Security Now! #973 - 05-07-24

Not So Fast

This week on Security Now!

What danger is presented by the world's dependence upon GPS? And why is that of any concern? Has the sky fallen on all VPN systems? And why does the tech press appear to think so? Today's myriad network authentication options are confusing and incomplete. What does the future promise? Why might Apple have been erasing iCloud Keychain data? And what's actually going on between Google and the United Kingdom regarding the sunsetting of 3rd-party cookies? What's the problem? Or is there one?

Oh, don’t mind us... we’re just putting the lid back on the cactus…

Security News

The vulnerability of GPS

I wanted to start off this week by sharing an important piece of interesting news that's not Internet security related, but is nevertheless potentially quite a big and serious issue in the real world. Last Thursday's headline in WIRED was *"The Dangerous Rise of GPS Attacks"* with the subhead *"Thousands of planes and ships are facing GPS jamming and spoofing. Experts warn these attacks could potentially impact critical infrastructure, communication networks, and more."*

> *The disruption to GPS services started getting worse on Christmas Day. Planes and ships moving around southern Sweden and Poland lost connectivity as their radio signals were interfered with. Since then, the region around the Baltic Sea—including neighboring Germany, Finland, Estonia, Latvia, and Lithuania—has faced persistent attacks against GPS systems.*
>
> *Tens of thousands of planes flying in the region have reported problems with their navigation systems in recent months amid widespread jamming attacks, which can make GPS inoperable. As the attacks have grown, Russia has increasingly been blamed, with open source researchers tracking the source to Russian regions such as Kaliningrad. In one instance, signals were disrupted for 47 hours continuously. On Monday, marking one of the most serious incidents yet, airline Finnair canceled its flights to Tartu, Estonia, for a month, after GPS interference forced two of its planes to abort landings at the airport and turn around.*
>
> *The jamming in the Baltic region, which was first spotted in early 2022, is just the tip of the iceberg. In recent years, there has been a rapid uptick in attacks against GPS signals and wider satellite navigation systems, known as GNSS, including those of Europe, China, and Russia. The attacks can jam signals, essentially forcing them offline, or spoof the signals, making aircraft and ships appear at false locations on maps. Beyond the Baltics, war zone areas around Ukraine and the Middle East have also seen sharp rises in GPS disruptions, including signal blocking meant to disrupt airborne attacks.*
>
> *Now, governments, telecom and airline safety experts are increasingly sounding the alarm about the disruptions and the potential for major disasters. Foreign ministers in Estonia, Latvia, and Lithuania have all blamed Russia for GPS issues in the Baltics this week and said the threat should be taken seriously.*
>
> *Jimmie Adamsson, the chief of public affairs for the Swedish Navy, told WIRED: "It cannot be ruled out that this jamming is a form of hybrid warfare with the aim of creating uncertainty and unrest. Of course, there are concerns, mostly for civilian shipping and aviation, that an accident will occur creating an environmental disaster. There is also a risk that ships and aircraft will suspend traffic to this area and thereby affecting global trade."*
>
> *Joe Wagner, a spokesperson from Germany's Federal Office for Information Security, toldW WIRED: "A growing threat situation must be expected in connection with GPS jamming." Wagner said there are technical ways to reduce its impact. Officials in Finland say they have also seen an increase in airline disruptions in and around the country. And a spokesperson for the International Telecommunication Union, a United Nations agency, told WIRED that the number of jamming and spoofing incidents have "increased significantly" over the past four years, and interfering with radio signals is prohibited under the ITU's rules.*

Attacks against GPS, and the wider GNSS category, come in two forms. First, GPS jamming overwhelms the radio signals that make up GPS and make the systems unusable. Second, spoofing attacks can replace the original signal with a new location—spoofed ships can, for example, appear on maps as if they're at inland airports.

Both types of interference have increased in frequency. The disruptions—at least at this stage—mostly impact planes flying at high altitudes and ships that can be in open water, not people's individual phones or other systems that rely on GPS.

Within the Baltic region, 46,000 aircraft showed potential signs of jamming between August 2023 and March this year, according to reports and data from tracking service GPSJam. Benoit Figuet, an academic at the Zurich University of Applied Sciences who also runs a live GPS spoofing map, says there have been an additional 44,000 spoofing incidents logged since the start of this year.

Earlier this month, more than 15,000 planes had their locations spoofed to Beirut Airport, according to data Figuet shared with WIRED. More than 10,000 were spoofed to Cairo Airport, while more than 2,000 had their locations showing in Yaroslavl in Russia, the data shows. Separate analysis from geospatial intelligence company Geollect shared with WIRED shows that on April 16, around 55 ships broadcast their location as being over the main runway at Simferopol International Airport in Crimea, Ukraine. The airport is around 19 miles inland from the Black Sea, where it's believed the ships were actually located.

Zach Clements, a graduate research assistant at the University of Texas at Austin said: "The biggest change in the past six months is definitely the amount of spoofing that's going on. For the first time, we're seeing widespread disruptions in civil aviation, especially in the Eastern Mediterranean, the Baltics, and the Middle East. In prior years, there were reports of spoofing impacting marine vessels, but not aviation."

Clements says there appear to be three spoofers that can be traced back to Russia. One open source intelligence analyst, going by the pseudonym Markus Jonsson, has located jamming in the Baltics, and that which impacted the Finnish airline this week, to Kaliningrad and other Russian locations. One research group has suggested disruption near Poland impacted Russia's own GNSS system less than others. Russia has a long history of interfering with GPS signals both within its borders and internationally. Russia's embassy in the UK did not respond to a request for comment.

The disruptions can cause uncertainty and potential safety issues for airline pilots and their passengers. A spokesperson for Eurocontrol, a European aviation organization with more than 40 countries as members, says its analysis shows disruptions are happening in the Eastern Mediterranean, areas around Ukraine and the Black Sea, as well as the Baltic states. During one week in March, 4,387 aircraft reported issues, the Eurocontrol spokesperson says, while for the same week last year, there were 2,646 flights reporting problems.

The Eurocontrol spokesperson says planes can fly safely without GNSS, but interference "puts a higher workload on pilots and air traffic control." A safety notice issued by the UK's Civil Aviation Authority this month says loss of GNSS can result in navigation issues, incorrect emergency "terrain" warnings that the plane is low to the ground, and the failure of various other systems.

In a NASA report detailing GPS incidents that was also published this month, one pilot said: "I have flown with crew members who were not fully aware of this problem." Other pilots said

> *they had received "false terrain warnings" that caused them to pull up and that pilots should have a "thorough review of jamming affects on the different aircraft systems" as part of their training.*
>
> *Jari Pöntinen, a director at Traficom, the Finnish transport and communications agency, says there has been an increase in disruptions both close to and in Finland since the beginning of this year. Pöntinen, who was formerly a pilot, says flight operators need to make sure they have made "comprehensive risk assessments" and to properly train pilots about what to expect if GPS disruptions occur.*
>
> *The increase in GPS disruptions has partly coincided with Russia's full-scale war in Ukraine and Israel's attacks in Gaza. Disrupting GPS as part of electronic warfare has become common on Russia and Ukraine's battlefields as a way to try to limit the operation of drones. And while Iran launched a barrage of missiles and drones against Israel on April 13, Israeli GPS disruption designed to limit the impact of the attack also impacted mapping and taxi services as well as food delivery.*
>
> *Kevin Heneka, the founder of cybersecurity company Hensec whose work includes detecting GPS disruptions, says jamming and spoofing technology has become cheaper and smaller over the years, to the extent that individuals can install them in their cars to hide their movements. However, Heneka says, more sophisticated attacks use equipment that can cost huge sums. "In conflict zones, in military terms, and in professional terms, this spoofing is very sophisticated, and it always goes hand in hand with jamming."*
>
> *As the number of jamming and spoofing incidents increases, there is a growing concern that the disruption of crucial services could become normalized. Multiple experts worry that the full extent of GPS interference is not known, and systems beyond those of airlines and shipping may be at risk if the disruption becomes more widespread.*
>
> *Maksim Barodzka, the CEO of GNSS-detection firm GPSPatron said: "Many do not realize that GNSS is not only used in your mobile phone for navigation, but is also a primary source of time synchronization for vital infrastructure: power grid systems, data centers, automatic train control systems, communication systems—especially 5G—financial services, and any distributed management and control systems. What's happening with this infrastructure is not widely reported in the public domain."*

Since both the jamming and the location spoofing disruptions are enabled through the use of very powerful local radio transmitters, which overwhelm the reception of the authentic signals being beamed down from the GPS satellite systems in orbit, so long as you're not in the region of the Baltics, where it appears Russia has taken to deliberately creating major disruptions, the attacks, being local, will not be a felt.

But the problem for those who **are** in the region is that GPS and the wider GNSS (which stands for Global Navigation Satellite System) have always been incredibly reliable resources. And as we know, when something is very useful and earns the reputation for also being very reliable, the result is a strong dependence upon that resource. So many modern, **non**-military, commercial systems have become so reliant upon GPS that deliberate disruptions for military purposes, such as Russia is likely perpetrating, can cause dramatic collateral damage.

The GPS system was conceived a little over 50 years ago, in 1973. Five years later the satellites began launching and encircling the Earth and today's 24-satellite system has been fully operational since 1993.

And talk about depending upon something that's more fragile than we might want. Our phones and automobiles only know where they are today largely thanks to GPS signals from space. We've talked about the militarization of space and the idea that having satellites attacking one another "up there" is no longer the territory of James Bond science fiction.

As global political tensions increase, we can hope that no major powers having space-based military capabilities, nor the ability to kill satellites from the ground, believe that denying the entire world these benefits would create an advantage for them.

Before GPS, the only way for something to know where it was, was through a system of inertial navigation. Inertial navigation is a closed system which relies upon the system's precise measurement of its own linear and angular accelerations. It integrates those over time to determine its velocities which are again integrated over time to determine its position. Even though inertial navigation systems are still in use due to the nearly instantaneous position, and especially angular, feedback they provide, the errors that tend to creep in over time can only be eliminated with the use of slower but far more accurate input from the global GPS system.

I suspect Russia's primary concern is with the use of autonomous military drones, which may rely upon GPS to determine their in-flight location. But since the risks presented by GPS jamming have been so well known for quite some time, I suspect that the latest technologies are much more immune to GPS outages than those in Russia might wish. Given all of the advances made in vision and its real-time recognition, I would be surprised if the latest autonomous technologies were not able to fly nearly as well by sight as they can by GPS. They might well use GPS as a first choice, but use vision to detect location spoofing while also being able to switch to pure vision if GPS should fail completely. Another likely strategy, since GPS signals will always be originating from above, would be to shield any GPS receiver and its antennas from ground-based interference. This can probably work for anything flying, but might not be practical for ships at sea.

Is the sky falling on all VPN systems?

Yesterday, Ars Technica got a little carried away in their reporting of a clever hack that a Seattle, Washington based penetration testing firm, the Leviathan Security Group, posted in their blog. The blog posting carried the headline "How Attackers Can Decloak Routing-Based VPNs For a Total VPN Leak" and, curiously, they've assigned a CVE number to their discovery even though nothing about this is a bug or a flaw. It's just a clever local exploit of a little-used feature of DHCP servers. Unfortunately, Ars Technica's headline for their story was "Novel attack against virtually all VPN apps neuters their entire purpose" which makes this sound more like the end of VPNs as we've known them. It isn't. Here's what's going on...

Our PCs all interact with both internal and external networks through network interfaces. Most systems typically have a single physical network interface, or NIC, but it's possible for a machine to have more than one physical network interface with each interface connected to different physical networks. In that case, it’s important for outgoing network packet traffic to know which physical interface any given packet should be routed out through. To answer that question, our machines contain a routing table. The routing table performs a “most specific match” function, based upon the destination IP address, to decide which interface should receive each packet being sent. Under Windows, opening a command prompt and entering the command “route print” will display a list of the system’s interfaces followed by the IPv4 and IPv6 routing tables, respectively.

This sort of network communication comes in so handy that in addition to true physical interfaces, many of our machines will have one or more virtual network interfaces. For example, the use of virtual machines has become very popular and they create virtual network interfaces to talk to their host machine and to the outside world.

And, here’s the main point: Many VPNs, like OpenVPN for example, operate by creating their own virtual interface in the hosting machine. It looks like and operates like any other network interface. But being a VPN – a Virtual Private Network – which is used to transact privately with encryption, any packets sent out of that virtual interface are first encrypted then re-routed out of an actual physical interface to be sent to the VPN’s matching endpoint. Since the typical VPN user, while using a VPN, wants all of their machine’s traffic to be “tunneled” through the VPN, when the VPN tunnel is brought up, the VPN software dynamically edits the system’s global routing table in such a way that instead of the system’s traffic being routed out through its normal actual physical interface, all of its traffic is, instead, routed to the VPN’s software-created virtual network interface. This is the way that, deep down inside the guts of our machines, all of the traffic that’s normally unencrypted suddenly becomes encrypted when we activate our VPN.

Essentially, the routing table is used to divert all outbound traffic to the VPN’s virtual network interface where it will be encrypted on the way out and decrypted on the way back in.

We need one other piece of information just to be certain that everyone is on the same page. DHCP stands for Dynamic Host Configuration Protocol. By default, when any networked machine boots up and gets itself going, it needs to be using an IP address for itself on its local network that’s unique for that network and it needs to know the IP address to which it should address packets bound for the outside world – in other words, the network’s gateway IP. It may also want to know the IP addresses of some DNS servers that will honor its requests. It’s the network’s inward facing DHCP server that answers all of these needs. When any networked machine starts up, by default, it will emit a broadcast packet onto the network announcing its presence and asking for any listening DHCP server to please provide it with all the information it requires to become a well-behaving citizen on the local network and to connect to the rest of the global Internet.

DHCP cleanly organizes the various types of information it can supply to requesting clients by number, each known as an option where the option number is a single byte, thus having a value from 0 to 255. The options are provided as a list of information terminated by option 255, which is a byte of all 1’s. So...

Option 1 provides the network’s subnet mask to the requesting client. Option 2 specifies the offset of the client's subnet in seconds from Coordinated Universal Time (UTC). Option 3 specifies a list of the IP addresses of routers on the client's subnet – what we know as the Gateway IP. Option 4 specifies a list of time servers available to the client. Option 6 provides a list of DNS servers for the client's use. There are some surprises among the list. For example, options 69 and 70 provide the IP addresses of SMTP and POP3 eMail servers. We're all used to specifying those ourselves, but back in 1997, that was information that DHCP was able to supply.

Something else that DHCP was able to provide is the source of today's trouble. The RFC's definition for Option 33 defines it as the "Static Route Option" and says: *"This option specifies a list of static routes that the client should install in its routing cache. If multiple routes to the same destination are specified, they are listed in descending order of priority. The routes consist of a list of IP address pairs. The first address is the destination address, and the second address is the router for the destination."* Now, if some of you just said "Oh, crap!" that would be the correct reaction. What this means is that the response from a DHCP server can be used to mess with a machine's routing table. And, as we noted earlier, a machine's traffic is routed to the VPN's virtual interface traffic through a dynamic modification of the machine's routing table.

Now, as it happens, Option 33 is not really the problem because, it was defined back in 1997 when IP networks were all class A, B, or C. That meant that networks were defined to always have exactly one, two or three bytes of host machine addresses. This was extremely wasteful of IP addresses for networks falling into intermediate sizes. So CIDR – C.I.D.R. – Classless Inter Domain Routing was adopted, which is what we have today, where the network mask can have any number of contiguous bits set. This change obsoleted Option 33, forcing its replacement five years later in 2002 by RFC 3442 which introduced option 121 which allows for the specification of classless static routes.

I mentioned that I was surprised that these Leviathan Security Group guys had arranged to get a CVE assigned for this, since technically this is a feature, not a bug. And all the way back in 1997 the fundamental vulnerability of DHCP was well understood. Section 7 of RFC 2131 dated March 1997 is titled "Security Considerations" and it reads:

> *DHCP is built directly on UDP and IP which are as yet inherently insecure. Furthermore, DHCP is generally intended to make maintenance of remote and/or diskless hosts easier. While perhaps not impossible, configuring such hosts with passwords or keys may be difficult and inconvenient.* ***Therefore, DHCP in its current form is quite insecure.***
>
> *Unauthorized DHCP servers may be easily set up. Such servers can then send false and potentially disruptive information to clients such as incorrect or duplicate IP addresses,* ***incorrect routing information*** *(including spoof routers, etc.), incorrect domain nameserver addresses (such as spoof nameservers), and so on. Clearly, once this seed information is in place, an attacker can further compromise affected systems.*

So, here's how the Leviathan guys describe the attack they've devised by abusing Option 121:

> *Our technique is to run a DHCP server on the same network as a targeted VPN user and to also set our DHCP configuration to use itself as a gateway. When the traffic hits our gateway, we use traffic forwarding rules on the DHCP server to pass traffic through to a legitimate gateway while we snoop on it.*

> *We use DHCP option 121 to set a route on the VPN user’s routing table. The route we set is arbitrary and we can also set multiple routes if needed. By pushing routes that are more specific than a /0 CIDR range that most VPNs use, we can make routing rules that have a higher priority than the routes for the virtual interface the VPN creates. We can set multiple /1 routes to recreate the 0.0.0.0/0 all traffic rule set by most VPNs.*
>
> *Pushing a route also means that the network traffic will be sent over the same interface as the DHCP server instead of the virtual network interface. This is intended functionality that isn’t clearly stated in the RFC. Therefore, for the routes we push, it is never encrypted by the VPN’s virtual interface but instead transmitted by the network interface that is talking to the DHCP server. As an attacker, we can select which IP addresses go over the tunnel and which addresses go over the network interface talking to our DHCP server.*
>
> *We now have traffic being transmitted outside the VPN’s encrypted tunnel. This technique can also be used against an already established VPN connection once the VPN user’s host needs to renew a lease from our DHCP server. We can artificially create that scenario by setting a short lease time in the DHCP lease, so the user updates their routing table more frequently. In addition, the VPN control channel is still intact because it already uses the physical interface for its communication. In our testing, the VPN always continued to report as connected, and the kill switch was never engaged to drop our VPN connection.*

So then they raise the question that I’ve had all along, by asking “Is TunnelVision a vulnerability?” and I appreciated their answer. They wrote:

> *This is debatable. We’re calling it a technique because TunnelVision doesn’t rely on violating any security properties of the underlying technologies. From our perspective, TunnelVision is how DHCP, routing tables, and VPNs are intended to work.*
>
> *However, it contradicts VPN providers’ assurances that are commonly referenced in marketing materials; in our opinion, TunnelVision becomes a vulnerability when a VPN provider makes assurances that their product secures a customer from an attacker on an untrusted network. There’s a big difference between protecting your data in transit and protecting against all LAN attacks. VPNs were not designed to mitigate LAN attacks on the physical network and to promise otherwise is dangerous.*
>
> *In our technique, we have not broken the VPN’s cryptographically secured protocol, and the VPN is still fully functional. An attacker is instead forcing a target user to not use their VPN tunnel. Regardless of whether we classify this as a technique, VPN users are affected when they rely on assurances that a VPN can secure them from attackers on their local network.*

And as for what systems are affected, the short version is everything except Android. They wrote:

> *In our testing, we observed that any operating system that implements a DHCP client according to its RFC specification and has support for DHCP option 121 routes is affected. This includes Windows, Linux, iOS, and MacOS. Notably, it does not affect Android as they do not have support for DHCP option 121.*

So just to be clear about the scope of the danger presented by the potential abuse of DHCP's Option 121, this is strictly a local LAN-side attack. It requires that an attacker arranges to setup and operate a malicious DHCP server on the same LAN network as anyone whose VPN will be intercepted. And the attacker needs some means of defeating the network's actual DHCP server. Since DHCP client's will accept the first reply to their query, being faster to reply is typically all that's needed.

So it is definitely conceivable that larger enterprise environments that may be depending upon VPN security for mission critical tasks could be targeted. And it also turns out that Option 121 is not the least bit obscure in the enterprise. It is heavily used.

A posting over on StackExchange says:

> *I'm running OpenVPN on a CentOS 7 server. The DHCP server on the LAN uses Option 121 to tell other devices to use this CentOS server if they want to get to the VPN subnets the OpenVPN server is connected to. This works great.*
>
> *The problem is that this CentOS server is getting these same routes from the DHCP server, which breaks things. If I manually remove these static routes from its routing table, everything works.*

And just last Tuesday, someone posted to the embarrassingly useless Microsoft answers forum:

> *When connected to my office network, its DHCP server will use option 121 to assign three different networks to be reached using a router which is not the default gateway. This works absolutely, the networks appear in my routing table, in active routes. Everything works, networks are reachable.*

He goes on at some length to explain that things don't work when he boots his PC without any network connectivity. Right. It's unclear what success this person was hoping to have without any network, but that's beside the point. I wanted to point out that this DHCP option comes in very handy and is apparently in quite heavy use within more complex corporate networks. This means that simply blocking or disabling it may not be feasible. This comes down to a useful and important feature that's in heavy use within the enterprise that is clearly subject to abuse.

Although it would be extremely unlikely for anyone at home to ever have anything to worry about, it's still instructive to paint the picture. The way I can see this might occur would be if some malicious device were connected to a residential network and wanted to capture all of the user's traffic, whether tunneled through a VPN or not. By being the first device to respond to any DHCP query, such a malicious device could first establish itself as the network's gateway to receive, inspect, and forward any traffic from the network's machines. And then by additionally using option 121, such a device could use that to insert entries into the user's routing table to prevent their VPN from tunneling the user's traffic, even though the VPN would show that everything was working and the user's traffic was protected. The VPN tunnel would be up and established, but it would not be carrying any of the user's traffic.

Since there are many environments where Option 121 is not needed and never used, it would be nice for our operating systems to provide the option to disable it permanently.

Closing The Loop

Dave Brenton / @Un_Woke_DB

> *Mr. Gibson. Quickly may I say as A Machine Language Coder - I admire your work in that area! I am a Spinrite owner / user, and longtime fan (since near the beginning) of SecurityNow! My question is about security keys. I hope this is not too long a question:*
>
> *I am about to make the transition to YubiKey and so I intend to purchase 2 to have a safe-fallback in case of loss. I am also planning to convert the wife over to the PassKey world. My question is: can the passkeys be paired across two user accounts - thereby ensuring recovery in case of loss with only 3 keys? - My mental model said it made sense but I do not know for sure ...*
> *1) Can the same key be applied to two different people?*
> *2) To assure full backup protection can all three keys be coded into both users?*
>
> *It may be a silly notion - but could it work or should I Just buy 4 keys to begin with? Thank you for all your good work and propeller-head installments! On to 999 and beyond! Dave*

I chose to share Dave's question because it so perfectly demonstrates the near total mess the user authentication world has fallen into today. I'm hopeful this may just be a transition phase. But truth be told, all of our collective experience also leaves me feeling somewhat skeptical. I worry that all we've done by having the FIDO group lower the bar for entry from requiring physical key dongles to allowing pretty much anything else – smartphones and PCs running simple software passkey clients – is to expand upon the number of available options, with an additional and, difficult as it is to believe in this day and age, not well thought out system. And we've added this new and not well thought out system, without removing any of the previous options. Have traditional username and passwords been replaced? No. Are they ever going to? Not in this lifetime. Have the "I forgot my password" links gone away? No. Are they ever going to? No. What about those time-based one-time passcodes? Are they going away? No. Any plan for that? No. What about OAuth which brings us the "login with your Google or Facebook or some other account?" Have those been obsoleted and removed? Nope. Can they be? Not easily, since many such sites only know their users thanks to their redirection through another web service's authentication.

And so, to this pile of existing half-baked remote network identity authentication solutions we're now adding Passkeys. A mysterious new solution that its designers all say is amazing and far more secure, and which works sort of like magic right up until it doesn't work at all. And when that happens, what do we do? We fall back to "send me an eMail".

What we have wound up with is the well-known and often observed phenomenon of "solution spread". We invent a better idea than what we had before. Perhaps it's because the times have changed and the older solutions are no longer adequate. Or perhaps we have more technology and available processing power than we had before, so new solutions are available than were previously. But the problem is, we rarely are able to kill off the things that came before. Why? Because, by the time we can do something more, too many people have come to depend upon the previous solution... and the one before that... and the one before it.

And this solution-spread doesn't just apply to the authentication domain. Just look at Windows.

Without getting bogged down in the details, every few years Microsoft comes up with a new and much improved way of writing applications for their Windows OS. And they promote the hell out of it, explaining how and why it's so much better than everything that came before. And do they then kill off the previous ways of programming Windows? No. Of course not. They can't. They were once promoting the hell out of each of ***those*** previous solutions and they got lots of people onboard using them then. So even though they no longer love them, and are urging everyone to use the new system, that never happens. I've heard Paul over on Windows Weekly saying that the original Windows API, Win32, should have died off long ago. That's what all of my Windows apps are written in; and not just mine; a gazillion others as well. (And that's "gazillion" with a 'G'.) I'm certain Paul knows that Microsoft will never abandon Win32. They can't, any more than websites will ever be able to stop offering username and passwords with an "I forgot how" eMail link.

So just to be clear, the industry has added a bright and shiny ***additional*** way for people to login to their accounts. But none of the existing ways are, or will be, removed. Remember that today in 2024, only one out of every three Internet users is using any password manager. I really don't know what the rest are doing. Perhaps these are the people whose iOS and Android support for passkeys is mostly aimed at. They don't know, understand or care about their online identity. So when Apple or Google comes along and asks "how would you like to logon instantly with passkeys and never worry about another password?" that sounds great.

But that's not Dave, our listener whose questions launched me into first taking a wider view of where we stand today. So let's look at Dave's situation. Dave says he's planning to convert his wife over to Passkeys. I'm sure he means that he would like to have his wife begin to use passkeys, since it's not possible to "convert over" to passkeys in any meaningful way when so few websites offer the option. The caution there, since we do not yet have passkey transportability, is to be careful about which app is holding a site's passkeys. As I mentioned last week, iOS, Windows, Android and now an increasing number of traditional password managers will all be vying to be **the** app that generates the passkey to be provided to a website. Since only **that** app will then be able to authenticate the user to that site with a passkey, the only sound strategy will be to only and always use a single platform for passkeys.

This issue, and Dave's other questions, require a quick bit of foundation about the operation of passkeys. When an application prompts its user about whether the user wishes to have it create a passkey, that's exactly what's happening. That application generates a cryptographically strong secret and private key – which never leaves the application and which the application guards carefully. From that closely held private key it then generates a public key, and only the public key is sent to and retained by the website. In the future, that website will use the public key it holds to verify the signature of a challenge that it sends to the user's passkey authenticator.

So my point here is that, today, there is no provision for these private keys, which were generated internally and have ever since been guarded by the application, to ***ever*** leave that application's control. And a security conscious organization like Apple can make the defensible claim that since all of the passkeys' security derives from the "secretness" of these private keys – which is critical – no other application, including its user, can or should be entrusted with their stewardship. Since this represents a powerful platform lock-in it's not at all clear to me that Apple will ever allow for passkeys export.

That being the case, I think that a very strong case can be made for only ever storing passkeys in a 3rd-party Passkeys client, such as a browser extension. In theory, it ought to be possible for a website to allow its user to replace one passkey with another. So if Apple or Android were to inadvertently become the generator and holder of a passkey, if a website supported passkey replacement it should be possible to migrate away from one passkey application to another.

Just to put a bit of frosting on this discussion before we talk about the problems with hardware authentication dongles, I wanted to share a few points from Google Chrome's FAQ on passkeys. They start off with all of the glowing bits:

> ***Manage passkeys in Chrome***
>
> *You can use a passkey to sign in easily and securely with just a fingerprint, face scan, or screen lock. Passkeys are a simple and secure way to sign in to both your Google Account and all the sites and apps you care about — without a password. You may be asked to sign into a website with a passkey or create one to improve your account's security.*
>
> *Tip: Passkeys are built on industry standards, so you can use them across many platforms.*

That all sounds terrific. What could possibly go wrong? Well, here's what Google has to say about that...

> ***Store passkeys in Windows***
> *If you have Windows 10 or up, you can use passkeys. To store passkeys, you must set up Windows Hello. Windows Hello does not currently support synchronization or backup, so passkeys are only saved to your computer. If your computer is lost or the operating system is reinstalled, you can not recover your passkeys.*
>
> ***Store passkeys in macOS***
> *You can save passkeys in your Chrome profile, where they're protected by a macOS Keychain.* ***Important:*** *Chrome can not save or use passkeys stored in iCloud Keychain. If your computer is lost or the Chrome profile is deleted, you can not recover your passkeys.*
>
> ***Store passkeys on a security key***
> *You can use a security key to store your passkeys.* ***Important:*** *Passkeys stored on security keys aren't backed-up. If you lose or reset the security key, you can't recover your passkeys.*

What a wonderful system! This clearly represents a huge step forward!

It's clear that, unfortunately, what we have at the moment is an extremely fragile system. The problem is the extreme secrecy surrounding the private keys which create passkeys. It's true that they do need to be guarded. Unfortunately, at the moment they're being jealously guarded. How Microsoft could possibly imagine that it's practical to have all of a user's passkeys locked up in a single machine, unable to synchronize with any of a user's other devices is beyond me.

But we're ready to entertain the second part of Dave's question where we asked:

> *Can the passkeys be paired across two user accounts - thereby ensuring recovery in case of loss with only 3 keys? - My mental model said it made sense but I do not know for sure ...*
> *1) Can the same key be applied to two different people?*
> *2) To assure full backup protection can all three keys be coded into both users?*
> *It may be a silly notion - but could it work or should I Just buy 4 keys to begin with? Thank you for all your good work and propeller-head installments! On to 999 and beyond!*

The answer is that not one of those operations is available. And what's more, I just double checked. As we learned last week, Yubico's Yubikeys have the most ample storage for passkeys in the industry, and even that is limited to a total of only 25. And they are utterly and absolutely non-exportable. A Yubikey is, at its heart, an HSM – a hardware security module. The internal Yubikey dongle hardware contains a very high-entropy random number generator that's used to synthesize a unique private key. That private key never leaves the device. There's no way to put a passkey in, and no way to take a passkey out. This would not be a problem if sites were to allow multiple passkeys to be registered for a single account. And there's no reason that would not be possible. But how many sites today support the use and management of multiple passwords for a single account? I've never seen one. So it's unclear why support for multiple passkeys would ever be created – even though nothing prevents it.

With Yubikeys having a 25-passkey limit, other than for experimentation, they seem most practical for higher-end enterprise-grade security application and perhaps for eventually signing into only a few of the most secure sites where the inconvenience of having an absolute hardware-lock is warranted by its ultimate level of security. And, as we noted last week, a Yubikey might be used to unlock a password manager, which is where, we would all have to conclude, all of a user's Passkeys should probably be stored.

The only sane conclusion we can draw is that while this is all very interesting none of this is yet ready for prime time. Poke at it, experiment with it, but wait until Bitwarden's passkey-supporting clients emerge from their current beta-testing state, at which point it will be practical to start depending upon passkeys. And even then, be very careful to only allow Bitwarden to generate and hold your passkeys even when other passkey clients on iOS or Android might be trying to.

Willie Scott / @WScottis1
… has some feedback and advice about the operation of the iCloud keychain:

> *Hi Steve,*
>
> *In regards to your discussion of Passkeys on last week's show, the part about the author's partner losing their iCloud Keychain passwords intrigued me. After the LastPass hack, I decided to switch to using iCloud Keychain for my passwords because I am in the Apple ecosystem and wanted to start using Passkeys instead of passwords wherever possible.*
>
> *I'm writing to mention that I too have had passwords and 2FA authentication codes wiped from my iCloud Keychain, although my Keychain has never been fully wiped unlike the poor partner's Keychain did. As near as I can tell, I believe I know the culprit of why it may be wiping credentials from iCloud Keychain and wanted to pass this along to anyone who might*

> *still be using iCloud Keychain to store their passwords (or knows someone who does).*
>
> *When I started changing all my passwords and adding accounts into iCloud Keychain, I noticed that an old Amazon password that I don't use anymore was already stored in there, probably from when the Amazon app asked, "Do you want me to remember your password?" It was an old password that I don't use anymore, so I deleted it. However, a couple days later, I noticed that even though I deleted that password (or so I thought), it had somehow reappeared in my iCloud Keychain. Not only that, but I also noticed that one or two accounts that I had recently added to the Keychain were missing and this process repeated itself a few more times. So that's when I started investigating.*
>
> *While digging through the settings, I went through my Apple ID account settings, and that's when I realized that my old iPhone 6S Plus, which was running an old version of iOS (iOS 14 to be exact), was still signed into my iCloud account and had iCloud Keychain turned on. I removed that old iPhone from my iCloud account and ever since I did that, no passwords have been wiped since. If you're in the Apple ecosystem, it's always a good idea to keep your devices up to date, but it might also be a good idea to do some "spring cleaning" and remove old Apple devices from your iCloud that you don't use anymore.*
>
> *Having said all of that, I sadly was agreeing with a lot of the points you were making about Passkeys and I think I've decided that I will probably switch over to BitWarden once Passkeys become officially supported in BitWarden (using https://bitwarden.com/twit of course).*
>
> *Thank you for a great show! I look forward to it each week! I'm also a proud SpinRite owner and can't wait to start using 6.1 on my SSDs and a troubled hard drive!*

This mysterious iCloud credential removal has the feel of something Apple would be deliberately doing out of their typical abundance of caution. I’ll bet there’s a security model behind it. For example, while an older iPhone is also signed into an account’s iCloud Keychain, Apple might be deliberately limiting what they’re willing to save into that shared Keychain while an older and presumably lower-security device also shares access. In other words, it’s a feature, not a bug!

Not So Fast

Today's podcast is titled "Not So Fast" because that's the absolutely best way to characterize what's going on in the United Kingdom with Google. As we know, during our podcast two weeks ago Leo dropped the news that Google's 3rd-party cookie deprecation would **not** be happening as had been long planned this summer. The abandonment and deliberate blocking of all 3rd-party cookies and other web-tracking hacks represents such a dramatic sea change for the web that many understandably skeptical observers doubt it can or will ever actually come to pass. So, self-confessed technology fanboy that I am, I wanted to determine what was going on. Were some stuffed-shirt bureaucrats somewhere going to screw this up? When I went to take a look at that for last week's podcast I quickly became lost in a paper shuffle. I decided that whatever was going on was worthy of understanding, since I consider this single forthcoming change to be one of the most important things that's going on today. As I've previously said, this represents a complete reconceptualization of the way the Internet finances itself.

The news that Leo had picked up on came in the form of an announcement that left more questions than it answered. On the 23rd of last month – which was, yes, Tuesday before last – on their PrivacySandbox.com sight, Google posted under the headline *"Update on the plan for phase-out of third-party cookies on Chrome"*:

Their brief introduction said: "The UK's Competition and Markets Authority (CMA) and Google publish quarterly reports to update the ecosystem on the latest status of Privacy Sandbox for the Web. As part of Google's Q1 2024 report, we will include the following update about the timeline for phasing out third-party cookies in Chrome in the April 26th report.

The update reads, simply...

> *We are providing an update on the plan for third-party cookie deprecation on Chrome.*
>
> *We recognize that there are ongoing challenges related to reconciling divergent feedback from the industry, regulators and developers, and will continue to engage closely with the entire ecosystem. It's also critical that the CMA has sufficient time to review all evidence including results from industry tests, which the CMA has asked market participants to provide by the end of June. Given both of these significant considerations, we will not complete third-party cookie deprecation during the second half of Q4.*
>
> *We remain committed to engaging closely with the CMA and ICO and we hope to conclude that process this year. Assuming we can reach an agreement, we envision proceeding with third-party cookie deprecation starting early next year.*

And then they conclude by noting: "Once published, you will be able to view both Google and the CMA's full reports." Those reports were published three days later, on April 26th.

This entire issue is described best by the following statement:

> *On 7 January 2021, the CMA commenced an investigation under section 25 of the Act in relation to Google's Privacy Sandbox proposals. The CMA subsequently informed Google that the CMA was concerned that Google's proposals, if implemented without regulatory scrutiny and oversight, would be likely to amount to an abuse of a dominant position.*

It's unclear and not really important to know the genesis of this inquiry. Since we're talking about the elimination of all 3rd-party cookies and the curtailment of what had become the widespread practice of tracking Internet users around the web as a means of determining their interests, it may well have been that advertising technology companies based in the UK were crying foul behind the scenes.

What ensued was about what you'd expect from any healthy and well established bureaucracy as old and wizened as the United Kingdom. Experts were found, neutral 3rd-party "monitors" were enlisted and Google created a document describing the **commitments** it was prepared to make. A document titled "Investigation into Google's 'Privacy Sandbox' browser changes" opens with the assertion that *"The CMA has accepted commitments offered by Google that address the CMA's competition concerns resulting from investigating Google's proposals to remove third-party cookies and other functionalities from its Chrome browser."* ... which begs the question, what exactly are these commitments that the CMA has accepted?

I found the points of concern in the description of the roles of the appointed technical expert that will be supporting the monitoring agent. The document states:

> *On 26 September 2022, the CMA approved the appointment of S-RM Intelligence and Risk Consulting Limited by the Monitoring Trustee (ING Bank N.V.) as an independent Technical Expert to support the Monitoring Trustee in monitoring compliance with the following provisions of the binding commitments accepted by the CMA on 11 February 2022:*
>
> *Google's use of data (paragraphs 25 to 27), non-discrimination (paragraphs 30 to 31) and (with respect to those provisions) anti-circumvention (paragraph 33). The role of the Technical Expert is to provide specialized knowledge to support the Monitoring Trustee, particularly in relation to monitoring of data flows, and understanding the possible impacts of the Privacy Sandbox changes on ad tech markets.*

Okay. So we have the ING Bank serving as the neutral monitor and this monitor has appointed another firm with the required technical expertise. And everything it focused upon a small handful of paragraphs somewhere. I found out where. They are in Appendix 1A of the latest version of the "Google's final commitments" document. The first set of paragraphs, 25 through 27, amount to Google promising not to use any Personal Data from a user's Chrome browsing history, a customer's Google Analytics account, or to in any way track users. So that's all pretty much what Google has explained to be its intentions and goals. So it appears that the CMA just wanted that very clearly and succinctly spelled out.

The non-discrimination, paragraphs 30 and 31 state that Google promises to create a totally level playing field. Having examined, explored and shared on this podcast the operation of Google's cookie-replacement technologies as they have evolved through the years, this was clear to me. It's implicit throughout Google's design – though it **has** also grown to be much better thanks to all of the feedback and criticism the various pieces have received. And I can understand how bureaucrats, who will never understand how Google's "TOPICS" API functions, need a simple "what does it mean" spelled out in English. Since this is crucial to the acceptance of Google's technology, I'm going to share Paragraphs 30 and 31. Paragraph 30 says:

> *30. Google will design, develop and implement the Privacy Sandbox proposals in a manner that is consistent with the Purpose of the Commitments and take account of the Development and Implementation Criteria, Google will ensure that it does not distort competition by discriminating against rivals in favor of Google's advertising products and services. In particular, Google will not:*
>
> *a. design and develop the Privacy Sandbox proposals in ways that will distort competition by self-preferencing Google's advertising products and services;*
>
> *b. implement the Privacy Sandbox in ways that will distort competition by self-preferencing Google's advertising products and services; or*
>
> *c. use competitively sensitive information provided by an ad tech provider or publisher to Chrome for a purpose other than that for which it was provided.*
>
> *For the avoidance of doubt, Privacy Sandbox proposals that deprecate Chrome functionality will remove such functionality for Google's own advertising products and services as well as for those of other market participants.*
>
> *31. Google will not change its policies for customers of Google Ad Manager, Campaign Manager 360, Display & Video 360 or Search Ads 360 to introduce new provisions restricting a customer's use of Non-Google Technologies before the Removal of Third-Pay Cookies, unless in exceptional circumstances (such circumstances to be discussed with the CMA) or as required by law. For the duration of the Commitments, Google will inform the CMA ahead of any such change to these policies.*

And this leaves us with the "anti-circumvention" paragraph 33 which is just a single line which reads: *"33. Alphabet Inc., Google UK Limited and Google LLC will not in any way, whether by actions or omissions, directly or indirectly, circumvent any of the Commitments."* This sort of language will be familiar to any businessman or anyone who's been involved in any contractual agreements where attorneys were engaged. And it's important to understand that both the United Kingdom government and Google's various corporations recognize those provisions to be contractually and legally binding.

So it has been upon those representations, which are enumerated as "Commitments" with a capital "C", that the UK then proceeded to carefully examine Google's proposal. So now we return to the timeline for phasing out 3rd-party cookies. That work appears in a document titled "CMA Q1 2024 update report on implementation of the Privacy Sandbox commitment" dated last month, April 2024. The document's Summary lays out the entire story and it's interesting enough and short enough to share:

> *This report sets out the CMA's updated views on the issues we identified in our January 2024 report concerning Google's proposed Privacy Sandbox changes (see Annex 1). Our analysis is based on the framework for assessment set out in the legally binding Commitments that Google made in February 2022 to address competition concerns relating to its proposals to remove third-party cookies from Chrome. The January 2024 report set out our provisional views on the impact of the Privacy Sandbox on competition, publishers and advertisers and user experience.*
>
> *We outline Google's response to the concerns we identified in that report and the steps it is taking to resolve pending issues. We have also considered the feedback received from market participants on these points. We have included a summary of this feedback in the sections below.*
>
> *This report also incorporates the preliminary assessment of the Information Commissioner's Office (ICO) on the privacy and data protection impacts of the Privacy Sandbox. Having consulted with the ICO, we set out our current views on these concerns for each of the APIs.*
>
> *Although there are a number of concerns to work through, based on the available evidence, we consider that from 1 January 2024 to 31 March 2024 (the relevant reporting period), Google has complied with the Commitments. This means that in our view Google has followed the required process set out in the Commitments and is engaging with us (and the ICO) to resolve our remaining concerns ahead of third-party cookie deprecation. However, further progress is needed by Google to resolve our competition concerns ahead of deprecation.*
>
> *We will continue to work with Google to resolve our concerns between now and the point at which Google triggers the Standstill Period. We will provide an update on progress in our next update report.*
>
> *Testing of the Privacy Sandbox tools is also currently underway. The test results will form part of a wider evidence base that we will use to assess the effectiveness of the Privacy Sandbox. The test period runs until the end of June this year.*
>
> *Given the time needed to resolve outstanding issues and take account of testing results, we have agreed with Google that there should be a limited delay to third-party cookie deprecation. Subject to resolving our remaining competition concerns, Google is now aiming to proceed with third-party cookie deprecation starting in early 2025. Under the Commitments, it is for Google to decide when the Standstill Period is triggered.*
>
> *We encourage market participants taking part in testing to submit their results directly to us by the end of June deadline. We also welcome any additional feedback from stakeholders on the concerns identified in this report. Our contact details are included at the end of this report.*

This made reference to a "Standstill Period" several times so I tracked that down in the earlier Commitments documents. It appears to just be more bureaucracy for its own sake. It says:

> *19. Google will not implement the Removal of Third-Party Cookies before the expiration of a standstill period of no less than 60 days after Google notifies the CMA of its intention to implement their Removal. Google may increase the length of such a standstill period at any time between giving such notice and the period's expiry. At the CMA's request, Google will increase the length of this standstill period by a further 60 days to a total of 120 days.*

What follows that document summary are 97 pages of interesting, but ultimately mind-numbing, back and forth detail, as every conceivable facet of the big change Chrome will be implementing in Chrome is examined under a bureaucratic microscope. The real concern is over Google's size and whether the changes it is making will disadvantage smaller ad tech players. But what becomes clear after reading some of it is that both parties are truly negotiating in good faith. No one is stonewalling, no one is being unreasonable, and true progress is being made, even though it's slow as molasses.

It does not appear to me that Google's Privacy Sandbox technology is in any trouble at all. The truth is, this does represent a massive change to the way the Internet pays for itself and it's also true that many companies whose revenue has been entirely derived from the oh-so-slimy practice of tracking users and aggregating their data, without their knowledge or permission, for the purse of then selling it to anyone with a wallet will be significantly impacted. And not in a good way.

So, having read through the documents, I can understand the process that's taking place. And in retrospect, though I would not have predicted this would happen, it's at least understandable and it appears that the world will, indeed, soon be receiving this dramatic change in the way interest-based advertising is carried out. It's clearly far superior to the status quo. And as we know, changing any deeply entrenched status quo, especially when people's livelihoods will be impacted, is never easy.

ZTDNS – Zero Trust DNS

Last Thursday, Microsoft published a preview of a forthcoming security solution they call "Zero Trust DNS". It's been clear for a long time that DNS represents both an Achilles heel of network security and a point where it's also very possible to introduce a significant new level of security. From my brief scan of the technology Microsoft has outlined, it appears that any of our listeners who may have followed-up on my discovery of Adam Networks' DNS solution, which they call "Don't Talk to Strangers", may already be enjoying the benefits of dramatically improved security thanks to leveraging the power of DNS. So, for next week's podcast, I plan to take a deep look into what Microsoft has announced. One thing that immediately stood out was that Microsoft might be attempting to use it as a way of driving enterprises to Windows 11 since they are explicitly labeling the clients as Windows 11 machines. Since no one actually wants Windows 11, since Windows 10 still commands more than twice the number of desktops as Windows 11 (and a much greater percentage within the enterprise), and since a huge installed base of machines won't even run Windows 11, if what Microsoft is planning to do is truly a Windows 11 only solution, then the client-agnostic system that the Adam Networks guys already have working and well-proven seems like a far more practical choice. But in any event, by the end of next week's podcast we'll all know exactly what's going on.