312 SCIENTISTS & RESEARCHERS RESPOND

Description: 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.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-975.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-975-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. Of course, as always, there's a ton to talk about. We will in just a little bit talk about the scientific response to the European CSAM proposals, those proposals that break Internet encryption. We'll also talk a little bit about three new zero-days in Google Chrome, what happened to Google Search, and why AI is not the answer. And just how long can an unprotected XP machine live on the Internet? The answer will surprise you. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 975, recorded Tuesday, May 21st, 2024: 312 Scientists & Researchers Respond.

It's time for Security Now!, the show where we cover the latest security news, privacy information, talk about hacks and hackers and, you know, just kind of shoot the breeze with this super intelligent human being we call Steve Gibson. Hi, Steve.

Steve Gibson: And continue to justify our existence, apparently.

Leo: Without Security Now! there's no justification, yeah.

Steve: So, okay. We're going to have some fun this week, not that we don't always. We're going to examine which browser has had a very rough week, and why. Which bodily fluid should you probably not drink despite Google's recommendation to the contrary?

Leo: Eww. Okay.

Steve: I know. It's freaky. And also, how can you tweak your browser so that you will be avoiding those recommendations in the future? What happens when a Windows XP machine is exposed to the unfiltered Internet? Duck and cover comes to mind.

Leo: Oh, boy.

Steve: How did a pair of college kids get their laundry washed for free? And what do we learn about that from the still-clueless corporations which clearly exist? And finally, after engaging with some terrific listener feedback that we have this week, we're going to examine the latest actually quite thought-provoking response to the EU's proposed Child Sexual Abuse Regulation, this time from their own scientific and research community. Thus the title of this podcast: "312 Scientists & Researchers Respond."

Leo: Awesome. And four out of five doctors agree that this is the only thing you should be doing on a Tuesday afternoon; all right?

Steve: That's right.

Leo: All right.

Steve: To justify your existence.

Leo: We will get to all that. I can't wait to hear the washing machine story. I saw the headlines on that, and I thought, oh, Steve has got to cover this.

Steve: Yeah, yeah, yeah, too fun.

Leo: Very interesting. Very, very interesting. And now I am prepared for the Picture of the Week.

Steve: So I gave this picture just a simple caption: "Uh, what?"

Leo: It says it all, really. You don't - you don't need...

Steve: Exactly. You know, I was tempted to give it the caption, "I don't think that means what you think it means." And this is another one of those, Leo, where you've just got to ask yourself, you know, somebody produced this. Somebody, like, created a plate to put on a door which you can only push in order to open the door. Yet prominently displayed at the top, beautifully engraved and then color filled, you know, etched in this plate it says "Pull."

Leo: I think it's for Jedi warriors to, like, practice the Force. Pull.

Steve: That would be good. Yes, exactly. You just, you know, work on your telekinesis.

Leo: Unbelievable. Wow.

Steve: Anyway, I just got just a kick out of that. Thanks to our amazing listeners. They find these things and send them to us so I get to share them with everyone.

Okay. So Google's much-beloved Chrome browser has had a very rough week. In just one week, the total number of exploited, in-the-wild, zero-day vulnerabilities to be patched so far this year jumped from four to seven.

Leo: Wow.

Steve: In other words, last week saw three newly discovered Chrome vulnerabilities receiving emergency Chrome patches. In 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, who said they were looking into it, and they were going to, like, work on fixing this thing, too, because of course Microsoft is also using the common Chromium engine.

So this latest trouble is rated as a high-severity zero-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. Now, these so-called "Type Confusion" bugs, we see them are arising often. They're more formally referred to as "Access of Resource Using Incompatible Type," which sort of says the same thing. This occurs when code misinterprets data types which can lead to unpredictable behavior - which is putting it mildly - that can allow attackers to manipulate program logic or access sensitive information.

We've talked about before how the values stored in a computer's registers or in memory might either be the actual data itself, or often can be a pointer to some other data. The use and manipulation of pointers is, wow, I mean, it's very powerful, but also very dangerous because the pointer can potentially point to anything. 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 for example the length of the data they've just sent, could be mistakenly treated by some buggy code as a pointer, in which case the attacker can control what the pointer points to and thus increase the amount of mischief that they're able to get themselves into. In theory, that would allow an attacker, for example, to do exactly the sort of things that we see happening.

So 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 that they're now fixed. So, you know, they say update your Chrome and you'll be okay. And, you know, all they say is "Access to bug details and links may be kept restricted until a majority of users are updated with a fix." And of course 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 since been fixed in Chrome. So they sort of say, oh, we're not going to tell you until later, and later never comes.

However, you were just talking about the Thinkst Canary. And this article talks about this because, or this event, because one thing that comes 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's end-point security solutions, and

those solutions are feeding the intelligence that they collect back to Kaspersky's mothership for monitoring and interpretation. So when one of Kaspersky's customers is targeted, red flags go up at Kaspersky central.

Okay, now, as I said, there were three this past week. The other two actively exploited Chrome zero-days patched this week are 4671 and 4761, which also double 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, once again, the V8 JavaScript engine. And it's worth noting that four out of the seven zero-day bugs Chrome has patched so far this year have all been located in Chrome's V8 JavaScript engine. This is not necessarily the JIT, the Just-In-Time compiler portion.

But recall that the observation has been previously made that the overwhelming majority through time of bugs in the common Chromium core were being found in V8's - in the JIT, the Just-In-Time compiler portion of 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, especially given how much faster our processors are today than when this was first implemented, back when they really did need all the speed they could get. Now it's like, well, you know, the processors are sitting around doing nothing most of the time anyway. So how about trading off some speed in return for cutting serious vulnerabilities by more than half.

Toward the end of last month, Microsoft explained the so-called "Enhanced Security for Edge" setting that they have in their browser. They wrote: "Microsoft Edge is adding enhanced security protections to provide an extra layer of protection when browsing the web and visiting unfamiliar sites." That's the keyword. "The web platform," they wrote, "is designed to give you a rich browsing experience," blah blah blah, "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."

They wrote: "Enhanced security mode in Microsoft Edge mitigates memory-related vulnerabilities by disabling Just-In-Time JavaScript compilation and enabling additional operating system protections for the browser. These protections include Hardware-enforced Stack Protection and Arbitrary Code Guard. 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 being selectively enabled when visiting unfamiliar sites. But this allows Edge running on, for example, Outlook 365 or Google properties to race ahead at full speed with those extra protective guards disabled. And given Chrome's past week of three newly exploited in-the-wild zero-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.

Okay. So this next piece, the fact that Leo has been driven to a paid search solution I think says important things about what has happened to search. We're going to see some additional evidence of that. One of the things I most loved about the early Google search was its search results' cleanliness and simplicity. They were remarkable, not only because they were relevant. I mean, it was astonishing back then. Anyway, 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 hoped to be able to announce this week that the subscription management

frontend was ready for the world. But it needs some additional testing, so that'll be next week's announcement. I wrote GRC's first email system back in the late 1990s, and it sent over the course of its life a grand total of 11 mailings. To my surprise, last week I stumbled upon the archive of those 11 emailings. And the second one, dated April 2nd - not April Fool's Day, fortunately - April 2nd, 1999 had the subject "Steve Gibson's News of a Stunning New Search Engine."

Okay. So the email that I sent to GRC's subscribers a little over 25 years ago read: "We've all experienced the problem. The automated search engines like Alta Vista return 54,321 items 'in no particular order'" - it actually used to say that - "many of which unfortunately were porn sites. But the human-indexed search services of the time like Yahoo," I wrote, "often cannot find what you want because they're only able to index a small fraction of the entire web since they're being indexed by people. So you're left with the uneasy, but probably accurate sense," I wrote, "that what you want is out there somewhere, but you're no closer to finding it."

And then I said: "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." And then I provided a URL that at that time no one had ever seen: http: - no "s" back then - http://google.com. And I wrote: "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."

I wrote, "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 I said: "And of course it's free, so give it a shot yourself." And then my email ended with a link to, again, Google.com, which 25 years ago when I sent this mail on April 2nd of 1999...

Leo: That's pretty impressive.

Steve: ...no one had ever heard of.

Leo: That's great.

Steve: So I just, I thought, I got such a kick out of that. So, you know, 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 really, you know, it was 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 oh, how times have changed. As I said at the start of this, the fact that you, Leo, have been driven to a paid search solution says some important things.

My own personal annoyance is that I never, I mean, literally, I never want to watch a video to receive an answer to whatever question I might have put into search. Yet Google promotes videos to the top of their search results, not because they provide the best answer, but because Google owns YouTube now.

Leo: Exactly, yeah.

Steve: Yeah. I'm writing my forthcoming email system's subscription management frontend because I'm very picky about exactly how I want it to work and how I insist that GRC treats its 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 backend 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.

Leo: Oh, wow.

Steve: And this was not because, as I originally wrote 25 years ago, those four alternative solutions are objectively better, but because they're paying Google to appear first.

Leo: Right. They're ads, yeah.

Steve: That's right. 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 the new 'Web' search mode, which we explain below. AI Overviews, also known as 'Search Generative Experience'" - and I might change it to Degenerative, but we'll get to that later - "is Google's new search feature," they wrote, "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 AI will rewrite content from independent websites and summarize it in its own words." They said: "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."

Leo: What?

Steve: I have a snapshot of the tweet from May 5th.

Leo: What?

Steve: It reads - it shows the person...

Leo: Every 24 hours.

Steve: Yes. How to pass kidney stones quickly. And the answer is "drinking plenty of fluids, such as water, ginger ale, lemon-lime soda, or fruit juice can help pass kidney stones more quickly." Next sentence: "You should aim to drink at least two quarters" - and it has helpfully in parens - "(two liters) of urine every 24 hours, and your urine should be light in color."

Leo: Okay, thank you, AI. Holy cow.

Steve: Thank you so much. Now, this was May 5th. And it noted that AI Overviews are experimental. This was the week before this was formally released. And I just so loved the comment of the guy who posted this who asked the question. He wrote in response to this: "Perfect. Ready to go. Ship it out."

Leo: Oh, my god.

Steve: I know. Wow. So BleepingComputer said: "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'" - well, yeah, because they don't get the answer the first time, they've got to use it some more. I don't think I want to drink pee, thank you very much. You got any other ideas?

So BleepingComputer continues: "That doesn't even seem to be the case on many Google support forums, that is, where people are more satisfied with their results. For example: '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 over on Google forums: 'Every single result I've received from the AI Overviews has been incorrect. I'm more capable of misinterpreting Internet articles on my own.'"

Leo: I don't need any help, thank you.

Steve: Don't need any AI to help me.

Leo: Wow.

Steve: "'And I can probably get at least slightly closer to actual understanding than the AI because I actually have cognitive processes.'" So BleepingComputer wrote: "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's no way to disable it now that it is out of Search Labs, and Google has quickly locked support threads for many people asking how to do so." So, whoops, you're not even allowed to ask anymore.

Leo: Wow.

Steve: And they said: "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" - get this, Leo - "a new 'Web' search option to return the old search feel." Wait. What? I thought that Google was web search.

Leo: Yeah, what are they talking about? What?

Steve: Right. 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, Google, no longer consider it to actually be web search. Okay. Google's search results list a series of search result "filters" in a line underneath the search field. They typically read All, then 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 now 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 then joined with an ampersand to the "q=" clause, causes the search to default to "Web," and you get much cleaner 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. At least, yeah, the AI won't.

So I have not encountered this default web search trick anywhere else, so I've placed a link to BleepingComputer's write-up in the show notes. And for ease of access I've also made this GRC's shortcut for the week. So for this podcast 975, it's grc.sc/975 will take you to this article at BleepingComputer if you want to, like, see them explaining how to do this. And you do need then to get into your browser and tweak, like to add a custom search technique and then select that as the default. But if you want to keep using Google, and you want to return to simpler times, you are able to get the Web search results. And you can quickly see for yourself how much better it looks just by going under those three dots and More and selecting Web, and you get a cleaned-up page.

So I'll just say, as an aside, what a mess. You know, 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 there are many people who 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. Wow.

Leo: I think there's actually a story behind all this, and we'll talk about it tomorrow on TWiG. It's unknown, obviously. Google's not talking. But a number of people from Google have said, you know, Google has just panicked. In the same way they panicked when they thought they were going to lose to Meta and created Google Plus, they panicked, thinking they were going to lose the farm to AI. And of course panic is not a good way to create new features. And it feels like that. They're just throwing stuff up against the wall at this point.

Steve: Especially something this complicated. Leo, this stuff is so, I mean, you know, arguably most people don't know how AI works. Right? It's like, well, it started sounding conscious. So...

Leo: I don't know how it works myself. It's kind of a mystery.

Steve: Exactly. I mean, the AI researchers are like...

Leo: They don't know, yeah.

Steve: ...what?

Leo: Yeah. Anyway, as you say, I don't use Google anymore, so I have missed this entire drama. You know, for a while I did use something called Neeva that had AI summaries at the beginning. But they were always footnoted, and I never found them to be wrong. It was done cleverly. But what I use now, Kagi, K-A-G-I, which is a paid search tool, doesn't do that. I think people have realized that these AI summaries are not that useful. And just give me the results.

Steve: Well, and think of the good that could be done, Leo, if we ever got to a point where all of the chaff could be separated, and instead of getting nonsense from an Internet search, no matter who does it, you actually got rigorous truth. That would be something.

Leo: Well, but let's also be fair to Google, part of this is because the Internet is full of crap.

Steve: Exactly.

Leo: I mean, the Internet...

Steve: So if you're training your AI on crap...

Leo: Right.

Steve: ...you know, crap in, crap out.

Leo: And even Google Search can only reflect the search contents. And if it's garbage, everybody's trying to game Google, it's going to reduce the search results anyway. So it's a mess right now. It's just a mess. All right. We're not alone as long as we've got Mr. G. here, help us out with the world at large. Steve?

Steve: So 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. Using a virtual machine, Parker set up a Windows XP instance and configured it to be fully exposed with no firewall and no antivirus software, just like the good old days."

Okay, now, just to remind everybody, even though XP always had a built-in firewall, it wasn't until Service Pack 3 that the firewall was enabled by default, like by the installation of that service pack, or installing XP that included SP3 after that point. And of course, thanks to the tyranny of the default, very few earlier pre-Service Pack 3 Windows XP machines were protected out of the box. And I remember those days. Remember that there was a big third-party market for firewalls. ZoneAlarm was the one that I found and liked a lot based on the way it operated. So Microsoft wanted to add a firewall to their Windows client platform, but they also didn't want to, you know, blatantly they'd already had a lot of problems with antitrust. They didn't want to just go obsoleting a whole class of software immediately. So they put it in, but they didn't turn it on.

Okay. So PC Gamer continues: "So how long exactly does it take for malicious software to appear on the PC? Parker returns to his PC" - you know, his virtual PC - "10 minutes later and, sure enough, there's something nasty running in Task Manager called conhoz [.]exe (C-O-N-H-O-Z dot 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" - you think? - "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" - and of course an FTP server - "the machine is already a complete and utter zombie."

And they said: "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." He says: "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," they write, "it just goes to show how effective those machines are that we can all be connected to the Internet now 24/7 and not give this stuff much thought. It's dangerous out there, boys and girls. Be careful," they conclude.

Okay, now, I would - I love that. I would edit that just a bit to observe that this vividly shows what's right now pounding away at the outside of our stateful NAT routers, those vital pieces of hardware all of our networks are blessedly 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 increasingly wild and wooly Internet, that have made it possible to use this crazy global network with any hope whatsoever of remaining safe. Presumably, I don't know what Eric's history is, but presumably the IP that his XP machine appeared on wasn't for any reason particularly high-profile; right? It was just some random guy.

And there's just that much crap hitting each of our IPs regularly enough that - and who knows? Was it all the same attacker who said, oh, my goodness, we just found a new victim, you know, let's get it? Or different attackers who were all randomly scanning the Internet and happened to lock onto this XP machine. I mean, I have to say I'm tempted to do this because this would sound like fun. Except you've got to be so careful. And it would be so easy to make a mistake. So, you know, if you do want to replicate what Eric did, then, you know, really, really, really be careful. For anyone who's curious to see Eric Parker's YouTube video described in this article, I posted the link in the show notes so it's easy to find. And also when I was there looking at it, I noticed that since then, and this got a lot of attention, he's done the same thing to Windows 2000. So I didn't make time to dig in and see if 2000 was similarly vulnerable.

Okay. TechCrunch reported that, thanks to the discovery made by a pair of curious students at the University of California at Santa Cruz, who to their credit did try to do the right thing by attempting to report the flaws they'd uncovered in the control software of 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."

Okay. So the company behind these widely deployed machines is called "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, yet 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 last 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. Since students will be students, experimenting further, they set one of their accounts to reflect a positive balance of several million dollars in 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 and Canada. Oh, and also Europe. You would 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, months later, despite 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."

Okay. So it seems to me that what might finally arouse CSC's attention - and apparently the only thing that will - may be a sharp and sudden drop in cash flow revenue as word of this spreads across college campuses in the U.S., Canada, and Europe. It's 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 discovery and 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 - get this - are only performed by the client app on the user's device.

Leo: [Laughing]

Steve: I know. And anything sent to CSC's servers are fully trusted.

Leo: Oh, my.

Steve: This allows...

Leo: Including things like I have a million dollars in my account.

Steve: Correct.

Leo: Unbelievable.

Steve: Correct. This allows fake payments to be posted to their accounts without ever putting any real world funds in the accounts. And Leo, it's worse. 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 - get this - 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. 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, four months, five months later, for users to freely give themselves any amount of money. Iakov said that he was disappointed that CSC did not acknowledge their vulnerability. He said: "I just don't get how a company that large makes those types of mistakes, then has no way of contacting them." He said: "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. You know? 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 about 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 does stand to reason that those who are not actively involved in this arena might not still be up to speed on today's etiquette.

Leo: Hey, we're laundry guys. What do we know about the Internet? You put in a quarter, you wash your laundry. This just shows you, though, it's really good for students to have to do their own laundry because that enforced period of boredom can really lead to some creative results. That and a lack of quarters. I love it.

Steve: Yeah. We actually, Leo, I have to confess, back, you know, when I was myself at Berkeley, we had coin-op washing machines.

Leo: Did you tie a string to the quarter and pull it back?

Steve: Actually there was a screw hole in the back of the quarter-accepting add-on to the washing machine. And it didn't take long before - and I'm not saying who.

Leo: Yeah, someone.

Steve: An enterprising student figured out that if you took a coat hanger - I've talked about how handy that coat hanger...

Leo: So useful.

Steve: How coat hangers are like, they are the perfect type of stiff wire. You could cut off a length and put a little hook in the end, snake it through the hole, and then you could reach in, hoping not to be electrocuted, by the way, and grab, you know, find the arm that gets pulled when the quarter is put on the little slider and pushed in, and give it a few tugs, and what do you know? The washing machine would start right up.

Leo: Unbelievable.

Steve: It is a good thing I'm not in college, you know, in this day and age.

Leo: You know, this, honestly, there's a subtext here that maybe CSC goes, yeah, so the students are ripping us off. What else is new? As long as we get most of the money, we're fine. What else is new?

Steve: Well, yeah. And exactly as you said, they're in the washing machine business; right? They probably contracted out to the lowest bidder...

Leo: Of course they did, mm-hmm.

Steve: ...to make themselves an app to put these machines on the Internet. And that guy is gone.

Leo: Long gone.

Steve: So they probably have no idea what to do. You know, they're able to monitor account balances and zero them when they get set to a million dollars. But other than that, eh.

Leo: So we had an intern for a while back in the Brick House days, the wonderful Jeff Needles. You may remember him.

Steve: I do remember him, yeah.

Leo: He was a fun guy. And he said, "I'll write a sales system for you." And so he wrote up a whole sales database system for us that we started using for every ad sale and so forth. And then he left. He got a better job. I think he went to that video company that he was so enamored of. And we said, well, can we give you a contract to maintain? He said, nah, you know. And he just left. And so we had this blob of software which breaks, well, let me put it this way. If two people try to use it at the same time, boom, it's dead. So it breaks all the time. But it's not really worth it for us to redo it. We just hired some guy who kind of looks at the code and pokes at it once in a while. And we just limp along with it. I think a lot of companies are like that. I don't think that's at all unusual. I mean, it's not a security issue because it's not public-facing in any way.

Steve: Right. When I released SpinRite 6.1, we made a decision that we would no longer maintain upgrades of SpinRite from before 6.0 because it's been 20 years.

Leo: Right. Right.

Steve: And, you know, we've been more than generous for decades. Anyway, Sue was greatly relieved because she was using what we called Dino, which actually was a dinosaur. It was the original GRC Novell NetWare database that was written in FoxPro by my second employee, who was a truly gifted coder. He went into writing gaming stuff. His name is Steve Rank, and neat guy. It's running today.

Leo: Yeah.

Steve: And so it was only a few weeks ago that I said to Sue, "Sue, you no longer" - and the reason, the point was that if someone said, hey, I bought SpinRite 3.1 in '98, I think it was, and so she would look it up in Dino and say, oh, sure enough, here you are.

Leo: Yeah. Well, you know, we just - the way it works here is when somebody's going to go use the sales system, they send out a company-wide beacon on Slack, says everybody out. It's funny, I mean, I think every company has something like that.

Steve: Yup.

Leo: It's just normal.

Steve: Yup.

Leo: It's the way it is. And technology moves so fast.

Steve: It's what kept IE from ever dying.

Leo: Exactly.

Steve: Is so many enterprises had written internal stuff that was dependent upon specific quirks in operation of Internet Explorer, that there was like, no, no, no, you can't take our IE away from us.

Leo: Now, back to the fun and games with Stevie Gibson, Little Stevie Gibson.

Steve: Okay. So, many of our listeners forwarded tweets from 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." So he's saying commands were being sent on Friday to switch from just sending data back as they have been to actually switching over to sending science back. Meaning the output of the Voyager 1's surviving sensor arrays.

Then Sunday on the 19th, two days ago, 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 in six months." So, yes, incredibly, I mean, it really is incredible.

Leo: Yeah, no kidding.

Steve: 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. And, wow.

Leo: Brilliant.

Steve: And also to our listeners, thank you, all you who tweeted that, making sure that I knew. Okay. Hakku. "Hello, Steve. Long-time follower and big fan of SN. Keep up the great work. One question following SN-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 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?" He says: "We're all 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."

Leo: Nice.

Steve: "To 999 and beyond." Okay. So 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. And I remember they mentioned OpenVPN. One of the problems is OpenVPN is open source and cross-platform. So what is cross-platform is using the routing table to manage rerouting. But if you're running OpenVPN on a Windows machine, the local firewall aspect is not cross-platform. So that's not something, I mean, whereas other platforms may also have local firewalls, they've all got their own. And Windows is certainly not compatible with anything over in the Linux world, or Mac.

So what they found was that many popular VPNs in widespread use today were true victims of the attack which 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. So let's hope that the popular VPN providers are being asked about their susceptibility to this particular form of simple routing table attack and then do take the effort to revise their solutions, if necessary. And I'm glad, for example, that this came up and that Hakku brought it up to his tech guys and that they're going to do some research to make sure that they're not vulnerable.

214normandy wrote: "Hi, Steve. I know you've been using" - oh - "the Netgate SG-1100 as well as the four-port Protectli Vault. I'm starting to see reports that the eMMC in the SG-1100 is starting to wear out for folks." He said: "I ran their suggested commands" - and he provides a link to their documentation - "to check the eMMC, and it says that my eMMC is end of life expected already." He says: "No big deal, I'll move on and try the four-port Protectli Vault instead. Hoping you can confirm that you are still happy with your Protectli. Thanks, Bob."

Okay. So this came as news to me, so I wanted to share it for any other Netgate SG-1100 users who may have followed my advice and my choice about that beautiful little Netgate appliance. 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 is currently churning away in the pfSense firewall. It's constantly writing logs to the file system, and eMMC memories do not have huge amounts of excess endurance. You know, they're meant more for embedded solutions that are not churning constantly. I still have a trusty SG-1100. I mean, right now the bandwidth for this podcast is passing across my SG-1100. And it's been giving me no trouble ever since I replaced its power supply. Remember that it was glitching, and it turned out to be the

power supply that was the problem. But it is sobering that it will have a lifetime limit due to the failure of an eMMC memory that cannot be replaced.

Bob also asked about my other favorite pfSense hosting device, which is the four-port Protectli Vault. That's what's running pfSense at my place with Lorrie. And yes, I'm still utterly happy with that choice, too. And in fact I have another of those standing by ready for deployment here if the SG-1100 should ever die, which unfortunately no longer seems as unlikely as it once did. You know, just giving it power apparently won't be enough in the long run, which is quite disappointing.

Okay. I have an important and interesting message from a listener who requested anonymity. He wrote: "Hello, Steve. I've been a listener of Security Now! for years, perhaps even a decade; a member of 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" - okay. "Managing a website with a significant audience of approximately one million visitors per day. Like many other websites, we rely heavily on advertising." Okay, now, you can imagine with that sort of website traffic what sort of revenue their site is able to generate from all 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" - now, we're hearing from a listener of ours who is over on the implementation side of all this. He says: "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" - and he cites the site that offers the service - "first-id.fr, while others leverage ISP data to identify users based on their Internet connection" - and then again he says - "like Utiq.com. Additionally, there are methods involving email or phone numbers for cross-website identification, like LiveRamp.com."

He said: "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" - meaning, you know, the amount of money they get - "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 being 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 and phone numbers. I believe it's crucial to inform your audience about this trend. 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."

Okay. So first of all, I thank our listener for this view from the trenches. It is disappointing, but unfortunately not surprising. It was the subject of our "Unforeseen Consequences" podcast back on February 6th of this year. Here's the way to think about this: Third-party cookies enabled tracking of users based only upon the ads that were being 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 now, 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 specifically 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 third-party advertisers, there is nothing whatsoever we can do about that. Anything a website knows about you will now be shared with third parties. In many cases, as we recently saw with Microsoft, which was forced to disclose this due to the GDPR, I think it was, what, more than 700 was it, or 500 and some odd, I don't quite remember. But a phenomenal number, you know, many, many hundreds of individual third parties they were confessing they would be sharing anything they had about their visitors with. We talked about websites beginning to want their visitor's email addresses.

And Leo, you pointed out that even if we give them our throwaway email, if we always give them the same one, it still identifies us as efficiently as if we were using 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 so that they are then able to turn that information over to their advertisers. How many are going to see this as a problem? I would venture probably not that many.

So what was once tracking being done without website assistance is evolving into collusion between websites and their advertisers. You know, pay us, and we will tell you everything we know about our visitors. 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 it's going on, and I suspect that many wouldn't care anyway.

Leo: And this is, by the way, exactly what's happening to podcasts, as well. The difference is we don't have any information to collude with advertisers. And when they do ask us to put tracking pixels in, or beacons of some kind, we just say no. And we try to constrain that. And it hurts us, which is the reason why most advertisers now move to places like Spotify because they can get that information. We're kind of out of luck.

Steve: Yeah, they want it, yeah.

Leo: That's why we want people to join the Club because ad support's just not going to do it in the long run.

Steve: Kevin van Haaren tweeted: "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 issue. 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." So yes, we had heard that support was in beta and coming soon, but I hadn't noticed that Bitwarden's support for mobile was out now, was 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 tweeted: "Regarding Passkeys, help me out here. I feel like you and Leo are missing the point." And Leo, it's actually more my fault than yours, so, you know.

Leo: Oh, I'm good at missing points. Go right ahead.

Steve: He said: "Or am I?" He said: "I thought Passkeys were to say, 'Hey, this device has already logged in properly so let's make future logins super easy but also secure.' So that would mean I don't," he says, "I don't want my Passkeys to be exportable. If I ever want to log in on another device or OS or ecosystem at all, I want to prove that it's me all over again with whatever way I do that on that website, hopefully with multifactor authentication. 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 own ecosystems, but only if it's really, really, really securable. Generally speaking, having Passkeys exportable is as bad as Firesheep days when grabbing someone's session cookie gave an opponent 100% impersonation of a victim. Yes? No? Thanks."

Leo: No. No.

Steve: "Listener from Episode 1. Rob." Okay.

Leo: Rob misunderstood.

Steve: So I think Robert makes a valid point. Although entirely different, although, well, okay.

Leo: Look. Passkeys are being proposed as a password replacement. Passwords are not specific to the device you use, nor should Passkeys be. It's the same thing.

Steve: So, okay. Another entirely different way to think of Passkeys is the way he does.

Leo: Yeah, but it's not right.

Steve: Okay. 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. Okay, I can see that as a workable model. But here's the critical factor, and this is what you're alluding to, Leo. That model only works in a world where every website allows for any number of Passkeys to be registered to any single web account. And as we've been saying in the last couple weeks, it is apparently the case

that any number of Passkeys can be registered to any single web account, maybe without limit? That's the question. Where's the limit?

So, you know, 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." You know, we don't know if that's ever going to happen, but we know that it could. And the experience that Kevin just reported of creating a single Passkey with Bitwarden on his iPad, then having Bitwarden later login for him using the same synchronized Passkey under Windows, well, that's pretty slick.

Leo: He just misunderstands what Passkeys are all about. He's saying he's thinking they're like session cookies, which is the problem Firesheep had. But Passkeys are not as easily accessed as cookies. I would hope they're better secured than that.

Steve: They're public key crypto.

Leo: Right. So they're secure.

Steve: And so they're not at all the same.

Leo: And furthermore, they're being proposed as a replacement for passwords. So that's not what he's just described. He's describing a replacement for session cookies. That's not what Passkeys are. So I think he just misunderstands what Passkeys are. They are a replacement for passwords and, as a result, are not tied or should not be tied to a specific device. And you're right. A password manager can be and probably should be the person that holds your Passkeys, just as they are the people who hold your passwords. Yeah.

Steve: Okay. Spencer Webb, he tweeted: "Enjoyed the eLORAN discussion. I know the guy at UrsaNav. 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."

Leo: He's probably a ham. I think that's a ham.

Steve: He is.

Leo: Okay.

Steve: Spencer is a serious radio guy.

Leo: That's what I thought.

Steve: We often exchange notes when something about radio comes up. I remember back in the days when you had to hold your iPhone in the proper way, he and I were having some conversations about antenna science. So anyway, 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.

Oh, this is interesting. Dr. Brian of London tweeted: "I integrated Passkeys into my own site as a secondary login 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 plugin wallet or something using QR codes that looks like SQRL to prove ownership of personal keys." So, okay. What he's saying is he has a website, and he rolled his own fancy login system which he's had for some time. But then he decided, hey, Passkeys is a standard. I'm going to add that to my site.

So he said: "To this I added the ability to associate one or many keys with an account and add/delete/rename them. One little gotcha which you probably only learn when implementing this." He said: "I store on my server a list of all public Passkeys, and every time I get a login 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."

He said: "I use that username to filter the list of public keys I send back to the client, which then figures out if the user's device has any of them." He says: "It works nicely with Apple Passkeys and other Passkeys which already sync across multiple devices nicely." So basically he's saying he rolled his own server-side Passkeys implementation. As a consequence, he has a bunch of accounts, and each of those accounts has Passkeys. He could send all of the public keys to the client, which would then say, oh, I found a match, which would tell him who it was that was wanting to log in. But instead he asks them for just their username, which allows him then to filter from all of his public Passkeys only those associated with that user and send those back in order to give it a chance to log in with a Passkey. So anyway, we have users who are implementing Passkeys on the server side, which is also very cool. Or rather listeners.

Shaun Merrigan said - remember that it was Shaun - oh, right. Okay. First of all, Shaun was the guy with the old LORAN receiver which woke up when eLORAN was turned back on late last year. Anyway, he heard us talking about him last week, and he followed up with a bit of more interesting information. He said: "To close the loop on this, my location is Edmonton, Alberta, Canada." Okay, that's where he is.

He said: "The three eLORAN stations that are currently testing are Fallon, Nevada; George, Washington; and Havre, Montana." So he's receiving signals from those three locations. He says: "This is my best information. Currently, my old Austron 2100F is showing 2.8E-12 seconds offset from GPS." Again, 2.8E-12 seconds offset from GPS. So, yeah, lots of accuracy in LORAN available timing data. And really it sounds like once this is turned back on, all of our clocks that used to synchronize on WWVB...

Leo: Oh. Lost LORAN.

Steve: Yeah.

Leo: Oh, that's cool.

Steve: Because, you know, that's not very reliable, that whole WWVB.

Leo: And they wanted to turn it off for a long time.

Steve: Yeah.

Leo: Yeah. Well, clever.

Steve: Talk about range. That's really cool range. Oh. Markus Daghall tweeted: "Hi, Steve. While looking at the PIN heat map graph, the number 1701 seems to be more prevalent than its surrounding numbers." Which I love because we know what that is.

Leo: Star Trek NCC 1701.

Steve: Exactly.

Leo: That's a good PIN. Wish I'd thought of that.

Steve: NCC 1701. That is great. Ed Ross tweeted: "Re Big Yellow Taxi," he said, "presumably that system helps in situations where 'you don't know what you've got till it's gone.'"

Leo: They paved paradise, put up a parking lot.

Steve: And that was your observation last week, Leo.

Leo: Yes, yes, yes.

Steve: Riny, and I can't even begin to pronounce this guy's last name, he's in Spain, H-E-I-J-D-E-N-D-A-E-L.

Leo: It's a Dutch name, Heijdendael.

Steve: Perfect, thank you, Leo. So he wrote: "As many, I started the FIDO1 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,' 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 two-factor authentication that would work as a standalone device when traveling, and even that is in their device. I just don't understand why YubiKey is still pushed as the de facto standard. What do you think?"

Now, he included in his note a link, and it's in the show notes at the top of page 15. And he finished, saying - and Leo, I should mention I have two on order now.

Leo: Yeah, I think I'm about to buy some, yeah.

Steve: Yup. He said: "Keep up the good work. By the way, I silently suspect that you were hired by the UK government to write their specs for them." And of course he's talking about the fact that we talked about the requirements that the UK had for their consumer IOT devices, and it did actually sound like, you know, they've been listening to this podcast.

Leo: It sure did, yeah, yeah.

Steve: Okay. So 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 will 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. And I should note that YubiKey has since announced a new key, and I don't remember the number. It might be 200. But even now it doesn't appear to still be 300. Or maybe 100. Anyway, the bad news is unfortunately these guys are in Switzerland, and the one we want is currently sold out.

Leo: But it says June 17th shipping, so that's not so...

Steve: Oh, that's good.

Leo: Yeah.

Steve: That's good. Although shipping, unless you choose postal mail, which they discourage, is twice the cost of the dongle. So, okay. Anyway. Anyway. The way I know that is that I've ordered two, and they're on the way. So anyway, thank you very much, Riny, for providing a direct link to the Token2 page, which as I said is in the show notes.

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

Leo: Clearly they put a fairly potent chip in there. So when it says 300, that could be 300 Passkeys?

Steve: It is, 300 FIDO2 WebAuthn Passkeys.

Leo: Oh, wow. 300 is probably a good start. I don't - at least for a while, yeah. All right, Steve. Let's talk about 200 doctors can't be wrong. Or something like that.

Steve: Our listener, Robin van Zon in the Netherlands, brought this recently produced letter to my attention. So 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." Now, we're interested in this, of course, because this is all about whether we're going to have backdoors and something is going to be monitoring communications for, you know, grooming and CSAM material and so forth. "So as of the 7th of May, exactly two weeks ago today, the letter has been signed by 312 scientists and researchers across 35 countries." I mean, it is the Who's Who of security and research.

So, and what's interesting is that there has been some very good, you know, good faith back-and-forth here. So this is not an open letter that's just being blown off and being ignored. The EU's regulators and legislators have changed their legislation in an attempt to solve the problems that were earlier voiced. As we're going to see, not only are they not there yet, but there's real good reason to believe, as we probably all know, you can't get there from here.

Okay. So 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 assembled. And it's not overly long, you know, because the devil as it turns out is in the details, and because there's probably no more important issue on the table at this moment, arguably in the world, than what the EU's political class will finally decide to do about this. And importantly, as we'll see, this is the technical response to the politicians' responses to the previous technical response. And as I said, what's heartening is that both sides so far appear to be negotiating here in good faith. And the politicians are at least listening.

So as we know, for their part, the UK was faced with the same problem and serious opposition to their similar proposal to require all private conversations to be monitored for content. What they did was wisely added the caveat "where this can be proven to be technically feasible without compromising security," which allowed the politicians to say that they had passed legislation and allowed all the messaging providers to continue offering fully private end-to-end encryption because it hadn't been and cannot probably be proven to be feasible without compromising security. So win-win, win-win-win.

Okay. 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're writing in response to the new proposal for the regulation introduced by the Presidency on the 13th of March, 2024." So 13th of March, right, just a couple months ago.

"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 of 2023" - so nearly a year ago - "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' decision in Podchasov v. Russia on the 13th of February of this year. We elaborate on these aspects below."

Now, just to interrupt here, 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 fully encrypted 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 backed that up. And I skipped all of the earlier paragraphs.

Here's what those two paragraphs say. 79 says: "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." In other words, yes, the intention is legitimate, but the only way you can do this is by weakening encryption for everybody. And that's not a proportionate response.

And then paragraph 80 says: "The Court concludes from the foregoing" - and that's all the other paragraphs that I'm sparing everyone - "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.

"Insofar 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 listeners' sanity, as I said, 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 policymakers 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." Again, "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."

So then they bring up five points. The first: "The proposed targeted detection measures will not reduce risks of massive surveillance." They said: "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," and they said, "as of the last draft, twice for known CSA material and three times for new CSA material and grooming."

They said: "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, on the order of millions."

So they then had a footnote which explains how they draw this conclusion. They said: "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, in other words, in a thousand times it incorrectly classifies non-CSAM as CSAM, which is much lower than any currently known detector." Right? So they're drawing like a best, absolutely beyond best possible case.

They said: "Given that WhatsApp users send 140 billion messages per day, even if only one in 100 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 five 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" - that is, you know, repeated hits on a given individual before you raise the alarm in order to bring down basically the rate at which alarms are being raised, you need to raise that number of repetitions, they say - "would grow significantly to the point of not effectively reducing the CSAM sharing capabilities." Meaning detection would be effectively neutered.

Then they said: "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, for example, photos to doctors, or legitimate sharing across family and friends." They said: "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."

Okay. So in other words, the politicians proposed to minimize false positive detections by requiring multiple detections for a single individual before an alarm is raised. But the science of statistics says that won't work because entirely innocent photographs of one's children will not be evenly distributed across the entire 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 because there is massively non-equal distribution of content that might falsely trigger CSAM detection.

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.'" I should explain that another amendment to the proposed legislation involves their attempt, the legislators' proposal of attempting to divide applications, that is, you know, like WhatsApp as an application, Telegram as an application, 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 what was previously low-risk and high-risk services. As a result, almost all services would 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'll 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 the 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's not technologically possible to limit its application to a subset of 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 blanket detection regulation." So basically these guys are just, you know, they're cutting through these proposals one after the other, very carefully backing up their statements with, you know, actual data.

The second is: "Detection in end-to-end encrypted services by definition undermines encryption protection." They go over this again, explaining why that's the case. Oh, and they note one of the other arguments is, and we've talked about this on the podcast, the idea of adding age discrimination. Well, they said: "Introducing more immature

technologies may increase the risk. And they note that their 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," they said, "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're meant to prevent."

So they conclude, saying: "With secure paths forward for child protection," and this is really good. They said: "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." So 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.

So I think the problem is that this is like an iceberg. CSAM is the tip of the iceberg that is the visible manifestation of something that is 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 a true, 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. And of course we have no idea what's going to finally happen, which is what makes all this so interesting. And it is obviously very important. So stay tuned.

Leo: Yeah. It's so much easier to go after the symptom than the cause; you know? So much - yeah, yeah.

Steve: Isn't it? Exactly right. That is exactly right.

Leo: And unfortunately there's huge side effects to going after the symptom that make for more problems. So it's not really a great solution.

Steve: Yeah. And nothing prevents the politicians from wanting to save face or look good by saying, "We did this."

Leo: We fixed it. It's all over.

Steve: Yeah.

Leo: But, see, it's not going to be all over. And really that's the nut of it.

Steve: And the best, yes, the most important reminder is that CSAM is the output of the practice, not the practice itself. And it's the practice that you want to curtail.

Leo: Right. Well, good stuff, as usual. You have no fear to go where angels fear to tread, and that's good. That's good. That's what we want. You're going to hear it here. You're going to hear it all.