, and cannot be used for international travel and border crossing in lieu of a U.S. passport.*
>
> *Jennifer Bailey, Apple's vice president of Apple Pay and Apple Wallet said: "With the launch of Digital ID, we're excited to expand the ways users can store and present their identity — all with the security and privacy built into iPhone and Apple Watch. Since introducing the ability to add a driver's license or state ID to Apple Wallet in 2022, we've seen how much users love having their ID right on their devices. Digital IDs brings this secure and convenient option to even more users across the country, as they can now add an ID to Wallet using information from their U.S. passport."*

The launch follows the capability for users to add an eligible driver's license and state ID to Apple Wallet. If users do not have a U.S. passport to create their Digital ID, they can still add an eligible driver's license to Apple Wallet.

Jake Peterson, Lifehacker's senior technology editor, offered a more balanced and less Apple-centric view. He wrote:

Back in 2021, Apple announced a new feature for the Wallet app, that allowed users to add their driver's licenses or state IDs to their iPhones. To me, it sounded like the beginning of the end for physical wallets. In reality, it was anything but: Not only are the applications limited, but even after all this time, only 12 states and Puerto Rico actually support the feature.

While the rest of us wait for our respective states to get on board, many might have another option for these virtual documents. On Wednesday, Apple announced "Digital ID," a new initiative that lets you create an ID in the Wallet app using your passport. This bypasses the waiting period for the 38 states that don't yet support these ID features: If you have a passport, you can try this feature out today. Even if your state supports driver's license and state ID uploads to the Wallet app, you'll miss out on features if you don't have a REAL ID. If you have a passport, however, you can use it instead, which opens up the wallet ID feature to even more users than before.

Like previous attempts at virtual IDs, however, don't expect to be able to use this Digital ID just anywhere you'd normally show documentation. Right now, the main use for Digital ID is for flying: According to Apple, Digital ID is launching in beta at over 250 airports to be used at TSA checkpoints. Importantly, this feature only supports domestic flights, even though it uses your passport. As such, do not rely on your Digital ID when flying outside of the U.S. You'll still need your physical passport in order to validate your identity. In the future, however, Apple says you'll be able to use this Digital ID for other purposes, such as booking flights and hotels, as well as opening new accounts.

So we still have a ways to go. In California, where Leo and I are, we have digital driver's licenses. But as we've noted before, support remains spotty. So Jake's point, that a passport can provide Apple's Digital ID with a verified identity source means that those people who live in a state that does not yet support a digital driver's license, but who may have a valid passport, now have an alternative means for robustly identifying themselves to their iPhone.

Many pieces of any complete solution for online age verification remain missing – and those are big pieces. But we need to start somewhere, and I was encouraged by last Wednesday's pleasant surprise of Apple's Digital ID, since this is likely the foundation which will develop into more in the future. From a foundation like this, Apple will be able to generate secure privacy-preserving assertions such as "over 18" without revealing a single additional fact about a device's user. And given everything we know of Apple, there's no company whose motivations surrounding the preservation of their user's privacy I would trust more.

Checkout.com also says "No" to extortion – Donates requested ransom to security research
Last Wednesday, Mariano Albera the Chief Technology Officer at Checkout.com who was previously CTO at Expedia, OVO Energy, and Thomas Cook, posted his company's decision to say no to the ShinyHunters extortion gang. In his posting, headlined "Protecting our Merchants: Standing up to Extortion", Mariano wrote:

Last week, Checkout.com was contacted by a criminal group known as "ShinyHunters", who claimed to have obtained data connected to Checkout.com and demanded a ransom. Upon investigation, we determined that this data was obtained by gaining unauthorized access to a legacy third-party cloud file storage system, used in 2020 and previous years. We estimate that this would affect less than 25% of our current merchant base. The system was used for internal operational documents and merchant onboarding materials at that time.

This incident has not impacted our payment processing platform. The threat actors do not have, and never had, access to merchant funds or card numbers. The episode occurred when threat actors gained access to this third party legacy system which was not decommissioned properly. This was our mistake, and we take full responsibility. We are sorry. We regret that this incident has caused worry for our partners and people. We have begun the process to identify and contact those impacted and are working closely with law enforcement and the relevant regulators. We are fully committed to maintaining your trust.

We will not be extorted by criminals. We will not pay this ransom. Instead, we are turning this attack into an investment in security for our entire industry. We will be donating the requested ransom amount to Carnegie Mellon University and the University of Oxford Cyber Security Center to support their research in the fight against cybercrime.

Security, transparency and trust are the foundation of our industry. We will own our mistakes, protect our merchants, and invest in the fight against the criminal actors who threaten our digital economy. We are here to assist our merchants in whatever way we can. As always, we are available through your regular Checkout point of contact for any further assistance or questions you may have.

This is the way to handle a data breach if there's any way to do so. Mariano's donation is meant to have the effect of backfiring on the attackers. Not only will they not be paid, but the security researchers who work to track them down and take them down will be strengthened by receiving the money that Checkout.com refused to pay to the criminals. Nice going.

Google's Private AI Compute

Last year Apple launched their Private Cloud Compute and Google is now offering a similar solution. Under the banner: *"Private AI Compute: our next step in building private and helpful AI"* last Tuesday they said *"Today we're introducing Private AI Compute to bring you intelligent AI experiences with the power of Gemini models in the cloud, while keeping your data private to you."* Removing the glad-handing marketing-speak and excerpting the best technical bits from their announcement, they explained:

Today, we're taking the next step in building helpful experiences that keep users safe with Private AI Compute in the cloud, a new AI processing platform that combines our most capable Gemini models from the cloud with the same security and privacy assurances you expect from on-device processing. It's part of our ongoing commitment to deliver AI with safety and responsibility at the core.

AI is evolving to become even more helpful, personal and proactive. It's moving from completing simple requests to AI that can anticipate your needs with tailored suggestions or handle tasks for you at just the right moment. This progression in capability requires advanced reasoning and computational power that at times goes beyond what's possible with on-device processing.

About that, I suspect there's likely universal agreement. The thing that appeals to me about this is that AI requires short but massive bursts of computation, often followed by long periods of quiescence. Anyone who's been around since the days of the mainframe will recognize that this was the original brilliant concept that became known as "time sharing". Time sharing changed the world. The idea was that no one needed the full-time services of a massively expensive and capable mainframe. Instead, hundreds of people could use little time slices of that big machine's power. The result was massive efficiency.

The minicomputer largely encouraged a one-on-one usage mode though there were certainly many minicomputers running time sharing operating systems. Though it might have been more fair to call them mini-mainframes.

What really drove the nail into the time sharing coffin was the microcomputer, where the costs had come down so far that it no longer made any sense to share the machine. What developed was a truly personal computer.

The Internet, massive connectivity and massive data storage has begun to shift this model back toward the shared-massive-resources model with "Cloud Computing." It's pretty clear that Microsoft would be delighted to be servicing everyone out of their data centers. And that would be terrific right up until everyone suffers a massive service outage as Microsoft and all of their cloud-dependent users recently did.

But as I first noted, everything about the usage model of today's AI suggests that time sharing is back, and for exactly the reasons it was first explored in the early 1960's – massive resources used only briefly and intermittently by a great many people.

This leaves us with the question of security. Google wants us to believe that this can be every bit as secure as running on-device – meaning locally – on some poor overworked array that we're pouring ice water over while trying to figure out what to do with all the steam billowing out.

Everything we know tells us that it cannot be ***as*** secure. The security model is not identical. But everyone knows that. The question is, can it be secure enough? The problem with local compute is that to be fast enough, it needs to be super-powerful. And being cost effective while super-powerful means somehow keeping the darn thing super-busy. So the test for Google or Apple or any other cloud-based AI time-sharing farm is not whether it's ***as*** secure as local. It's not. The question is, it is secure enough. Here's what Google explains to convince us that theirs is:

> *We built Private AI Compute to unlock the full speed and power of Gemini cloud models for AI experiences, while ensuring your personal data stays private to you and is not accessible to anyone else, not even Google. Private AI Compute allows you to get faster, more helpful responses, making it easier to find what you need, get smart suggestions and take action.*
>
> *Private AI Compute is a secure, fortified space for processing your data that keeps your data isolated and private to you. It processes the same type of sensitive information you might expect to be processed on-device. Within its trusted boundary, your personal information, unique insights and how you use them are protected by an extra layer of security and privacy in addition to our existing AI safeguards.*
>
> *Private AI Compute is built on a multi-layered system that is designed from the ground up around core security and privacy principles:*
>
> - *One integrated Google tech stack: Private AI Compute runs on one seamless Google stack*

powered by our own custom Tensor Processing Units (TPUs). World-class privacy and security is integrated into this architecture with Titanium Intelligence Enclaves (TIE). This design enables Google AI features to use our most capable and intelligent Gemini models in the cloud, with our high standards for privacy and the same in-house computing infrastructure you already rely on for Gmail and Search.

- *No access: Remote attestation and encryption are used to connect your device to the hardware-secured sealed cloud environment, allowing Gemini models to securely process your data within a specialized, protected space. This ensures sensitive data processed by Private AI Compute remains accessible only to you and no one else, not even Google.*

So, does this make sense for your application? That's not for me to judge. What I can judge is that the concept of sharing massive AI Compute in the cloud makes ***all*** kinds of sense. That's what rings so true here. I would also note that doing this in a truly privacy-preserving fashion is not for fly-by-night outfits. Stick with brand names, here. If you're using Cloud AI to quickly compute gambling odds for near real time betting, then it might not matter who you use. But if you're using Cloud AI to form your publicly traded Fortune 500 company's ten-year product development plan, well ... frankly I'd have a difficult time letting that anywhere near the public Internet. But if you need to, use a security-first host.

Google stepped in it and is now scraping it off

Last Wednesday, Google tacitly acknowledged that they had been wildly overzealous with their pronouncement that all Android developers would henceforth be required to register using their real-world identities, and pay for the privilege, before they'd be able to publish their apps on the Google Play Store for Android devices.

Their update first reiterated the crucial importance of tightening down the security of Android's apps. On the subject of students and hobbyists, they wrote:

We heard from developers who were concerned about the barrier to entry when building apps intended only for a small group, like family or friends. We are using your input to shape a dedicated account type for students and hobbyists. This will allow you to distribute your creations to a limited number of devices without going through the full verification requirements.

That sounds like a terrific option to address the needs of those who are not looking or needing to reach a mass of users. But what about users of, for example F-Droid, who are advanced and security aware users? Google has made a carve-out for them, too, explaining:

While security is crucial, we've also heard from developers and power users who have a higher risk tolerance and want the ability to download unverified apps.

Based on this feedback and our ongoing conversations with the community, we are building a new advanced flow that allows experienced users to accept the risks of installing software that isn't verified. We are designing this flow specifically to resist coercion, ensuring that users are not tricked into bypassing these safety checks while under pressure from a scammer. It will also include clear warnings to ensure users fully understand the risks involved, but ultimately, it puts the choice in their hands. We are gathering early feedback on the design of this feature now and will share more details in the coming months.

Again, this sounds like the right approach, and it solves the F-Droid dilemma that we talked about previously.

But, so far, Google's primary goal of knowing and holding accountable the developers of apps for Android is not disappearing. It's onerous, but I don't see any other way to deal with this problem that has grown out of control. Something needs to change. Any developer who wishes to offer apps to the wider Android user base, and who doesn't want to subject their potential users to Google's deliberately terrifying advanced installation flow will need to register. Otherwise, for any app published by an unregistered developer, Google will be saying something like: "Please acknowledge that you understand that by downloading this Tiddly-Winks-In-Space game, you may be placing your life, and the lives of everyone you hold near and dear, at significant risk." Developer registration will prevent those notices, which would probably dramatically increase download counts. And at some point it's also conceivable that the Google Play Store may offer a filter to only show apps from "known" registered developers.

Windows 11 creates a Passkeys API that 1Password and Bitwarden support

This month's November update to Windows 11 added an API that allows third party password managers to deeply integrate with Windows 11. Under the heading *"Windows 11 expands passkey manager support"*, here's what Microsoft explained last week:

> *Windows is committed to making sign-in simpler, quicker, and more secure for every user. Today, we're excited to announce a major step forward in passwordless authentication: native support for passkey managers in Windows 11. This new capability empowers users to choose their favorite passkey manager — whether it's Microsoft Password Manager or trusted third-party providers. It's generally available with the Windows November 2025 security update.*
>
> *By partnering closely with third-party managers, we're delivering a more flexible, secure, and intuitive experience for Windows users everywhere, starting with 1Password and Bitwarden today and other passkey managers coming soon.*

I had to smile when I saw that the two top password managers we're affiliated with are the two that are enough on the ball to be participating with Microsoft on this out of the gate. Microsoft's announcement then quoted Travis Hogan, End User Group Product Manager for 1Password, saying:

> *"Working alongside the Windows Security team on the development of the passkey plugin API for Windows 11 has been a rewarding partnership. As the first password manager to offer native passkey support in Windows 11, we're proud to give customers a seamless passwordless experience inside and outside the browser. Together, we've ensured that 1Password and other third-party passkey providers can deliver a secure, standards-based experience natively on Windows, marking another major step towards a passwordless future."*

However, it appears that 1Password is actually tied for first with Bitwarden, since we also have:

> *"Bitwarden is delighted to collaborate with Microsoft on bringing native passkeys to Windows 11. This partnership enables more organizations and users to embrace passkeys confidently, knowing they can manage their credentials securely on Windows and across all their devices and platforms."*

Microsoft then asks the rhetorical question: *"Why plugin passkey managers?"* which they answer:

> *Passkeys are phish-resistant, less vulnerable to data breaches, and easier and faster to use than passwords. With plugin passkey manager support, you get:*
>
> - *Choice and flexibility: Use your preferred passkey manager natively on Windows.*
> - *Easy authentication: Create and sign in with passkeys using Windows Hello.*
> - *Passkeys everywhere: Your passkeys are synced between your Windows PCs and mobile devices. They go where you go.*
>
> *With plugin passkey manager support, packaged credential managers can integrate directly into Windows. Users can save, manage, and use passkeys across browsers and native apps — thanks to the new plugin provider capability. Setting up your credential manager is part of the passkey creation flow. Authentication uses Windows Hello — whether that is PIN, face, or fingerprint — so only you can access your credentials.*

And as for their own Microsoft Password Manager, they remind us of its benefits, writing:

> *We've integrated Microsoft Password Manager from Microsoft Edge natively into Windows as a plugin. That means you can use it in Microsoft Edge, other browsers, or any app that supports passkeys. This integration of Microsoft Password Manager from Microsoft Edge comes with added security benefits:*
>
> - *Passkey operations (creation, authentication, and management) are protected by Windows Hello.*
> - *Passkeys stored in Microsoft Password Manager will be synced and available on other Windows devices where the user is logged into Microsoft Edge with the same Microsoft account.*
> - *Syncing is protected by your Microsoft Password Manager PIN and a cloud enclave solution.*
> - *Azure Managed Hardware Security Modules (HSMs) help protect encryption keys.*
> - *Sensitive operations are performed inside a hardware-isolated environment in Azure Confidential Compute.*
> - *There is tamper-proof recovery with Azure Confidential Ledger.*
>
> *In other words, your passkeys are securely stored and easy to use.*

I'm sure Microsoft did not create this open API out of the goodness of their heart. They must have realized that there was no way they could force everyone to use their Windows-centric solution. The good news for those of us who are using 1Password or Bitwarden, is that those password managers, which are both comprehensively cross-platform, can now deeply and fully integrate into Windows 11, to serve as that operating system's native manager for its passkeys needs.

Last week, Microsoft, 1Password and Bitwarden published synchronized new releases about the appearance of this new Win11 support. Bitwarden wrote:

> *Now available in beta, the Bitwarden desktop application integrates with Windows 11 for an OS-native passkey experience. When creating or using a passkey in Windows, you can now select Bitwarden as the storage destination. Any passkey created and securely stored in the*

vault is synced to all of your devices providing you access from anywhere.

This works both ways, allowing for passkeys already saved in the Bitwarden vault to be used in Windows 11, for applications outside of the browser, and for the use of passkey website logins inside the browser - even without needing to have the Bitwarden extension installed. Simply select the Bitwarden desktop application when Windows prompts you to choose a passkey provider.

Bitwarden worked closely with Microsoft to develop the Windows component required for this functionality. In this beta release, the feature requires installing the desktop application from the Github repository. It will later be widely available through the standard desktop application install.

And Travis with 1Password wrote:

After 6 months in beta and working hard to address all your feedback, today's the day we finally bring desktop-level support for passkeys on Windows 11. No browser, no problem. You'll be able to seamlessly sync and manage passkeys on Windows, with 1Password as your credential manager. We're also introducing an improved onboarding flow to enable 1Password passkeys on Windows 11 to better meet you where you are.

This integration requires the MSIX version of 1Password for Windows. It uses the MSIX technology to better support all the functionality Windows 11 offers, including system-level passkeys. We've already begun the process of migrating nightly and beta users to the MSIX build, and we're starting to migrate those on stable today. If you'd like to get a jump start, you can download the latest version of 1Password for Windows.

To try out the new passkeys feature on Windows:

- *Ensure you are on the most up-to-date version of Windows 11.*
- *Download the latest version of 1Password for Windows here.*
- *Enable the passkey feature in your desktop app through the new onboarding prompt, or with Settings > Autofill and enabling the Show passkey suggestions setting.*
- *You should be redirected to enable 1Password as the system authenticator. If not, enable System Settings > Account > Passkeys > Advanced options. Then, enable 1Password using the toggle.*

As of today, the ability to use passkeys is available to all Windows 11 users. We'd again like to thank the Windows Security team for partnering with us so closely in order to get this out the door. Try it out and let us know what you think!

https://techcommunity.microsoft.com/blog/windows-itpro-blog/windows-11-expands-passkey-manager-support/4467572
https://www.1password.community/discussions/announcements/official-release-1password-now-supports-windows-11-passkeys/164379
https://bitwarden.com/blog/bitwarden-launches-passkey-management/

I've placed links to all three announcements in the show notes for anyone who wishes to follow up.

Newly appearing SIM cards in Russia are forced to "cool down"
I ran across an interesting piece of news that weirdly tied in with today's topic about cellular technology. Here's what the news report published:

> *The Russian Ministry of Digital Development, Communications and Mass Media announced on Monday that Russian authorities have begun blocking mobile phones being brought back into the country from abroad for 24 hours in an attempt to undermine Ukrainian drone strikes. The Ministry said that the measure had been applied in test mode on Monday, with mobile internet and SMS messages being blocked for 24 hours for anyone returning to Russia from abroad or for those who had not used their SIM card for 72 hours.*
>
> *While users should receive a notification informing them of the block via SMS, the ministry did say that internet access could be restored before the end of the 24-hour cooling-off period by completing a captcha sent by individual telecom operators.*
>
> *Investigative news outlet Verstka said that two of Russia's largest telecom operators, MegaFon and Beeline, were already warning their customers about the temporary suspension of their mobile data, but said that the Beeline link to restore access to data services did not appear to work.*
>
> *The Ministry said the measure had been introduced to avoid SIM cards being used to navigate Ukrainian drones. The cooling-off period was first reported by Russian business daily Kommersant on Friday, with some experts warning that technical glitches could mean that SMS notifications warning clients about the measure would not arrive.*
>
> *Last month, the Russian authorities began blocking foreign SIM cards from accessing data networks and texting services for 24 hours after entering the country to enable them to distinguish genuine foreign SIM cards from those being used to navigate Ukrainian attack drones, according to tech-specialist media outlet Rozetked.*

I thought that was a pretty clever idea. It's likely to have the tendency to false-positive somewhat, but Russian citizens will just need to put up with that inconvenience. And since most in-country SIM cards will be persistently connected to Russia's internal network, the idea would appear to be broadly workable. So the idea of not giving any newly appearing SIM card Internet or SMS messaging access for a period of time after it appears, is clever. What's not clear, though, is how this would prevent an enemy drone from using cellular towers for navigation. Navigation and communication would appear to be separate, and it was my impression from the reporting I've seen that Ukraine's drones were using Russia's cellular towers to determine their location. But perhaps there's also active command and control happening as well.

Google sues Lighthouse PhaaS (Phishing as a Service)
Google has filed a lawsuit against a Chinese Phishing-as-a-Service platform called "The Lighthouse." The Lighthouse is believed to be behind recent waves of SMS spam that targeted users across the world, posing as Google, USPS, and other services. The service has compromised over 1 million victims across 120 countries. Google is seeking a court order to shut down Lighthouse's infrastructure and is seeking injunctions against 25 individuals.

Every week or so, Lorrie shows me an SMS message and asks whether it's legitimate. But more than one million victims! Wow.

Listener Feedback

Anon the Moose

To Mr. Steve Gibson,

I needed to stop the most recent podcast and pen a reply because I must respond. I also must disclose that I am unable to speak for my employer, or anyone else including myself so please refrain from using my name.

I won't say I am old , but I have been messing with computers for only five decades or so. I don't remember when I first started listening, I know I can't claim number 1, but it was definitely in the low one or two hundreds. I am also a fan of the TWiT network and several of the past and present podcasts.

But your rant today about the demise of XSLT today is what made me respond. This is not the first time you talk about some format or another about how bad it is and that is and that how great it is to have it _finally_ going away. If I had a bingo card with the stuff I work on it would mostly be filled in.

Yes, I understand that interpreters are gaping security holes. I would also like to point out that the metrics the various companies are not accurate when you talk about things in closed environments. With respect to XLST I work on an international standards group that has tens of thousands of XSLT code that converts files of many different formats into a HTML deliverable that describes the exchange model that is a fairly complex graph. As a side note there is a significant amount of Leo's favorite language LISP in play too.

Ironically tomorrow morning I will be in a meeting where my proposal for a new technology stack will be questioned by folks that started working on this project before XSLT was invented... The refrain I always get is if it works why change because it will cost millions of dollars and several years to change...

In the past you spent several episodes dancing on the grave of flash. It turns out it can create a very good UI in a PDF file and when Adobe finally pulled the plug it killed a set of PDF applications that were used to maintain some expensive hardware.

Even further back you had a minor go at CGM where people were hiding malware in graphic files. I work in an industry that uses those files for illustrations and know that they are only uncommon on the web. Internally I need 9+ digits to count them especially when data retention and archiving is needed. I once wrote a tic-tac-toe game to demonstrate the things you could do with a graphics file.

If you want a crystal ball of what will be the next thing to be disallowed, I could look at the other things I have in my development directories to give you an idea of what the world will deprecate next. :) Keep the content coming it might be what is keeping me sane.

signed - Anon the moose

I think our anonymous moose intends to make the point, by way of grumbling about past work that occasionally needs to be tossed into the dust bin of history, that the world is changing and that development is a moving target. That has probably never been more clear than with the relatively sudden rise of AI coding agents. It seems clear that production coding tomorrow is not

going to look anything like production coding yesterday.

Matt

Steve, As you have stated, many listeners likely run their own mail server as you do, and have ventured into the world of SPF, DKIM, the extra hoops Google and Microsoft require and all of that work. As you know, having an email domain can have lots that needs done and a great tool I found in keeping mine running is sending an email from the domain you care about to this email address: check@dmarcly.com. Even just a blank email to that domain will result in an email that will return to where you emailed from containing all of this information:

- *Header from the Domain using RFC 5322*
- *DMARC: Pass/Fail*
- *DKIM: Pass/Fail, Alignment, domain and your selector record*
- *SPF: Pass/Fail, Alignment, Domain*
- *BIMI: I don't do this one so mine just says No Record found, unsure what it would show if you had it enabled.*
- *MTA-STS and TLS/RPT: Again I don't do these so mine just says record/policy not found.*
- *Blacklists: Checks your IP and mail server if it has hit any Blacklists*
- *Spam Score: Tells you what SpamAssain sees as the score of your email,*
- *mine is at a pleasant -0.1*

Darmcly themselves of course offers paid tiers of email support, but this email check service is completely free, and I have a weekly tasker set for myself to send it an email to just see how my email server is doing in the real world. I probably should automate that to a script that way the email just shows up once a week to me.

I also think this tool could be helpful for those who DON'T run their own server, just to see how the provider they are using is keeping their email deliverability something of a priority. As we all know, setting up an email server is dead simple. Getting emails to deliver from it, is a whole other matter. / Matt

I went over to DMARCLY.COM – be sure to spell it DMARC and not DMARK since that's a different email service (don't ask me how I determined that). It appears that they're wanting to collect business email accounts. I took Matt's suggestion and sent eMail to check@dmarcly.com but I didn't receive a reply. I figured that I might need to create an account, which Matt might not have known since he may have already had a free account. So I did that, created an alias address at GRC.COM and created a free account. I received the expected email confirmation "click here to confirm your email and account", then I sent another "check@dmarcly.com" email from that alias and that did the trick. I received a very nice and thorough analysis of GRC's SPF, DKIM and DMARC status. As well as GRC's BIMI email logo. That DMARCly.com site also has a bunch of very nice tests, advice and educational resources. So it looks like a very reasonable place to learn about DMARC and to test out one's email setup. Thanks, Matt.

But what matters to me most, is that Google now blames GRC for ZERO – ***absolutely ZERO*** – of the spoofed email that is continually flowing into them pretending to have been sent by GRC. Since I showed that flat-line at zero chart last week, we have another week of absolutely zero crap being blamed on GRC. With those SPF, DKIM and DMARC DNS records in place, no email sent from anywhere else will ever be seen as valid by any recipient who cares to check.

Scott Ulrich — Subject: *"Still getting Windows 10 Updates"*

Hey Steve, I made a point (on principle) of not doing anything Microsoft required to obtain Windows 10 updates past October: no storing settings in the Cloud, no payments, and I don't have enough Microsoft brownie points to get the extra year. I'm not in Europe, and I still seem to be getting Windows 10 updates; see attached. Curious if others are seeing the same. Cheers, Scott Ulrich

Scott attached a screenshot of his Windows update showing two seemingly contradictory things: his Windows 10 machine is reporting that it's receiving a November 2025 Cumulative Update for Windows 10 Version 22H2 for x64-based systems. And it notes that it's KB5071959. So this is clearly what Scott was referring to when he noted that his machine was still receiving updates.

But then below that, in the screenshot Scott thoughtfully provided, we see the familiar notice "Enroll in Extended Security Update" with the explanation "Your device is no longer receiving security updates. Enroll now to stay protected and productive for another year:

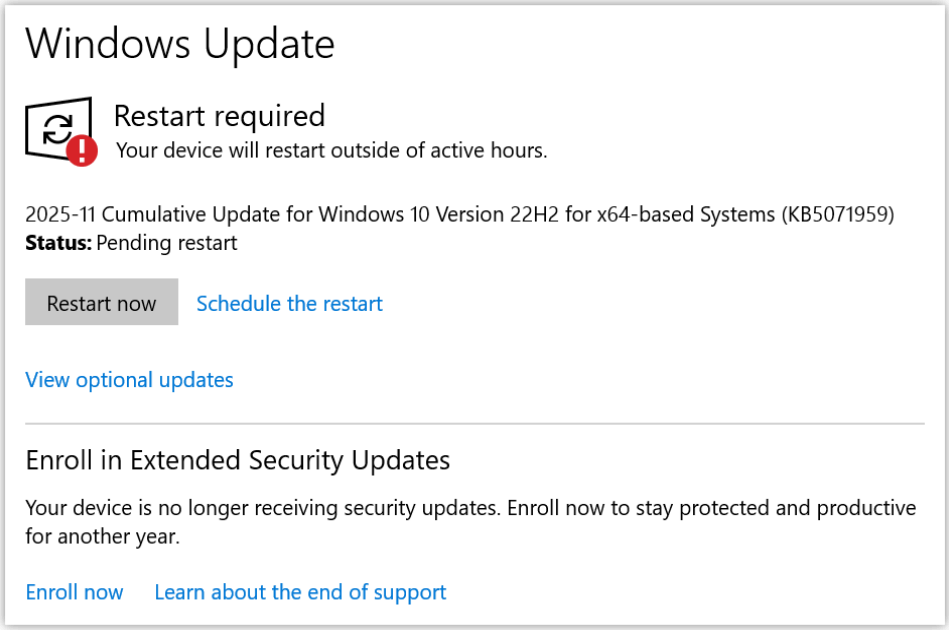

So what's going on here? The key is that specific Knowledge Base number KB5071959. It turns out, that's not what it might appear to be at first glance. Despite its November 2025 date, it is ***not*** providing November's security fixes for that machine. Instead, it's repairing a known ESU bug that has been preventing machines from successfully enrolling in Microsoft's ESU program.

Some of the reports of ESU failure are somewhat comical, since Windows will simply report: "Something went wrong", which is not very satisfying for someone who's panicked about keeping their Windows 10 system up to date. I found some terrific reporting of this over on "The Guru of 3D" site, where it's author wrote:

Windows 10 users sticking with the older operating system have one remaining lifeline for security updates: ESU, Microsoft's Extended Security Updates program. It's designed for systems that cannot move to Windows 11 or for users who simply prefer to stay on Windows 10 a little longer. ESU provides up to three more years of critical security patches, but you

need to be enrolled to receive them. Depending on the device, enrollment can be either paid or linked to Windows Backup, and it also requires a Microsoft account. The problem is that not everyone could enroll. Over the last few months, a mix of bugs made ESU activation unnecessarily difficult. Some users in the EU saw messages claiming the service was temporarily unavailable, even though the program was active. Others, trying to use the free activation method through Windows Backup ran into a generic "Something went wrong" message that stopped the process entirely. These issues appeared right as Windows 10 transitioned out of standard support, which created more confusion during a time when many users were already dealing with upgrade decisions.

There were also earlier cases where Windows 10 insisted the system had reached end-of-life even when ESU was active. The odd part was that this affected not only standard Windows 10 installations but also Enterprise LTSC 2021 and LTSC IoT 2021 — editions that still have years of official support ahead of them. Microsoft patched those cloud-based configuration errors earlier, but the enrollment bugs continued to cause trouble.

Microsoft has now addressed the remaining issues with the out-of-band update KB5071959. This patch fixes the EU enrollment failures and the sign-up errors tied to Windows Backup activation. If your device couldn't enroll in ESU before, this update is required to restore the system's ability to join the program. On the other hand, if ESU already works on your machine, the patch isn't mandatory. It mainly targets systems that were blocked by the earlier bugs.

With KB5071959 now available, all known ESU enrollment problems should be resolved. Windows 10 users who rely on extended support can finally complete the process without running into misleading warnings, regional availability errors, or dead-end messages. Nothing about ESU's requirements has changed, but at least the sign-up path is no longer impeded by these software faults.

If you're still running Windows 10 for the long haul, installing this update is worth doing before attempting ESU enrollment again. It ensures the last security-update window Microsoft offers for Windows 10 actually works as intended — especially important for anyone keeping older hardware in service.

So that's what's going on. Microsoft's page for this explained that this out-of-band, somewhat emergency update, not only fixed these well-known and persistent ESU problems, but it also included all of the security updates up through October 14th, when all non-ESU security updating ended.

I wanted to take the occasion of Scott's note to let everyone know what was going on. This fix, which became available last week will be automatically installed into all Windows 10 machines and should then resolve any remaining ESU enrollment problems.

Scott added:

P.S. There has been talk in recent weeks about going back and listening to previous episodes. I have some experience with this. I found your podcast in 2019 while I was studying for my CISSP during a career change toward a focus on security. I started listening weekly to all fresh episodes then went back to start listening from episode #1 while exercising and working on projects around the house.

It took several of the early episodes before I realized Leo was the same guy I used to watch in high school on Tech TV. The episode I started with was Episode 723 from July 16th, 2019 and I finally caught up with all prior episodes on October 28, 2023 ... so 4 years, 3 months, 12 days.

While I cannot credit you with achieving my CISSP at the time, I do thank you for keeping me interested in InfoSec and on top of the current topics ever since. Listening to all those old episodes was a great refresher on various IT topics and the evolution of security over the past 20 years. You helped to reinvigorate my career in technology and I've been happy to support you and Twit as a Club Twit member since 2021. Please keep up the great work.

That's a great backstory, Scott. Thanks for sharing that!

Larry Wilson

"Indeed, only about 0.02% of web page loads today actually use XSLT at all, with less than 0.001% using XSLT processing instructions." While I agree that these percentages indicate that XSLT is a small minority of web page loads, I have to imagine that the raw number of loads, per day, say, is actually tremendously large. Not to say that this changes the security concerns, but I don't interpret these numbers as saying that it sees little use. It seems to me that it says that it's being used, what, hundreds of thousands of times a day? Millions?

Larry's point is that even 0.02% of all web page loads, while representing a small fraction of the total, still represents a large actual value. And he's right, of course. And...

Lisa Lombardo

I hate to say it, but I am aware of enterprise product use of XSLT. Thanks for sharing this, so I can forward this news. Thank you /Lisa

As I noted last week, I suspected that within the reach of our listeners would be people who were actually using and still depending upon XSLT — or knew of others who were. And when I say "still" depending upon, that's sort of unfair, because it has been a universally supported standard from its original release, so there's no reason for anyone to not "still" be depending upon it ... except, of course, that everyone listening is now aware that some reengineering of those existing aging solutions will be required.

What's going to happen for those who are not listening to this podcast or who are not tapped into some source of similar information, is that come March of next year, there will be a rude awakening to the coming demise of XSLT when Google flips the "default ON" switch to "off". Suddenly, all of those facilities serving pages that are being displayed only thanks to the XML-to-HTML translation provided by those built-in browser features will fail. After some panicked scurrying around, everyone will figure out that the switch needs to be flipped back "ON", at which point those still using XSLT will have at most eight months to redesign their perfectly working system around more modern solutions.

Hearing firsthand from some of our listeners who will be directly touched by this, Google quite apologetic pronouncement is a bit more understandable.

John G. Ata

Looks like the Apple podcast subscription has doubled. I'm in the business of supporting GRC, not Apple. So I cancelled. Suggest you talk about this on your next episode.

I have no idea what's going on with Apple's podcast pricing, so perhaps Leo knows what's up?

David Lemire

Hi Steve, Your recent coverage of AI-related topics caused me to realize I have zero clue how an AI shopping agent works (full disclosure: I've yet to deliberately try out any AI tools). I found this brief article about a Columbia Business School study that offered some interesting insights. A paragraph that stood out:

One of the study's most striking conclusions is how differently AI models behave. Claude Sonnet 4, GPT-4.1, and Gemini 2.5 Flash frequently made divergent choices when asked to choose among identical assortments. For example, Claude favored one brand in the fitness watch category nearly twice as often as the other models. These preferences were consistent and measurable, suggesting that each AI model effectively creates its own miniature market with its own demand patterns.

As always, love your work. /David

Thanks, David. David's quote from Columbia and the surprise that might at first be felt, causes me to note that with AI we're no longer working with the sorts of "computers" that we always have. With a computer, we assume that there's one right answer. So we might at first be inclined to imagine that asking three different AI models to select the treadmill that offers the most value for the price, they ought to all converge to the same conclusion. But we know better. If we ask three different **people** that question we'll likely get three different answers. Today's AI models are individually hand crafted by their designers and the modeling data they train on – and the details of the way they train and are reinforced – may be similar but it's different in detail. And we know that even if they were all given identical training data, the differences in their internal design and operation would likely still cause them to reach different conclusions ... just like those three different people we might have asked.

Simon Zerafa

Hi Steve, For podcast listeners who are tech support for their non-techie friends and family, it is possible to disable the Windows Run dialog through Group Policy or the Registry as follows...

- *Navigate to HKEY_CURRENT_USER\Software\Microsoft\Windows\CurrentVersion\Policies\Explorer.*
- *If you don't have the Explorer key under "Policies", right-click on Policies and click New > Key and name that key Policies.*
- *Right-click in on the right side and click New > DWORD (32-bit Value) and name it NoRun.*
- *Double-click NoRun and change the Value Data to 1. Change the Value Data to 2 to enable the Run dialog or delete the NoRun registry key if you no longer need it.*

For users who don't need access to the Run dialog (many don't) this might be an effective solution to the problem of pasting "captcha" commands and unwittingly compromising their systems. Of course, this will disable the Run dialog for maintenance purposes to ensure you have access to the tools you might need via some other route!

Global Cellphone Tracking

We need to credit today's topic to a listener of ours named **Amir Katz** who wrote the following:

> *Hello Steve & Leo, Long time subscriber & SpinRite owner, etc.*
> *This is about a different type of phone hacking, so you may find this story very interesting:*
> *https://www.lighthousereports.com/methodology/surveillance-secrets-explainer/*
> *It was pointed out in Bruce Schneier's monthly newsletter, to which I'm sure you subscribe as well. /Thank you*

I do subscribe. But I subscribe to more than I can consume. And when I'm intensely focused upon coding I fall much further behind. So I didn't see Bruce's pointer to this, but I did see Amir's pointer to where Bruce was pointing.

As usual, I'm primarily driven by technology. That's what I find most interesting. And that appears to be the main reason our listeners keep listening and find this podcast worth their very precious time. So when I understood the enabling technology underlying this global cellphone tracking, I closed my eyes and shook my head – it was so insidious and obvious in retrospect. And I knew that everyone would feel the same way and would "get it" as I did.

So I've trimmed the original report where I could to keep its length under control, but it contains so much interesting detail that I'm sure everyone will find this fascinating. The piece's title is *"How First Wap Tracks Phones Around the World"* where *"First Wap"* is the name of a private company. And the article's teaser reads: *"From telecom protocols to a 1.5 million row dataset, here's how we uncovered the reach and tactics of a mercenary phone-tracking company."*

Before I share the edited-down version of their reporting, stop for a moment to ask yourself exactly how something that we all take entirely for granted works: How does the global telephone network "know" where everyone is all the time? Sure, we know that as we roam around, our handsets are pinging and logging into various nearby cell towers, and that relative signal strengths are compared to determine which cell tower base station should handle our connection. But what underlying protocol is used and who, exactly, has access to it? And more importantly, can anyone, anywhere, query the instantaneous location of anyone else?

We talked about that instance last June where, as a security precaution, senior Iran officials were deliberately ***not*** carrying mobile phones, because they were acutely aware of their trackability. But they failed to insist upon the same level of care from their bodyguards, who were carrying cellphones and who were, being bodyguards, in close proximity to their protectees. We assumed at the time that the bodyguards must have been practicing poor personal phone security hygiene and had their phones infected with some form of tracking spyware. But what if the reality is far worse? What if the underlying global cellphone network is so poorly designed and so insecure that anyone's location can be known at any time by anyone else, without the aid of any spyware, just by virtue of it being a cellular phone?

Here's what the team at Lighthouse Reports wrote:

> *In the spring of 2024, Lighthouse found a vast archive of data on the deep web. It contained thousands of phone numbers and hundreds of thousands of locations from nearly every country in the world.*

The data came from a little-known surveillance company called First Wap. Headquartered in Jakarta but run by a group of European executives, First Wap has quietly built a phone tracking empire spanning the globe. There have been leaks of telecom network targeting data in the past, but none of them has included this amount of successful targeting of individual phone numbers. The team found material inside the archive for dozens of stories, including how the company's tracking tech was used against Rwandan dissidents targeted in an assassination campaign, a journalist investigating corruption in the Vatican, and a businessman being investigated for compromising material. Unlike top-tier spyware firms, such as the notorious NSO Group, phone-tracking firms like First Wap have flown under the radar.

It's possible to view the surveillance industry as a pyramid. At the top are the elite spyware companies, selling expensive, highly targeted and invasive tools like NSO Group's Pegasus or Intellexa's Predator. At the bottom sit the preliminary tools that help enable surveillance operations: OSINT and social media scraping tools to develop profiles of targets, internet infrastructure to spin out lists of honeypot domains, and vulnerability vendors trading identified weaknesses in operating systems and other software. Sandwiched in between these is a middle layer — firms that track locations or intercept communications at scale, like First Wap.

With the top of the pyramid grabbing the most attention, the middle tier has managed to operate with less scrutiny, despite enabling surveillance on a far broader scale. A key player in this middle tier is First Wap, a little known phone-tracking firm headquartered in Jakarta.

First Wap's primary product is a surveillance tool called Altamides, an acronym for "Advanced Location Tracking and Deception System." While Altamides boasts a number of capabilities, its flagship feature is the ability to track a phone number anywhere in the world without leaving a trace on the device. Besides location tracking, Altamides also has the ability to intercept text messages and phone calls, spoof messages and even breach encrypted messaging apps like WhatsApp.

Okay, now, they rattled that off but that is the key: <quote> *"its flagship feature is the ability to track a phone number anywhere in the world without leaving a trace on the device."* This Altamides system leaves no trace because it's leveraging the fabric of the global telecom system to do all of the work. As I noted earlier, though it's obvious once you stop to consider it, the global cellphone network somehow always knows where every cellphone in the world is located. It must know this in order to function at all.

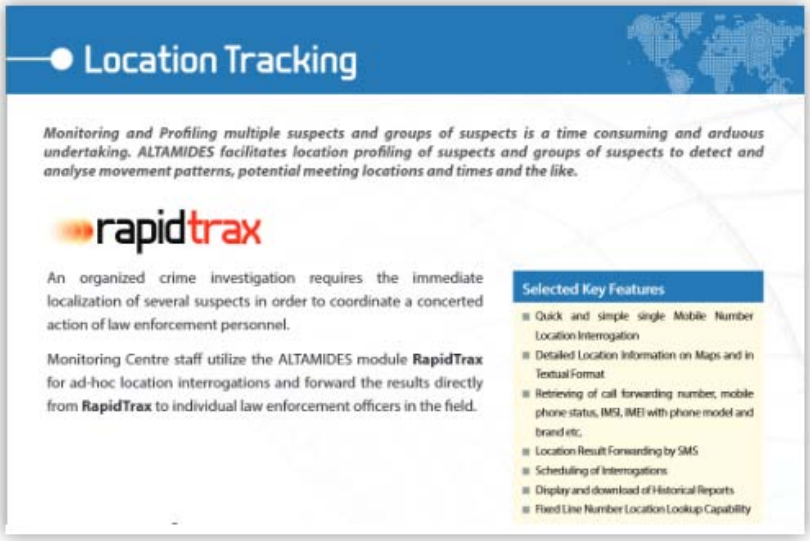

Their report included a snippet from the First Wap brochure describing Altamides. It says (above) Their report continues:

The investigation started with a 1.5 million row archive of surveillance operations carried out via First Wap's systems. Within the dozens of columns we found a relatively straightforward taxonomy of data – times and dates, latitude and longitude, phone numbers, country and phone operator names, map URLs – alongside fields that were at first glance less obvious, such as query methods, cell identifiers and other technical details. Numerous internal references in the dataset demonstrated its ties to First Wap and the Altamides tool.

What was clear was that this was a record of years of location tracking, targeting thousands of phone numbers in a vast range of countries. What was less clear, at first, was how to make sense of this mass of data.

On any given day the dataset might exhibit activity in dozens of places. On initial analysis, we saw that the majority of targets were tracked a small number of times, while a minority were tracked heavily or regularly. Similarly, while nearly every country in the world featured in the dataset, certain regions emerged as clear hotspots, either in terms of total volume of tracking, or in terms of number of devices tracked.

We wanted to understand who was being targeted. So we ran all of the more than 14,000 numbers through a combination of OSINT (Open Source Intelligence) tools which link phone numbers to internet accounts. We mapped the links between numbers and people using Maltego, and then connected this to the diachronic tracking data with an interactive user interface developed by team member Christo Buschek.

"Maltego" is a potent OSINT and link-analysis tool used to discover relationships between people, organizations, websites, domains, social media accounts, IP addresses, breaches, and many other entities. It's able to integrate with Shodan, VirusTotal, HaveIBeenPwned, Whois, social media lookups, public breach databases and many cybersecurity tools. In other words, it automates a lot of legwork. It's a commercial tool used by professionals, and there is a free rate-limited Community Edition.

Although this automated process surfaced thousands of potential matches between phone numbers and names, we only considered identifications to be valid if more than one datapoint connected the number to a person (beyond simply a matching name). A team of more than ten reporters at Lighthouse and paper train media spent months building up a high-confidence list of targeted individuals, which at time of publication included over 1,500 phone numbers.

Looking for outliers in the dataset led us to cases of harm and obvious misuse. Among the most heavily featured numbers we came across Anne Wojcicki (wo-JIT-skee)*, co-founder of 23andMe and at the time married to Google's Sergey Brin, who was tracked more than 1,000 times as she moved across the San Francisco Bay Area. We also detected cases where tracking was automated, with timestamps at the same minute of each hour, as was the case for Gianluigi Nuzzi* (john-loo-EE-jee NOOT-see)*, a well known Italian journalist who had uncovered a corruption scandal inside the Vatican.*

While we could see who was being tracked, we could not determine which Altamides user was carrying out the tracking.

Understanding the broader patterns of surveillance, and ultimately their motivation, required searching for "clusters" of targets — networks of people whose tracking was connected in time or space. A series of Nigerian election officials, for example, were all tracked in the city of Bauchi (BOW-chee) *ahead of Nigeria's 2011 election. In 2012, meanwhile, the wife of General Faustin Kayumba and the bodyguard of Patrick Karegeya, two founders of the Rwanda National Congress — an opposition movement operating in exile in South Africa — were tracked within minutes of one another. Both men had been targeted for assasination, with Karegeya found strangled in a Johannesburg hotel room 18 months after his bodyguard was targeted by Altamides.*

As we continued to identify phone numbers, we homed in on a portion of the dataset that indicated use in customer demonstrations. This data showed how First Wap's executives, or middlemen they had contracted to market their technology, tracked themselves and their associates so that potential clients could experience Altamides in action. In turn, these records allowed us to see the movements of First Wap's salesmen as they hopscotched the globe, interacting with potential customers – who themselves were sometimes exposed in the data, either by identity or location.

And now we come to the technology, answering the question what made all of this not only possible, but feasible and functional:

So how did First Wap connect the numbers in the dataset to locations? And why did some of the data contain blank locations or unsuccessful location attempts? In contrast to top-tier spyware like Pegasus, First Wap's Altamides does not infect a phone, it operates entirely at the level of the telecom network. First Wap's late founder, Josef Fuchs, realized before almost anyone that by exploiting an antiquated communication system he could trick phone networks into revealing the locations of their users.

Signalling System 7, or SS7, is a decades-old set of protocols that allows phone networks to communicate with one another, routing messages and calls across borders.

And here comes the phrase we have so often mentioned on this podcast: ***"It was never designed with security in mind"*** Right. Just like the Internet that came later, in the early days it was a miracle that all of this worked at all. Wikipedia tells us:

Signalling System No. 7 (SS7) is a set of telephony signaling protocols developed in the 1970s that is used to set up and teardown telephone calls on most parts of the global public switched telephone network (PSTN). The protocol also performs number translation, local number portability, prepaid billing, Short Message Service (SMS), and other services.

The protocol was introduced in the Bell System in the United States by the name Common Channel Interoffice Signaling in the 1970s for signaling between No. 4ESS switch and No. 4A crossbar toll offices. The SS7 protocol is defined for international use by the Q.700-series recommendations of 1988 by the ITU-T. Of the many national variants of the SS7 protocols, most are based on variants standardized by ANSI and the European Telecommunications Standards Institute (ETSI).

Then, Wikipedia adds, right on cue:

> SS7 has been shown to have several security vulnerabilities, allowing location tracking of callers, interception of voice data, intercept two-factor authentication keys, and possibly the delivery of spyware to phones.

In other words, First Wap has weaponized and commercialized the world's dependence upon the original insecure telephony system. What about improvements to security since then? The report continues:

> ***It was never designed with security in mind,*** *and while operators have moved to more secure evolutions with 4G and 5G, they still need to maintain backwards compatibility with SS7. This is likely to remain the case for years if not decades to come.*
>
> *Phone networks need to know where users are in order to route text messages and phone calls. Operators exchange signalling messages to request, and respond with, user location information. The existence of these signalling messages is not in itself a vulnerability. The issue is rather that networks process commands, such as location requests, from other networks, without being able to verify who is actually sending them and for what purpose.*
>
> *These signalling messages are never seen on a user's phone. They are sent and received by "Global Titles" (GTs), phone numbers that represent nodes in a network but are not assigned to subscribers. Surveillance companies have often leased GTs from phone operators and used them to send unauthorised signalling messages into other networks, benefitting from the fact that the signalling messages appear to be coming from the legitimate operator which owns the GT.*
>
> *First Wap primarily works via in-country installations of Altamides. In this setup, a government client uses Altamides via an SS7 link belonging to a local phone operator. The local phone operator provides the GTs and Altamides uses these GTs to conduct location tracking domestically and internationally.*
>
> *But the company also offered customers SS7 connectivity through Liechtenstein's national operator, Telecom Liechtenstein (formerly Mobilkom). The First Wap archive shows Altamides using GTs from Mobilkom to carry out hundreds of thousands of location tracking queries.*

Their report then digs into the details of the data they obtained, explaining the operation of the various commands that were issued to the global network. They address the question of abuse of the system, writing:

> *Over time, more phone operators have started to install firewalls to counter this type of threat. But maintaining them is complicated, and spotting this type of location tracking request within the millions of legitimate queries sent to an operator's subscribers on a daily basis is challenging. The more legitimate the source, the more likely it is that the operator on the receiving end of the query will let it through. Examination of the dataset shows that a considerable proportion of the activity in it was sent via Mobilkom Liechtenstein, which has excellent world wide links to other networks and, operating in the heart of Europe, also appears to be a trustworthy traffic source.*

In response to this investigation, Telecom Liechtenstein (formerly Mobilkom Liechtenstein) said it was unaware of any misuse of its network by First Wap. The phone operator said it had immediately "suspended its business relationship with First Wap" and that "If the allegations are substantiated, the collaboration will be terminated without notice, and the company reserves the right to take legal action."

First Wap stated in response to this investigation that it has "fully complied with the statutory and legal requirements and have also imposed this on our business partners". The company stated that it has "never attempted to hack an SS7 stack or similar" and has "not offered or sold our products and solutions to repressive systems or sanctioned countries or individuals."

As for the determination of location, they wrote:

The SS7 commands used do not themselves return longitude and latitude coordinates. Instead, they return a Cell ID, which is a unique number assigned to a cell in a mobile network and physically designating a tower or base station. A complete ID is made up of four parts: the country (MCC), network (MNC), area (LAC) and finally the cell.

Cell IDs can be mapped to a longitude and latitude using proprietary or public databases. Governments and operators will maintain their own lists, while there are also publicly available crowd sourced databases such as OpenCellID. When First Wap installed a system in a country, it requested the client to provide an up-to-date mapping of its domestic cell towers so that Altamides could convert Cell IDs to locations. But, as a brochure we obtained demonstrates, the company also offered to facilitate foreign Cell ID mapping for its customers, thus allowing them to carry out tracking operations abroad.

In the case of the First Wap sales representative cited in the data, the Cell ID was successfully mapped, and the phone was tracked right next to the headquarters of Nigeria's State Security Service (SSS).

The accuracy of such mapping depends on the density of cell towers in an area. In urban areas, such as Union Square in San Francisco, the high density of towers means that individual Cell IDs can be quite precisely located. In rural areas, there might be only one tower servicing a much larger area. So the accuracy of the map data depends on real-world physical context as well as technical issues of signalling queries.

Across the span of the archive, it is clear that First Wap's database of Cell IDs was still evolving. This meant that in many cases Altamides successfully obtained a Cell ID but was unable to map it into longitude and latitude. In these instances, the tool would either provide no coordinates, or would provide an estimated centre of a much larger area.

And they conclude their investigative report, writing:

Most countries have a legal mandate to carry out domestic phone network surveillance. The First Wap archive demonstrates, however, how phone network connections can be leveraged to allow tracking all over the world, without authorisation from the targeted networks.

In recent years, a number of investigations have explored the ways in which surveillance companies gain access to phone networks to enable this type of tracking. Lighthouse and its partners have previously written about how SS7 abuses were linked to the murder of a reporter in Mexico and a crackdown on an activist in Congo, and how they were enabled via leasing of Global Titles.

Although in recent years some phone operators have acted to improve security, global imbalances remain. Stopping unauthorised signalling is expensive and as such tends to be better implemented by wealthier operators. When operators in the global minority world allow surveillance companies to access their systems and send traffic, they are increasing risk to people in the global majority.

Last year, the UK's telecom regulator took the step of banning Global Title leasing on the grounds that phone operators had failed to manage risk and UK networks were carrying large amounts of malicious signalling traffic, including surveillance traffic, into networks worldwide. But for real change to occur, other regulators in the EU need to follow suit.

So today we've learned that without ever have any prior access to someone's cellular phone and without any necessity of installing any sort of spyware or malware, once someone's cellphone number is known, which can often be accomplished through some digging or a bit of skull-duggery, it is then possible to track their global movements with the granularity of cellphone towers.

In past podcasts we've seen how much damage this form of metadata can do even lacking any communication's content. For example, by simultaneously tracking multiple individuals who may be affiliated it would be possible to determine when and where they meet by monitoring the convergence of their locations.

As individuals, there's little we can do, though I doubt there's little that we should do. It's only very high profile people who probably have anything to worry about, though the wake-up call here is that no amount of cellphone hygiene will prevent this tracking. Nothing can prevent it. It's part of the fabric of the cellular radio-based system we all use today. And I suppose this reality does argue for the use of cheap "burner phones" by anyone who wishes to prevent any subsequent forensic analysis of their movements. We don't know how much logging of our locations is being done by the providers in our area for after-the-fact forensic data mining.

In any event, I wanted to make sure that everyone listening was at least aware that a malware infection is not a prerequisite to being tracked on the Internet and that there's nothing Apple or Google or Samsung or anyone else can do to prevent it. Rotating WiFi MAC addresses or using ephemeral MACs until associated with a WiFi access point won't help. Switching a phone to Airplane mode, or completely switching a phone's cellular radio off, so that it drops from the global cellular network, is the only way to disappear.