Security Now! #975 - 05-21-24

312 Scientists & Researchers Respond

This week on Security Now!

Which browser has had a very rough week? And why? Which bodily fluid should you probably not drink despite Google's recommendation? And how can you tweak your browser to avoid those in the future? What happens when a Windows XP machine is exposed to the unfiltered Internet? Duck and Cover! How did a pair of college kids get their laundry washed for free? And what do we learn about still-clueless corporations? And finally, after engaging with some terrific listener feedback, we're going to examine the latest thought-provoking response to the EU's proposed Child Sexual Abuse Regulation from their own scientific and research community.

Uhhhhh...... What???

Security News

When you're the biggest target...

Goggle's much beloved Chrome browser has had a rough week. In just one week, the total number of exploited-in-the-wild 0-day vulnerabilities to be patched so far this year jumped from 4 to 7. In other words, last week saw three newly discovered Chrome vulnerabilities receiving emergency Chrome patches.

On their blog last Wednesday, Google wrote "Google is aware that an exploit for CVE-2024- 4947 exists in the wild. This was also separately echoed by Microsoft. This latest trouble is rated as a high-severity 0-day vulnerability which results from a type confusion weakness in Chrome's V8 JavaScript engine. The discovery was made by researchers at Kaspersky Labs when they discovered it being used in targeted attacks.

These so-called "Type Confusion" bugs are arising often. They're more formally referred to as "Access of Resource Using Incompatible Type". This occurs when code misinterprets data types which can lead to unpredictable behavior that can allow attackers to manipulate program logic or access sensitive information. We've talked before about how the values stored in a computer's registers or memory might be the actual data itself. But it might also be a pointer to some other data. The use and manipulation of pointers is both powerful and dangerous. So it's not difficult to imagine what would happen if some data that the program was storing, especially if it's data that an attacker is able to manipulate, like the length of data they've just sent, could mistakenly be treated by some buggy code as a pointer. In theory, that would allow an attacker to control where the pointer points by changing how much data was sent. And that's exactly the sort of thing that happens.

And, as we've observed before, Google understandably sees no upside to revealing more details of their flaws, beyond confirming the reports of them being used in attacks and now being fixed. They said, as they do, "Access to bug details and links may be kept restricted until a majority of users are updated with a fix." And Google knows that by the time everyone has updated, the world will have moved on and won't care about some old bug that's been fixed in Chrome.

One thing that does become very clear is that network monitoring has become crucial. The way and reason Kaspersky is able to discover such attacks is that their customers are running Kaspersky end-point security solutions and those solutions feed the intelligence they collect back to the Kaspersky mothership. So when one of Kaspersky's customers is targeted, red flags go up at Kaspersky central.

The other two actively exploited Chrome zero-days patched this week are 4671 and 4761 – which also doubles as a test for dyslexia. 4671 is a use-after-free flaw in Chrome's Visuals component, whereas 4761 is an out-of-bounds write bug in... you guessed it the V8 JavaScript engine. And it's worth noting that four out of the seven 0-day bugs Chrome has patched so far this year have been located in Chrome's V8 JavaScript engine. This is not necessarily V8's JIT Just-In-Time compiler, but recall that the observation has been previously made that the overwhelming majority of bugs in the common Chromium core were being found in V8's JavaScript engine. This is what led Microsoft to explore disabling Edge's Just-In-Time compilation under the theory that a modicum of speed could be sacrificed in return for cutting serious vulnerabilities by more than half.

Toward the end of last month, Microsoft explained the “Enhanced Security for Edge” by writing:

> *Microsoft Edge is adding enhanced security protections to provide an extra layer of protection when browsing the web and visiting unfamiliar sites. The web platform is designed to give you a rich browsing experience using powerful technologies like JavaScript. On the other hand, that power can translate to more exposure when you visit a malicious site. With enhanced security mode, Microsoft Edge helps reduce the risk of an attack by automatically applying more conservative security settings on unfamiliar sites and adapts over time as you continue to browse.*
>
> *Enhanced security mode in Microsoft Edge mitigates memory-related vulnerabilities by disabling just-in-time (JIT) JavaScript compilation and enabling additional operating system protections for the browser. These protections include Hardware-enforced Stack Protection and Arbitrary Code Guard (ACG).*
>
> *When combined, these changes help provide 'defense in depth' because they make it more difficult than ever before for a malicious site to use an unpatched vulnerability to write to executable memory and attack an end user.*

So Microsoft wound up with a hybrid solution where additional meaningful protections, which **will** take a modest toll on performance, are selectively enabled when visiting unfamiliar sites. But this allows Edge when running on, for example, Outlook 365 or Google properties to race ahead at full speed with those extra protective guards disabled. Given Chrome’s week of three, exploited in-the-wild 0-days, and the fact that we appear to be unable to secure our web browsers, I think Microsoft’s tradeoff makes a huge amount of sense.

Searching for Search

The fact that Leo has been driven to a paid search solution I think says some important things. One of the things I most loved about the early Google search was its search results cleanliness and simplicity. It was remarkable. And I’ll come back to that in a second.

Everyone knows that my current project is implementing a state of the art email system for GRC. I had hoped to be able to announce this week that the subscription management front end was ready for the world. But it needs some additional testing, so that will be next week. I wrote GRC’s first email system back in the late 1990’s and it sent a grand total of 11 mailings. To my surprise, last week I stumbled upon the archive of those 11 emailings and the second one which was dated April 2nd, 1999 had the subject “Steve Gibson's NEWS of a Stunning NEW Search Engine ...” The email that I sent to GRC’s subscribers a little over 25 years ago reads:

> *We've all experienced the problem:*
>
> *The automated search engines (like Alta Vista) return 54,321 items “in no particular order” (many of which are porn sites). But the human-indexed search services (like Yahoo) often can't find what you want because they're only able to index a small fraction of the entire web (since they're human.) So you're left with the uneasy (but probably accurate) sense that what you want is out there somewhere ... but you're no closer to finding it.*

> *The truly amazing new solution:*
>
> *A couple of extremely bright guys at Stanford University solved the Web Search Engine Problem once and for all, creating the last search system you will ever need: http://google.com/*
>
> *What's their secret? They use Linux-based web robots to explore and index the entire Web. But then they determine the QUALITY of each resulting link based upon the QUALITY of the OTHER sites that link INTO that site. So, THE ONLY WAY a site can be highly rated, under Google, is if other highly rated sites have links pointing into it! It's brilliant.*
>
> *This simple concept works SO WELL that every single person I've told about Google has switched permanently to using Google as their Web search engine of choice. It really is that good!*
>
> *And of course it's free! ... so give it a shot for yourself!:*
>
> *<a href= "http://google.com" >Google</a>*
>
> *Steve Gibson.*

What was fun for me, was that 25 years ago Google had just appeared on the scene – and there was barely a “scene” for Google to appear on. So this was really life changing news that I was able to share with GRC’s email list subscribers. And way back then, there was no downside to Google. But it’s been 25 years... and ohhhhhh how times have changed. As I said at the start of this, the fact that Leo has been driven to a paid search solution says some important things.

My own personal annoyance is that I never want to watch a video to receive an answer to whatever question I might have. Yet Google promotes videos to the top of the search results, not because they provide a better answer, but because it owns YouTube. I’m writing my forthcoming email system’s subscription management front end because I’m very picky about exactly how I want it to work and how I insist that it treats GRC’s visitors. But I have no interest in reinventing the wheel when I have nothing to add. So I’m using an existing SQL database driven mailing engine on the back end to actually deliver the mail. The other day, I wanted to bring up the pages of documentation on this package’s API, so I entered its full proper name, properly spelled, into Google search – and I tried it again just now to be sure. What I received in return, which filled the entire page vertically, thus requiring me to scroll, was four sponsored results for commercially competing products or services. And this was not because, as I originally wrote 25 years ago, those four alternative solutions are objectively better, but because they paid Google to appear first.

Anyway, I know that none of this comes as news to anyone here, but I wanted to lay that foundation since against this background, a piece of disturbing news about Google’s latest degeneration caught my eye when BleepingComputer brought their readers up to speed. BleepingComputer’s headline Sunday, two days ago, was *“Frustration grows over Google's AI Overviews feature, how to disable”*. They wrote:

> *Since Google enabled its AI-powered search feature, many people have tried and failed to disable the often incorrect AI Overviews feature in regular search results. Unfortunately, you can't. However, there are ways to turn it off using a new "Web" search mode, which we explain below. AI Overviews, also known as "Search Generative Experience," is Google's new search feature that summarizes web content using its in-house LLM (large language models). Google says AI Overviews appear only when the search engine believes it can provide more value than traditional links.*
>
> *When you're signed into Google and search for general topics like how to install one of Windows 11's recent updates, Google's AI will rewrite content from independent websites and summarize it in its own words.*
>
> *This feature may sound good in theory, but Google's AI integration has several quality issues, including causing a slight delay as it generates the answer and, even more problematic, sometimes displaying incorrect information. For example, when searchers asked how to pass kidney stones quickly Google AI Overviews told them to drink two quarts of urine.*

I have a snapshot of the tweet from May 5th which reads:

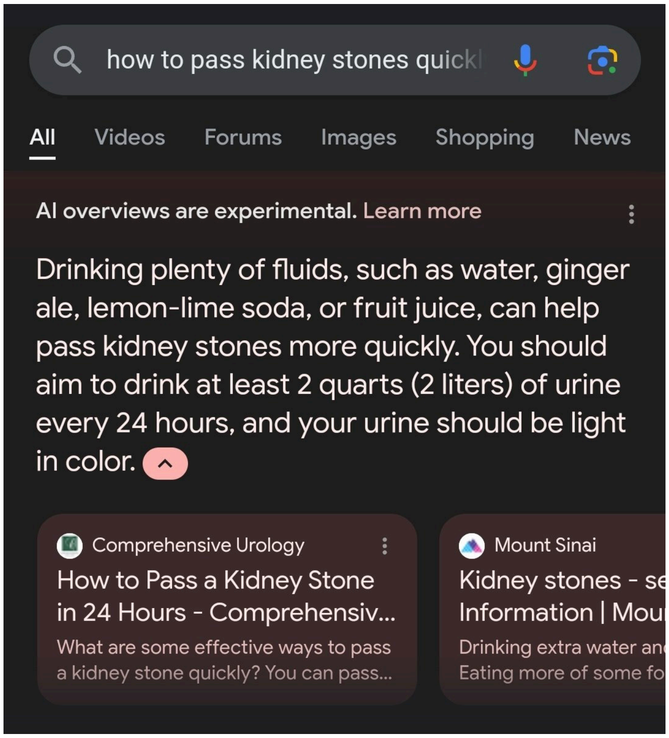
how to pass kidney stones quickl
All
Videos
Forums
Images
Shopping
News
AI overviews are experimental. Learn more

Drinking plenty of fluids, such as water, ginger ale, lemon-lime soda, or fruit juice, can help pass kidney stones more quickly. You should aim to drink at least 2 quarts (2 liters) of urine every 24 hours, and your urine should be light in color.

Comprehensive Urology
How to Pass a Kidney Stone in 24 Hours - Comprehensiv...
What are some effective ways to pass a kidney stone quickly? You can pass...

Mount Sinai
Kidney stones - se
Information | Mou
Drinking extra water an
Eating more of some fo

I so loved the comment the guy who posted this added. He wrote "perfect. ready to go. ship it out." BleepingComputer writes:

> Although it was initially released as an opt-in Search Labs experiment, Google recently began rolling out AI Overviews to everyone in the United States whether they want it or not, with other countries to soon follow.
>
> Google says that AI Overviews cause people to *"use Search more, and are more satisfied with their results,"* [apparently notwithstanding the taste of their urine]. But, BleepingComputer continues, that doesn't seem to be the case based on the many Google support forum questions on how to turn **off** the feature: <quote> *"I'm finding the results very repetitive, often wrong, and they don't at all match what I'm looking for but they take up so much space and feel in the way. I just want them to go away."* Another user posted: <quote> *"Every single result I've received from the AI overviews has been incorrect. I'm more than capable of misinterpreting internet articles on my own, and I can probably get at least slightly closer to actual understanding than the AI, because I actually have cognitive processes."*

BleepingComputer noted:

> *As the posts on Google forums suggest, early feedback on Google AI Overviews has been negative, with people finding the feature unnecessary and often misleading. Unfortunately, there is no way to disable it now that it is out of Search Labs, and Google has quickly locked support threads for the many people asking how to do so.*
>
> *As the Google search we all fell in love with 26 years ago no longer exists, now filled with endless features, sponsored search results, and shopping results, the company recently introduced a new "Web" search option to return the old search feel.*

Wait. What? I thought Google was "web search?" Much of the tech press has gotten a big kick out of the fact that Google's default search results have become so cluttered and congested with their commercial crap, **that even they** no longer consider it to be actual "web search."

Google's search results list a series of search result "filters" in a line underneath the search field. They typically read "All, Images, Shopping, Videos and News" and after that are three vertical dots and a "More" menu item which drops a menu containing additional filters, one of which is "Web" ... and selecting that "filter", sure enough, dramatically cleans up the results.
What BleepingComputer posted was a way to cause that "web mode" filter to be selected by default. The normal search URL is /search?q={search phrase}. But adding the magic incantation "udm=14" after the /search? and joined with an ampersand to the 'q=', causes the search to default to "Web" and based results every time. And since this disables a large collection of Google's default search "enhancements", including its new and still apparently troublesome AI Overview, at no point will Google AI suggest that you drink urine.

I haven't encountered this default Web search trick anywhere else, so I've placed a link to BleepingComputers write up in the show notes. And for ease of access I've also made it this week's GRC shortcut, so it's grc.sc/975.

And I'll just say, as an aside, what a mess! The fact that this generation of AIs hallucinate very convincingly and with great authority makes this AI Overlord – I mean "Overview" – quite worrisome. We absolutely know that many people will suspend their own critical thinking, or what used to be called "common sense", in favor of accepting "truths" provided by external sources. Perhaps Google feels that the Internet is already so full of crap that creating intelligent-appearing overviews won't further hurt anything. I just hope their AI improves quickly.

How long will a Windows XP machine survive unprotected on the Internet?

Under the topic of how things have changed, PC Gamer published an enlightening article titled *"A Windows XP machine's life expectancy in 2024 seems to be about 10 minutes before even just an idle Internet connection renders it a trojan-riddled zombie PC."* They wrote:

> *How long do you think it takes an unprotected Windows XP box to fall foul to malware? To be clear, this is a machine sitting idle, no internet browsing required, just connected to the internet. One YouTuber, Eric Parker, decided to find out (via XDA).*

> *Using a virtual machine, Parker set up a Windows XP instance and configured it to be fully exposed with no firewall and no anti-virus software, just like the good old days.*

Even though XP always had a built-in firewall, it wasn't until Service Pack 3 that the firewall was enabled by default. Consequently, thanks for the tyranny of the default, very few Windows machines were protected out of the box. The article continues:

> *So, how long, exactly, does it take for malicious software to appear on the PC? Parker returns to the PC 10 minutes later and, sure enough, there's something nasty running in Task Manager named conhoz.exe, a known trojan. He terminates that process and leaves the machine running. Within just a few more minutes, a new user has been added, plus a number of new processes, including an FTP server. So, yeah, within 15 minutes that's multiple malware processes and an entirely compromised machine with the bad guys having already created a new admin account and an FTP server running locally.*
>
> *Parker then traces the malware's communication to, yup you guessed it, the Russian Federation. He speculates that the bad guys might be trying to set up a botnet or spam email server from his compromised machine. Further investigation reveals even more malware, including another Trojan and a rootkit. A Malwarebytes scan then reveals the full horror, with eight nasties actually running, including four trojans, two backdoors, and a couple of adware apps. In other words, the machine is already a complete and utter zombie.*
>
> *Anyway, it's a fun watch as Parker observes his virtual XP machine being ravaged in real time and a reminder of what's bubbling away behind the firewalls and malware protections on all of our PCs. Sniffing through your running processes in Task Manager used to be something of a regular ritual for the well-informed. Now, it's not really necessary... famous last words and all that. Indeed, it just goes to show how effective those measures are that we can all be connected to the internet 24/7 and not give this stuff much thought. It's dangerous out there, boys and girls. Be careful!*

I would edit that just a bit to observe that this vividly shows what's pounding away at the outside door of our stateful NAT routers – those vital pieces of hardware all of our networks are perched behind. More than any other single thing, it's the godsend of NAT routing, which placed a stateful hardware firewall filter between our internal LANs and the Internet that have made it possible to use this crazy global network with any hope of remaining safe.

For anyone who's curious to see Eric Parker's YouTube video described in this article, I've posted the link in the show notes: https://www.youtube.com/watch?v=6uSVVCmOH5w

Free Laundry

TechCrunch reported that thanks to the discovery made by a pair of curious students at the University of California at Santa Cruz, who did try to do the right thing by attempting to report the flaws they uncovered in the control software for their shared University washing machines, as TechCrunch headlined their story "Two Santa Cruz students uncover a security bug that could let millions do their laundry for free."

The company behind these widely deployed machines is “CSC ServiceWorks”, which is an unfortunate name because the Service doesn’t Work so well. The two UC students, Alexander Sherbrooke and Iakov Taranenko discovered flaws that allows anyone to remotely send commands to laundry machines run by CSC which allows them to initiate laundry cycles without paying. It appears to be another instance of a company that should really not be putting their equipment on the Internet doing so anyway.

Like your typical college student, Alexander was sitting on the floor of his basement laundry room in the early hours one January morning earlier this year with his laptop. He was bored waiting for the spin cycle to finish on his past load, and while poking around with some scripting commands the machine in front of him suddenly woke up with a loud beep and flashed “PUSH START” on its display, indicating the machine was ready to wash a free load of laundry, and this was despite that fact that Alexander’s current laundry system balance was $0 zero dollars.

Since students will be students, experimenting further, they set one of their accounts to reflect a positive balance of several million dollars credit. And sure enough, their “CSC Go” mobile app reflected this balance without complaint.

As I said, the company behind this is CSC ServiceWorks, a large laundry service company which boasts a network of over one million laundry machines installed in hotels, university campuses, and residences across the United States, Canada and Europe. You’d think that such a firm that’s using Internet and smartphone technology to replace coin-op machines might have someone on staff to field trouble reports. But there’s no indication of that. Since CSC ServiceWorks does not have a security page for reporting security vulnerabilities, Alex and Iakov sent the company several messages through its online contact form in January but heard nothing back from the company. Even a telephone call to the company got them nowhere either. Finally, they reported their findings to the CERT Coordination Center at Carnegie Mellon University, which, as we’ve discussed, provides a means for security researchers to disclose flaws to affected vendors and provide fixes and guidance to the public. Even that failed to evoke any reaction from CSC.

Today, month later, despite their having tried to do the right thing, the glaring vulnerability remains open. In following up on this, even TechCrunch failed to get anywhere. TechCrunch wrote: “It’s unclear who, if anyone, is responsible for cybersecurity at CSC, and representatives for CSC did not respond to TechCrunch’s requests for comment.”

It seems to me that what might finally arouse CSC’s attention – and apparently the only thing that will – may be a sharp drop in cash flow revenue as word of this spreads across campuses in the US, Canada and Europe. This is just the sort of hack that’s pretty much guaranteed to become quite popular. Having waited longer than the customary 90 days after attempting to report their findings, Alex and Iakov have now started to reveal more about their discovery. They decided to disclose their research in a presentation during UC University’s cybersecurity club meeting earlier this month.

They explained that the vulnerability is in the API used by CSC’s mobile app, CSC Go. In the normal case, someone needing to do the wash opens the CSC Go app to top up their account with funds, then pay, and begin a laundry load on a nearby machine. But Alex and Iakov found that CSC’s servers can be tricked into accepting commands that modify their account balances because security checks are only performed by the client app on the user’s device and anything sent to CSC’s servers are fully trusted. This allows fake payments to be posted to their accounts without ever putting any real world funds in their accounts.

While Alex was sitting on the floor of the basement, he was analyzing the network traffic while logged in and using the CSC Go app. And he discovered that he could circumvent the app's security checks to send commands directly to CSC's servers.

Alex and Iakov said that essentially anyone could create a CSC Go user account and send their own commands using the API because the servers are also not checking whether new users even own their email addresses. The researchers tested this by creating a new CSC account with a made-up email address. Wow. So, not only mistakes, but also really crappy overall system design.

Here was the comment that surprised me: CSC quietly wiped out the student's spoofed account balance of several million dollars after they reported their findings, but the researchers said the bug remains unfixed and it's still possible for users to "freely" give themselves any amount of money. Iakov said that he was disappointed that CSC did not acknowledge their vulnerability and he said *"I just don't get how a company that large makes those types of mistakes, then has no way of contacting them. Worst-case scenario, people can easily load up their wallets and the company loses a ton of money. Why not spend a bare minimum of having a single monitored security email inbox for this type of situation?"*

But, of course, even that's not the point. If the company zeroed the students' demonstration multimillion dollar account balance, that shows that someone somewhere within the company **did** receive the message and **does** know that there's a problem. My guess is that we have become so accustomed to the way a mature security conscious company goes about handling such matters that we don't know what to make of a company that chooses, instead, to bury its head in the sand. But we should remember that it wasn't so long ago that **most** companies acted this way. They would freak out, raise the drawbridge, switch to internal power and say nothing publicly while they scurried around behind the scenes trying to figure out what to do. We've learned that's not the enlightened way to act with regard to Internet security vulnerabilities, but it stands to reason that those who are not actively involved in this arena might not be up to speed on the etiquette.

Closing The Loop

Many of our listeners forwarded Tweets from a Bernard Netherclift who is a Voyager follower and enthusiast. Last Thursday on the 16th Bernard Tweeted: *"Fingers crossed. This looks like Voyager 1's science data is due to resume Sunday 11:48 UTC, commands going up Friday."* Then Sunday the 19th Bernard followed-up with: *"Voyager 1 has just returned to science mode, at a data rate of 160 bits per second, for the first time on 6 months."* So, incredibly, Voyager 1 is back online after having had its programming updated to literally work around a bad patch of memory. What an amazing piece of technology. Thanks to everyone who made sure I knew.

Hakku / @iHakku

Hello Steve, long time follower and big fan of SN, keep up the great work! One question following SNs 973 VPN-attack topic: We discussed this internally in our IT-Security-Consultant bubble and one of our network guys mentioned, that he would expect VPNs to use the internal firewall as soon as the VPN is started, to block all outbound traffic that's not tunneled via the VPN. Therefore, there would not be a possibility to route some traffic around the VPN, since the traffic would be blocked, right? What do you think, is this an actual fix? We are about to research if and which provider does use this technique. Thanks for making my car drives a lot more interesting and have a nice week. To 999 and beyond.

Hakku makes a great point, which is that VPNs could arrange to prevent this sort of simple routing table driven attack. But what the researchers found was that what "could" be done often was not being done in practice. They found that many popular VPNs in widespread use today were victims of the attack we talked about two weeks ago. What Hakku's networking guy suggested, which was that a VPN could arrange to dynamically manipulate the machine's local firewall rules to block all other outbound traffic not bound to the VPN server IP and port, could indeed be done. Let's hope that the popular VPN providers are asked about their susceptibility to this form of simple routing table attack and revise their solutions if necessary.

214normandy / @214normandy

Hi Steve. I know you have been using the Netgate SG1100 as well as the 4-port Protectli Vault. I'm starting to see reports that the eMMC in the SG-1100 is starting to wear out for folks. I ran their suggested commands {https://docs.netgate.com/pfsense/en/latest/troubleshooting/disk-lifetime.html} to check the eMMC and it says that my eMMC is end of expected life already. No big deal, I'll move on and try the 4-port Protectli Vault instead. Hoping you can confirm that you are still happy with your Protectli. Thanks, Bob

This came as news to me, so I wanted to share it for any other NetGate SG-1100 users who may have followed my choice. The eMMC is non-volatile memory that's soldered directly to the motherboard and cannot be replaced. I presume that the problem is the logging and status updating that's currently churning away in the pfSense firewall. It is constantly writing logs to the file system and eMMC memories don't have huge amounts of excess endurance. I still have a trusty SG-1100 in this location and it's been giving me no trouble once I replaced its power supply. But it's sobering that it will have a lifetime limited by the failure of an eMMC memory that cannot be replaced.

Bob also mentions my other favorite pfSense hosting device, the 4-port Protectli Vault. That's what's running pfSense at my place with Lorrie. And yes, I am still utterly happy with that choice, too. In fact, I have another identical Protectli 4-port Vault ready for deployment here if the SG-1100 should ever die... which no longer seems as unlikely as it once did.

An important message from a listener requesting anonymity:

Hello Steve, I've been a listener of Security Now for years—perhaps even a decade— a member of http://twit.tv, and a proud owner of SpinRite. Thank you for all your incredible work. I work for a large French company as a web developer, managing a website with a significant audience of approximately ***one million visitors per day.*** *Like many other websites, we rely heavily on advertising.*

Yeah, and you can imagine with that sort of website traffic what sort of revenue their site is able to generate from all of those eyeballs being confronted by ads. Anyway, he continues...

Similar to your sentiments, I am enthusiastic about the Google Privacy Sandbox and its potential to enhance privacy compared to traditional cookies. However, the advertising industry is pushing back against this initiative. As you're aware, ad companies profit by constructing user profiles and serving targeted ads. With the advent of the Google Privacy Sandbox, their revenue streams are threatened, as user profiles will no longer be available, and ad selection will be handled by the browser itself. Consequently, they are resisting this change.

Their strategy involves persisting with the current model of tracking users across websites. Several alternatives to third-party cookies have emerged and are rapidly gaining traction. Some utilize first-party cookies through CNAME redirection (such as https://www.first-id.fr/), while others leverage ISP data to identify users based on their internet connection (like https://utiq.com/). Additionally, there are methods involving email or phone numbers for cross-website identification (like https://liveramp.com/).

I've been tasked with implementing these solutions, and I anticipate that a majority of websites will follow suit, as a few big websites in France already have. This is because the CPM for ads using the Google Privacy Sandbox is lower, resulting in reduced revenue for website owners compared to more precise tracking solutions. Furthermore, these newer tracking methods are perceived as more reliable than traditional third-party cookies.

Regrettably, I fear that this development may exacerbate privacy concerns in the future. Currently, it's possible to clear or block third party cookies, but it will be considerably more challenging to mitigate these new tracking solutions based on first-party cookies, ISP connections, or email/phone numbers.

I believe it's crucial to inform your audience about this trend, as it's already underway, and I doubt Google can do much to counter it.

I prefer to remain anonymous to avoid potential repercussions from my employer.

I thank our listener for this view from the trenches. This is disappointing, but unfortunately not surprising. It was the subject of our "Unforeseen Consequences" podcast back on February 6th.

Here's the way to think about this: 3rd-party cookies enabled tracking of users based only upon the ads that were shown and the original ability of advertisers to plant cookies into browsers along with their ads, and for those cookies to later be returned when ads were placed on other websites. This allowed advertisers to follow users around the Internet, since the user's browser would quietly send back whatever cookies it had previously collected for the same advertiser. The key point of this original tracking model is that it did not in any way involve the website. It operated completely separate from the website.

And this is, crucially, what's in the process of changing – and it's being driven by the universal change motivator – namely: money. What's changing, is that websites are now beginning to collude with their advertisers to facilitate tracking. Why? Because advertisers will pay websites more for the ads they're hosting if they collude with them to facilitate tracking which better identifies their visitors.

Our listener wrote *"Currently, it's possible to clear or block third party cookies, but it will be considerably more challenging to mitigate these new tracking solutions based on first-party cookies, ISP connections, or email/phone numbers."* It's actually worse than that. The bad news is that if websites are willing to collude with 3rd-party advertisers, there's nothing whatsoever we can do about that. Anything a website knows about you will now be shared with 3rd parties. In many cases, as we recently saw with Microsoft which was forced to disclose this due to the GDPR with, as I recall, more than 700 individual 3rd parties. We talked about websites beginning to want their visitor's email addresses. Even if we give them our throwaway email, it still identifies us just as efficiently as if we were to use our primary email. Money is the great motivator. We saw what the ability to extract extortion payment by cryptocurrency did for the ransomware world. It exploded overnight. Websites are now being shown how to make more money by asking their visitors more about themselves, then turning that information over to advertisers. How many are going to see that as a problem? I would venture, very few.

So what was once tracking being done without website assistance is evolving into collusion between websites and their advertisers. I think it's clearly inevitable, and there's nothing we can do about it. As with most things which are abhorrent but invisible, as tracking always has been, most people will have no idea this is going on and I suspect that many wouldn't care anyway.

Kevin van Haaren / @kvanh

> *I'm not sure anyone's mentioned this to you yet but Bitwarden's passkey implementation is available now. I was able to create a passkey for a site on my iPad, go into work and use that passkey from my Windows computer without issues. When I went to add a passkey to the account on the iPad, Bitwarden popped up automatically asking if I wanted to create the passkey in Bitwarden.*

We had heard that support was in beta and coming soon... but I hadn't noticed that Bitwarden's support for mobile was now out of beta. That's great news. (And, as we all know, Bitwarden is a sponsor of the TWiT network and we're very glad they are.)

Robert Harder / @rharder

> *Regarding passkeys, help me out here - I feel like you and Leo are missing the point. Or am I? I thought passkeys were to say, "Hey, this device has already logged in properly so let's make future logins are super easy but also super secure." So that would mean I *don't* want my passkeys to be exportable. If I want to log in on another device or OS or ecosystem at all, I want to prove that's me all over again with whatever way I do that on that website (hopefully with MFA). Only then is that device and that device only secured and proven. It's a nice bonus that Apple or Microsoft or Google have internal synchronizing for their ecosystems but only if it's really, really, really securable. Generally speaking, having passkeys exportable is as bad as the FireSheep days when grabbing someone's session cookie gave an opponent 100% impersonation of a victim. Yes? No? Thanks! Listen from episode one! -Rob*

Okay. So I think Robert makes a valid point. Another entirely different way to think of passkeys is the way he does. In that case, the existing username and password login is used one last time on each device, which then receives a passkey to forevermore authenticate that user and device to that website. I can see that as a workable model – but here’s the critical factor: That model only works in a world where every website allows for any number of passkeys to be registered to any single web account. If, at some point, a website were to reply: “We’re sorry but you’ve reached the limit of passkeys that can be assigned to your account. If you wish to add another, please review and remove some that have already been added.” We don’t know that’s ever going to happen – but we know that it could, and Kevin’s experience of creating a single passkey with Bitwarden on his iPad then having Bitwarden later login for him using the same synchronized passkey under Windows is pretty slick, too.

Spencer Webb / @SpencerWebb

> *Enjoyed the eLORAN discussion. I know the guy at URSA Nav, we had discussions about some projects a few years back. When the USG turned off LORAN, I thought it was incredibly stupid. It *does* work indoors, and in caves, and without an ionosphere. And yes, you can read into the above some interesting scenarios. Remember to feed your antenna. Best,*
> *Spencer*

Spencer is a serious radio guy. We often exchange notes when something about radio comes up. So it was nice to have him add to the eLORAN discussion. I think it’s clear that having a system that’s fundamentally terrestrial has many applications, even when GPS is working well.

Dr Brian of London / @brianoflondon

> *I integrated Passkeys into my own site as a secondary log in system which in some cases is easier to use, especially on mobile devices than the primary method (which is cryptographic signing of a challenge with either a browser plug in wallet or using something with QR codes that looks like SQRL to prove ownership of personal keys). To this I added the ability to associate one or many passkeys with an account and add/delete/rename them. One little gotcha which you probably only learn when implementing this: I store on my server a list of all public Passkeys, and every time I get a log in request from a client, I* ***could*** *send every public key I have and the client would figure out which if any it holds. But in reality I don't do that, I associate each of the public keys with a username (this is part of my primary system anyway*

> *but that username is the ONLY thing I hold, I don't have emails or passwords). I use that username to filter the list of pubkeys I send back to the client which then figures out if the user's device has any of them.*
>
> *It works nicely with Apple passkeys, 1password passkeys (which already sync across multiple devices nicely) and Yubikeys I have for testing.*

Shaun Merrigan / @shaun645D
Remember that it was Shaun with the old LORAN receiver that woke up when eLORAN was turned back on. He followed up with a bit of interesting info:

> *To close the loop on this: My location is Edmonton, Alberta, Canada. The three eLoran stations that are currently testing are Fallon, NV, George, WA and Havre, MT. This is my best information. Currently my old Austron 2100F is showing 2.8E-12 seconds offset from GPS. Thanks! Shaun M.*

So cool! And talk about range!

Markus Daghall / @daghall

> *Hi, Steve! While looking at the PIN heatmap graph, the number 1701 seems to be more prevalent than its surrounding numbers.* 🖖

What a terrific observation! Thanks, Markus!

Ed Ross / @edaross

> *Re "Big Yellow Taxi" - presumably that system helps in situations where "you don't know what you've got Till it's gone"?*

Riny Heijdendael via grc mailbag

> *Location: Spain*
> *Subject: Passkeys: A Shattered Dream - Time to get over Yubikey Limitations?*
> *Date: 01 May 2024 14:55:02*
>
> *As many I started the FIDO 1 journey with Yubikey, but even then I was splattered by the messy software support, implementation guides and it was at that level that I thought it was a no go for regular users : slot selection, HMAC, keyboard emulation, all cool... a bit too cool.*
>
> *But when FIDO2 came along we had to switch tokens anyway, and I switched to "Token2" keys, a Swiss made token that manages selective key removal, up to 300 keys, and enforced PIN complexity, all for a better price than the Yubikey.*
>
> *Furthermore I needed TOTP for 2fa that would work as a standalone device when traveling, and even that is in their portfolio. I just don't understand why Yubikey is still pushed as the*

> *defacto standard. What do you think?*
>
> https://www.token2.swiss/shop/product/token2-t2f2-pin-release2-type-c-fido2-u2f-and-totp-security-key-with-pin-complexity-feature
>
> *Keep up the good work, BTW I silently suspect that you were hired by the UK gov to write the specs for them!*

I needed to let all of our listeners know about these Token2 passkeys dongles... they look fantastic, and supporting 300 passkeys, individually manageable and deletable, with both USB-A and -C connection options, they look fantastic. I'll certainly admit to feeling some proprietary intellectual connection to YubiKey as the guy who happened to come along at the right time and had the perfect audience for them with this podcast. But that's the limit of it. I would like them to succeed in the long term, but that requires them to keep up in what has obviously become a very competitive market. The huge advantage they've been enjoying is having been first. And that's a big deal. But to remain first they need to remain competitive, and we've all been scratching our heads over why they would still have a 25-key limitation when such limitation pretty much relegates them to the enterprise or password manager unlocking role. To be a consumer's primary passkeys container requires that they be able to retain and selectively manage hundreds of keys. So I'll say it again, these Swiss-made Token2 dongles look fantastic.

The bad news is they're in Switzerland, the one we want is currently sold out, and shipping, unless you choose postal mail, which they discourage, is twice the cost of the dongle. And ask me how I know all that. Riny kindly provided a direct link to the Token2 page, which is in the show notes for anyone who wishes to follow.

Also, another listener, Andreas in Germany, also pointed to the Token2 solution which, by the way, is FIDO, FIDO2/WebAuthn, TOTP, USB and NFC. It really does look slick.

312 Scientists & Researchers Respond

Our listener, **Robin van Zon in the Netherlands**, brought this recently produced letter to my attention. Thank you, Robin.

The letter opens by introducing itself: *"The text below is an open letter on the position of scientists and researchers on the recently proposed* ***changes*** *to the EU's proposed Child Sexual Abuse Regulation. As of the 7th May 2024 – so exactly two weeks ago, today – the letter has been signed by 312 scientists and researchers from 35 countries."*

It turns out that what scientists and researchers have to say is quite refreshing because it actually engages science, math, statistics and... yes, reality... as opposed to the politicians' statements of "this is what we want and what we're preparing to demand."

So I want to share what these 312 scientists and researchers collectively wrote and signed, because it's not overly long, because the devil, it turns out, **is** in the details, and because there is probably no more important issue on the table at this moment in time than what the EU's political class will finally decide to try to do. And importantly, as we'll see, this is the technical response to the politician's responses to the previous technical response. What's heartening to see is that both sides, so far, appear to be negotiating in good faith. The politicians are, at least, listening.

As we know, when the UK was faced with serious opposition to their proposal to require all private conversations to be monitored for content, they wisely added the caveat "where this can be proven to be technically feasible without compromising security" which allowed the politicians to save face, enact their legislation, and all messaging providers to continue offering fully private end-to-end encryption because it hasn't been and cannot be proven to be feasible without compromising security.

But the European Union is not there yet. So here's the latest feedback from the EU's technical experts which is intended to inform the politicians of reality. The undersigned, wrote:

> *We are writing in response to the new proposal for the regulation introduced by the Presidency on 13 March 20241. The two main changes with respect to the previous proposal aim to generate more targeted detection orders, and to protect cybersecurity and encrypted data. We note with disappointment that these changes fail to address the main concerns raised in our open letter from July 2023 regarding the unavoidable flaws of detection techniques and the significant weakening of the protection that is inherent to adding detection capabilities to end-to-end encrypted communications. The proposal's impact on end-to-end encryption is in direct contradiction to the intent of the European Court of Human Rights's decision in Podchasov v. Russia on 13 February, 2024. We elaborate on these aspects below.*

I tracked down that decision. The case surrounded Russia's FSB demanding that Telegram turn over the decrypted communications of six individuals who the FSB alleges were involved in terrorism against the Russian state. Telegram refused, explaining that since all of the subjects involved had enabled Telegram's optional end-to-end encryption mode, Telegram's default ability to store unencrypted conversation data in their servers was thwarted. And, indeed, paragraphs 79 and 80 of the decision of the European Court of Human Rights reads:

> *79. The Court concludes that in the present case the ICO's statutory obligation to decrypt end-to-end encrypted communications risks amounting to a requirement that providers of such services weaken the encryption mechanism for all users; it is accordingly not proportionate to the legitimate aims pursued.*
>
> *80. The Court concludes from the foregoing that the contested legislation providing for the retention of all Internet communications of all users, the security services' direct access to the data stored without adequate safeguards against abuse and the requirement to decrypt encrypted communications, as applied to end-to-end encrypted communications, cannot be regarded as necessary in a democratic society.*
>
> *In so far as this legislation permits the public authorities to have access, on a generalized basis and without sufficient safeguards, to the content of electronic communications, it impairs the very essence of the right to respect for private life under Article 8 of the Convention. The respondent State has therefore overstepped any acceptable margin of appreciation in this regard.*

So what this tells us is that separate from whatever political pressures the EU's politicians may be under, when the issues at stake are very carefully and thoroughly examined by the European courts, their decisions never support the application of wholesale surveillance. For the sake of our listener's sanity, I skipped over the first 78 paragraphs. But those paragraphs make it very clear that the courts really do very clearly understand the issues. They clearly understand that the phrase "selective backdoors" is an oxymoron.

Okay. So continuing with the technologists' latest rebuttal response to the politicians' attempt to mollify them following their first surveillance proposal, they all wrote and signed:

> *Child sexual abuse and exploitation are serious crimes that can cause lifelong harm to survivors; certainly it is essential that governments, service providers, and society at large take major responsibility in tackling these crimes. The fact that the new proposal encourages service providers to employ a swift and robust process for notifying potential victims is a useful step forward.*
>
> *However, from a technical standpoint, to be effective, this new proposal will also completely undermine communications and systems security. The proposal notably still fails to take into account decades of effort by researchers, industry, and policy makers to protect communications. Instead of starting a dialogue with academic experts and making data available on detection technologies and their alleged effectiveness, the proposal creates unprecedented capabilities for surveillance and control of Internet users. This undermines a secure digital future for our society and can have enormous consequences for democratic processes in Europe and beyond.*

> ***1. The proposed targeted detection measures will not reduce risks of massive surveillance***
>
> *The problem is that flawed detection technology cannot be relied upon to determine cases of interest. We previously detailed security issues associated with the technologies that can be used to implement detection of known and new CSA material and of grooming, because they are easy to circumvent by those who want to bypass detection, and they are prone to errors in classification. The latter point is highly relevant for the new proposal, which aims to reduce impact by only reporting "users of interest" defined as those who are flagged repeatedly (as of the last draft: twice for* ***known*** *CSA material and three times for new CSA material and grooming). Yet, this measure is unlikely to address the problems we raised.*
>
> *First, there is the poor performance of automated detection technologies for new CSA material and for the detection of grooming. The number of false positives due to detection errors is highly unlikely to be significantly reduced unless the number of repetitions is so large that the detection stops being effective. Given the large amount of messages sent in these platforms (in the order of billions), one can expect a very large amount of false alarms (in the order of millions). Note 2*

At this point there's a footnote which serves to support this contention. It says:

2 Given that there has not been any public information on the performance of the detectors that could be used in practice, let us imagine we would have a detector for CSAM and grooming, as stated in the proposal, with just a 0.1% False Positive rate (i.e., one in a thousand times, it incorrectly classifies non-CSAM as CSAM), which is ***much*** *lower than any currently known detector. Given that WhatsApp users send 140 billion messages per day, even if only 1 in one hundred would be a message tested by such detectors, there would be 1.4 million false positives every single day. To get the false positives down to the hundreds, statistically one would have to identify at least 5 repetitions using different, statistically independent images or detectors. And this is only for Whatsapp - if we consider other messaging platforms, including email, the number of necessary repetitions would grow significantly to the point of not effectively reducing the CSAM sharing capabilities.*

> *Second, the belief that the number of false positives will be reduced significantly by requiring a small number of repetitions relies on the fallacy that for innocent users two positive detection events are independent and that the corresponding error probabilities can be multiplied. In practice, communications exist in a specific context (e.g., photos to doctors, legitimate sharing across family and friends). In such cases, it is likely that parents will send more than one photo to doctors, and families will share more than one photo of their vacations at the beach or pool, thus increasing the number of false positives for this person. It is therefore unclear that this measure makes any effective difference with respect to the previous proposal.*

In other words, the politicians proposed to minimize false positive detections by requiring multiple detections for a single individual before an alarm was raised. But the science says that won't work because entirely innocent photographs of one's children will not be evenly distributed across the population of all communicating users. People who have young families and like to share photos of their children frolicking at the beach in their bathing suits will generate massive levels of false positive CSAM detections.

The scientists explained:

> *Furthermore, to realize this new measure, on-device detection with so-called client-side scanning will be needed. As we previously wrote, once such a capability is in place, there is little possibility of controlling what is being detected and which threshold is used on the device for such detections to be considered "of interest". We elaborate below.*

I should explain that another amendment to the proposed legislation involves attempting to divide applications into high-risk and low-risk categories so that only those deemed to be high-risk would be subjected to surveillance. The techies explain why this won't work. They write:

> *High-risk applications may still indiscriminately affect a massive number of people. A second change in the proposal is to only require detection on (parts of) services that are deemed to be high-risk in terms of carrying CSA material. This change is unlikely to have a useful impact. As the exchange of CSA material or grooming only requires standard features that are widely supported by many service providers (such as exchanging chat messages and images), this will undoubtedly impact many services. Moreover, an increasing number of services deploy end-to-end encryption, greatly enhancing user privacy and security, which will increase the likelihood that these services will be categorized as high risk. This number may further increase with the interoperability requirements introduced by the Digital Markets Act that will result in messages flowing between low-risk and high-risk services. As a result, almost all services could be classified as high risk.*
>
> *This change is also unlikely to impact abusers. As soon as abusers become aware that a service provider has activated client side scanning, they will switch to another provider that will in turn become high risk; very quickly all services will be high risk, which defeats the purpose of identifying high risk services in the first place. And because open-source chat systems are currently easy to deploy, groups of offenders can easily set up their own service without any CSAM detection capabilities.*
>
> *We note that decreasing the number of services is not even the crucial issue, as this change would not necessarily reduce the number of (innocent) users that would be subject to detection capabilities. This is because many of the main applications targeted by this regulation, such as email, messaging, and file sharing are used by hundreds of millions of users (or even billions in the case of WhatsApp).*
>
> *Once a detection capability is deployed by the service, it is not technologically possible to limit its application to a subset of the users. Either it exists in all the deployed copies of the application, or it does not. Otherwise, potential abusers could easily find out if they have a version different from the majority population and therefore if they have been targeted. Therefore, upon implementation, the envisioned limitations associated with risk categorization do not necessarily result in better user discrimination or targeting, but in essence have the same effect for users as a blanket detection regulation.*

2. Detection in end-to-end encrypted services by definition undermines encryption

protection

The new proposal has as one of its goals to "protect cyber security and encrypted data, while keeping services using end-to-end encryption within the scope of detection orders". As we have explained before, this is an oxymoron. The protection given by end-to-end encryption implies that no one other than the intended recipient of a communication should be able to learn any information about the content of such communication. Enabling detection capabilities, whether for encrypted data or for data before it is encrypted, violates the very definition of confidentiality provided by end-to-end encryption. Moreover, the proposal also states that "This Regulation shall not create any obligation that would require [a service provider] to decrypt or create access to end-to-end-encrypted data, or that would prevent the provision of end-to-end encrypted services." This can be misleading, as whether the obligation to decrypt exists or not, the proposal undermines the protection provided by end-to-end encryption.

This has catastrophic consequences. It sets a precedent for filtering the Internet, and prevents people from using some of the few tools available to protect their right to a private life in the digital space; it will have a chilling effect, in particular to teenagers who heavily rely on online services for their interactions. It will change how digital services are used around the world and is likely to negatively affect democracies across the globe. These consequences come from the very existence of detection capabilities, and thus cannot be addressed by either reducing the scope of detection in terms of applications or target users: once they exist, all users are in danger. Hence, the requirement of Article 10(aa) which states that "a detection order should not introduce cybersecurity risks for which it is not possible to take any effective measures to mitigate such risk" is not realistic, as the risk introduced by client side scanning cannot be mitigated effectively.

3. Introducing more immature technologies may increase the risk

The proposal states that age verification and age assessment measures will be taken, creating a need to prove age in services that before did not require so. It then bases some of the arguments related to the protection of children on the assumption that such measures will be effective. We would like to point out that at this time there is no established, well-proven technological solution that can reliably perform these assessments. The proposal also states that such verification and assessment should preserve privacy. We note that this is a very hard problem. While there is research towards technologies that could assist in implementing privacy-preserving age verification, none of them are currently in the market. Integrating them into systems in a secure way is far from trivial. Any solutions to this problem need to be very carefully scrutinized to ensure that the new assessments do not result in privacy harms or discrimination causing more harm than the one they were meant to prevent.

4. Lack of transparency

It is quite regretful that the proposers failed to reach out to security and privacy experts to understand what is feasible before putting forth a new proposal that cannot work technologically. The proposal pays insufficient attention to the technical risks and imposes - while claiming to be technologically neutral - requirements that cannot be met by any state-of-the-art system (e.g., low false-positive rate, secrecy of the parameters and algorithms when deployed in a large number of devices, existence of representative simulated CSA material).

> *We strongly recommend that not only should this proposal not move forward, but that before such a proposal is presented in the future, the proposers engage in serious conversations about what can and cannot be done within the context of guaranteeing secure communications for society.*

And they conclude with their 5th point which, despite the fact that it tends to be overlooked and forgotten, is so important...

> ***5. Secure paths forward for child protection***
>
> *Protecting children from online abuse, while preserving their right to secure communications, is critical. It is important to remember that CSAM content is the output of child sexual abuse. Eradicating CSAM relies on eradicating abuse, not only abuse material.*
>
> *Proven approaches recommended by organizations such as the UN for eradicating abuse include education on consent, on norms and values, on digital literacy and online safety, and comprehensive sex education; trauma-sensitive reporting hotlines; and keyword-search based interventions. Educational efforts can take place in partnership with platforms, which can prioritize high quality educational results in search or collaborate with their content creators to develop engaging resources.*
>
> *We recommend substantial increases in investment and effort to support existing proven approaches to eradicate abuse, and with it, abusive material. Such approaches stand in contrast to the current techno-solutionist proposal, which is focused on vacuuming up abusive material from the internet at the cost of communication security, with little potential for impact on abuse perpetrated against children.*

In other words, you politicians are aiming at the wrong target anyway. So even if you got everything you want by effectively eliminating security and all privacy, it won't actually solve the problem that you're hoping to solve.

I think the problem is that it's like an iceberg. CSAM is the tip of the iceberg that is the visible manifestation of something that's abhorrent. And because we see it, the tip of that iceberg, we want to get rid of it. But these authors remind us that CSAM is the output, it's the result of these abhorrent practices — less so the practices themselves.

What I'm heartened by, as I said at the top, is that we appear to be seeing an honest back and forth negotiation in good faith between European Union politicians and European scientists and researchers. Given that the original proposed legislation was significantly amended after their first round of objections and feedback, it appears that the politicians are heeding what their technocrats are explaining.

We have no idea what's going to finally happen, which is what makes all this so interesting. Stay tuned.