LastMess

Description: This week we share some exciting and hopeful news about the UK's Online Child Safety legislation. What does it suggest for the future? How was it that Microsoft's super-secret authentication key escaped into the hands of Chinese attackers who then used it to breach secure enterprise email? What, if any, lessons did Microsoft learn? Why am I more glad than ever that I'm driving a 19-year-old car after the Mozilla Foundation shared what they learned about all of today's automobiles? And then, after sharing and exploring some feedback from our listeners, we're going to examine the horrifying evidence that the data stolen from the LastPass breach is being successfully decrypted and used against LastPass users.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-939.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-939-lq.mp3

SHOW TEASE: Coming up next on Security Now!, it's me, Jason Howell, filling in for Leo Laporte. But of course you check out Security Now! because you want to hear Steve Gibson share everything that there is to know this week about, well, security news. And actually there's news about the UK's Online Child Safety legislation that might be good news for encryption. Steve tells you all about that. Why Steve's happier than ever before to be driving a 19-year-old vehicle makes me rethink my purchase of a new vehicle just last year. Tons of feedback, some really great feedback from you, the listeners and viewers of this show. And then finally Steve ends things off by sharing why it's pretty clear that the data stolen during that devastating LastPass hack is actually being decrypted and cashed in on. You won't want to miss it. Security Now! is next.

JASON HOWELL: This is Security Now! with Steve Gibson, Episode 939, recorded Tuesday, September 12th, 2023: LastMess.

It's time for Security Now!. I am Jason Howell, filling in for Leo Laporte, who is, I think, on the other side of the country right now. But, you know, you know who's going to be here with me, of course, none other than Steve Gibson. It's good to see you again, Steve.

Steve Gibson: Jason, it's great to be with you. Leo is with his mom.

JASON: That's right.

Steve: Who everyone thinks is the cutest thing they have ever seen.

JASON: Yeah. She's been on the network a number of times, and it's always an enjoyable experience when she's on. So, yeah, he's assisting her. So, it's good to see you.

Steve: So we're here for Security - Security Now!.

JASON: Security Not Now.

Steve: I was going to say Security Nine because it's Episode 939, on 09/12/23. So, yeah. This week we're going to share some exciting and hopeful news about the UK's Online Child Safety legislation, and we're crossing our fingers. And once again the tech press kind of launched a little prematurely, but then got corrected, and there's some fun story there. And then we're going to explore what that suggests for the future. We also learn how it was that Microsoft's super-secret authentication key escaped into the hands of Chinese attackers, who were then able to use it to breach the secure enterprise email to some dramatic effect. Also, what did Microsoft learn from that, if anything?

Also we're going to look at why I am more glad than ever that the car I'm driving is 19 years old, still goes great. And this is after the Mozilla Foundation shared what they learned about all of today's automobiles from a privacy standpoint. It's, well, there was only one thing they've ever encountered that was worse, and I don't remember what it is, but it's in the show notes, so we'll get there. And then, after sharing and exploring some feedback from our listeners, believe it or not, we're going to examine the horrifying evidence that supports the belief that the data stolen from the LastPass breach is being successfully decrypted and being used against LastPass users.

JASON: Oh, dear.

Steve: As a consequence, this podcast is titled "LastMess."

JASON: Probably not the last mess from the LastPass mess.

Steve: So, yeah.

JASON: But holy moly. That just, yeah, it just keeps on giving, unfortunately. Bleah. All right. Well, we will get there. That's a whole lot of interesting stuff to talk about today on Security Now!. Before we get to all the news, we've got a picture that makes a whole lot of sense when you look at it. I mean, I guess it made sense to someone to do it, but, yeah, this is...

Steve: Yeah, we may have shown this before. But it came up again, and it's just such a perfect picture for the podcast. I gave this picture the title "The plumber's contract didn't say anything about moving any rocks." And so we have this, it's not clear which way the water is flowing through this pipe. But the pipe is trying to go - it comes from offscreen on the left, and it wants to go to a pipe which disappears into the ground over on the right. Unfortunately, there is a big honking rock, like, blocking the pipe. And the rock looks like, it doesn't look like it's glued down or anything. I mean, it looks like it's a big rock. You can presumably move the rock. For whatever reason, the plumber did not want to move this rock. Instead, he did - okay.

Those of us who are old enough to remember "The Three Stooges" will remember that famous plumbing episode where, I don't know, Moe or Curly or somebody was in the bathtub, and they tried to plumb around a leak. And anyway, so this pipe basically circumnavigates the rock in pretty much as few pieces as possible. So the water is flowing, and the rock stays where it originally was. So anyway, one of our fun pictures, I think.

JASON: I mean, it really does look like that rock could be moved. But maybe...

Steve: Doesn't it? It's like, what...

JASON: I mean, maybe it's not movable. Maybe it's like firmly embedded into the ground in a way that would require it to be destroyed. And I'm a plumber. I'm not a rock destroyer.

Steve: And you can kind of - you can see behind it. There's like a rock ledge or like a wall.

JASON: Yeah.

Steve: But this rock does not appear to be part of that.

JASON: No.

Steve: It looks like it's separately sitting on the ground. And I don't know.

JASON: I mean, apparently there was some reason why that rock wasn't going anywhere. Maybe it's a...

Steve: Jason, that's exactly where I was - I was about to say, there's a story here.

JASON: Yes, that we do not know the answer.

Steve: And we will never...

JASON: But I want to know the answer.

Steve: We will never know. Now, given the spread and breadth and ingenuity of our listeners, I wouldn't be surprised if one of them finds this at some point and goes and tries to move it. So if any of our listeners encounters this particular rock, we want to know.

JASON: Yes.

Steve: Can it be moved? Or...

JASON: Yeah. If you run into this rock, if you find yourself out in the world, and you cross paths with this rock, just do us a favor and see if you can nudge it. Does it even shake in place? And if it shakes in place, then we've got our answer.

Steve: Give it a kick. Maybe get out a crowbar if you have one handy.

JASON: Sure.

Steve: We want to know.

JASON: Sure.

Steve: So just a little follow-up on my current side project. I was getting ready. Remember I talked about ValiDrive last week, which arose from one of the SpinRite 6.1 testers having SpinRite rejecting one of his many smart - one of his smart drives. One of his many thumb drives. And then we dug into it, and we figured out that this was a fake drive. So I decided - and Leo was really moved by this revelation, and he said, whoa, you know, we need to know about these.

So I created this little freeware called ValiDrive, which I had expected I would finish with last week. But just as we were kind of getting ready to get there, one of our listeners

showed that a drive which we believe is fake, and I've now, since then, I've absolutely confirmed it, is passing ValiDrive's test when it should not.

So what's happening is ValiDrive jumps around the drive in a random sequence, checking 576 equally spaced out locations to verify that there is actual storage there. I think what's happening is that some read caching that exists in the chain between my code and the actual USB stick somewhere is generating false positives. So we're seeing ValiDrive is showing green when it should be showing red. Anyway, as a consequence I didn't finish, I didn't publish this thing yet. We'll have it next week. In fact, I was cheating a little bit, Jason, while you were reading our first advertiser sponsor. I was wondering why an undocumented command I'm using wasn't being - anyway. I'm at work on it literally as we speak.

JASON: You're amazing, Steve. You're podcasting while you're validating code. That's amazing.

Steve: So it'll be soon, and it'll be working correctly. Okay. So who blinked? What can only be called wonderful and welcome news surfaced in the middle of last week from the UK. Now, the short version is, the UK appears to have blinked. And in the face of all secure messaging apps - and now as we've just recently covered, Apple taking their stand, which was the topic of last week's podcast, saying firmly, unh-unh, we're not doing this - everybody said they were going to pull their services from the UK rather than sacrifice the privacy of their users, rather than compromising in any way. The UK apparently said, "Oh. Well, we never said that we wanted you to do that." Uh-huh. Right.

But of course nothing involving politicians and bureaucracies is ever clean and simple. And the details here are at least somewhat interesting. What first happened was that last Tuesday the Financial Times was the one who broke the story. And then the tech press jumped on it because this was big news.

9to5Mac's headline was: "Future of iMessage safe in the UK as government backs down on encryption." Wired carried the headline: "Britain Admits Defeat in Controversial Fight to Break Encryption." And their subhead was: "The UK government has admitted that the technology needed to securely scan encrypted messages sent on Signal and WhatsApp doesn't exist, weakening its controversial Online Safety Bill." Computerworld's story began with "UK rolls back controversial encryption rules of Online Safety Bill." CyberScoop headlined their coverage "UK lawmakers back down on encryption-busting 'spy clause'." And Infosecurity Magazine's headline was "UK Government Backs Down on Anti-Encryption Stance." All right? Okay. So anyway, everybody gets the idea. This is what all the headlines were.

Unfortunately, much as those were all attention commanding and welcome headlines, none of that was true. Well, or at least they were all probably deliberate oversimplification and exaggeration click-bait which, predictably, did not sit well with the UK. The politicians there didn't like those headlines. So the following day, last Thursday, we see follow-up headlines such as "UK tries to claim it hasn't backed down on encryption at all," and Reuters' headline was "UK has not backed down in tech encryption row, minister says."

And so anyway, here's what Reuters explained because their coverage is short and succinct. So LONDON, Sept 7 (Reuters): "Britain will require social media companies to take action to stop child abuse on their platforms, and if necessary work to develop technology to scan encrypted messages as a last resort, technology minister Michelle Donelan said on Thursday." And we've talked about dear Michelle in the past. She's the one who's in charge of this.

So Reuters said: "Platforms including Meta's WhatsApp and Signal have been fighting Britain's Online Safety Bill, which is currently being scrutinized by lawmakers because they say it could threaten the end-to-end encryption that underpins their messaging services. Junior minister Stephen Parkinson appeared to concede ground to the tech companies' arguments on Wednesday, saying in parliament's upper chamber that the Ofcom" - that's their communications regulator - "would only require them to scan content where 'technically feasible.'" Okay, now, that's the first time anyone had heard that. Of course, and then Reuters reminds us that: "Tech companies have said scanning messages and end-to-end encryption are fundamentally incompatible."

Okay. So in other words, "not technically feasible." So essentially, by admitting and facing reality, this junior minister Stephen Parkinson set off a firestorm. Was the fire set deliberately? Did the senior minister set this up to have junior drop this and then hide? You know? Maybe I'm being too cynical. I don't know. Reuters continues their coverage, saying senior technology minister Michelle Donelan, however, denied the following day, on Thursday, that the bill had been watered down in the final stages before it becomes law.

She told Times Radio: "We have not changed the bill at all." Okay, which doesn't seem to be true, but that's what she said. "If there was a situation where the mitigations that the social media providers are taking are not enough" - and, okay, we already know they won't be. She said: "And if after further work with the regulator they still cannot demonstrate that they can meet the requirements within the bill, then the conversation about technology around encryption takes place." Okay. Huh? What does that mean? Anyway, she said: "Further work to develop the technology was needed, but added that government-funded research had shown it was possible." Okay, all of this is new information; right? So, okay.

So that's the official CYA story from the UK's senior technology minister. But that's not the whole story because the Online Safety Bill actually was amended, despite the fact that Michelle Donelan just said it had not been changed at all, and she's desperately trying to obfuscate that fact. So here's the way AppleInsider explained what happened. They said: "Despite introducing a clause that means its Online Safety Bill is no longer a concern for Apple, WhatsApp, or users, the UK government is insisting with a straight face that it's still exactly as tough on Big Tech as before.

"On Wednesday, the UK Parliament debated an Online Safety Bill that, in its original form, would have seen Apple, WhatsApp, Signal, and more shutter their messaging and social media services in the country. Bowing to that pressure," wrote AppleInsider, "the UK regulator Ofcom introduced a face-saving clause that effectively stopped the country's nonsensical demands to break end-to-end encryption. Except the Conservative government that was pushing for this against the advice of security experts and even an ex-MI5 head insists that it has not even blinked."

As spotted by Reuters, UK technology minister Michelle Donelan told Times Radio the same thing I just shared: "We haven't changed the bill at all. If there was a situation where the mitigations that the social media providers are taking was not enough, and if after further work with the regulator they still can't demonstrate that they can meet the requirements within the bill, then the conversation about technology around encryption takes place." Anyway, I don't think anyone's ever going to listen to her or take her seriously again.

But AppleInsider says: "Ofcom's amendment to the bill said that firms such as Apple would be ordered to open up their encryption only 'where technically feasible and where technology has been accredited as meeting minimum standards of accuracy in detecting only child sexual abuse and exploitation content.'" In other words, you know, not really. But we had to say something; right?

So AppleInsider opines, saying: "There is no technology today that will allow only the good guys to break end-to-end encryption, and there never will be." Period. "Consequently," they write, "the Tory government can argue and is arguing that no word has been changed in the bill." Ah, but words have been added, and they neuter the entire - yeah, we didn't change anything.

JASON: We didn't change anything. We just...

Steve: But we added a few more words there to the end.

JASON: ...sprinkled a little on top, yeah.

Steve: Yeah, exactly, they just kind of like, oh. And AppleInsider says they neuter the entire nonsensical and unenforceable plan. Okay. So that's the story. And it's big news because of the critically important precedent that this sets. For their coverage of this, Wired interviewed Signal's quite outspoken president, Meredith Whittaker, who we've also often quoted here. So here's what Wired wrote of what Meredith had to say. They said: "Meredith Whittaker, president of the Signal Foundation which operates the Signal messaging service said: 'It's absolutely a victory. It commits to not using broken tech or broken techniques to undermine end-to-end encryption.'"

And then Wired said: "Whittaker acknowledges that, okay, it's not enough that the law simply won't be aggressively enforced. But it's major. She said: 'We can recognize a win without claiming that this is the final victory.'"

Then Wired continues, saying: "The implications of the British government backing down, even partially, will reverberate far beyond the UK, Whittaker says. Security services around the world have been pushing for measures to weaken end-to-end encryption, and there is a similar battle going on in Europe over CSAM, where the European Union Commissioner in Charge of Home Affairs has been pushing similar, unproven technologies. Whittaker said: 'It's huge in terms of arresting the type of permissive international precedent that this would set. The UK was the first jurisdiction to be pushing this kind of mass surveillance. It stops that momentum. And that's huge for the world.'"

And yes. I believe this is authentically a huge deal. No one has any real problem with face-saving legislation being created to allow the politicians to tell their CSAM activists that they now have powerful new legislation on the books which will, the moment it can be shown to be technically feasible to do this with the required level of accuracy, compel all encrypted messaging providers to protect the children. And those politicians can truthfully state that this was the strongest legislation they were able to obtain. Because indeed it was.

We know that this will in no way pacify Sarah Gardner, whom we talked about last week, after she threatened Tim Cook at Apple with her forthcoming pressure campaign, which starts this week, to compel Apple to perform client-side scanning for known CSAM imagery. But Sarah appears to be a lost cause. She has no problem demanding whatever concessions to everyone else's security and privacy might be needed to even incrementally offer improved protection for children. Everyone is for improving child protection, but there's no way to do that without compromising everyone's security and privacy, including the children's.

So the free world appears to have just taken the first big step toward the resolution of the encryption dilemma. It's going to be interesting now to see what the European Union does. You know, maybe they'll also just put the same sort of caveat into their legislation, and everyone can continue ignoring it, which would be wonderful.

JASON: And then does that end up trickling down here into the U.S. as the U.S. government tries to pursue a no-encryption policy?

Steve: You know, what we see with things like the EU and the GDPR, annoying as that GDPR is, it's helping us, I think, and U.S. politicians to say, oh, yeah, I guess, you know, privacy really is good. You know, maybe we should have some of that here, too.

JASON: It's at least forcing more of a conversation and actual, you know, taking a look, yeah, at these issues.

Steve: Yeah. And all of the encryption providers can say, when the U.S. tries to do this, hey, you know...

JASON: Over there.

Steve: Yeah, exactly. Over there they worked it out. They're okay with this. So just cool your jets.

JASON: Yeah, yeah, interesting.

Steve: Okay, so as we know from July, a Chinese-based attacker known as Storm-0558 somehow managed to acquire one of Microsoft's, what was supposed to be a very secret key which then allowed it to forge login tokens which they were able to use to access the private email of OWA and Outlook.com users to, like, serious effect. In the wake of these revelations, the entire security world has been left wondering exactly how Microsoft had managed to fumble the crucial protection of this very important key.

So last Wednesday, after a series of preliminary blog postings, Microsoft finally provided what may be the conclusion of their investigation, and it did answer some of these questions. Here's what Microsoft shared. They wrote: "Microsoft maintains a highly isolated and restricted production environment. Controls for Microsoft employee access to production infrastructure include background checks, dedicated accounts, secure access workstations, and multifactor authentication using hardware token devices. Controls in this environment also prevent the use of email, conferencing, web research, and other collaboration tools which can lead to common account compromise vectors such as malware infections or phishing, as well as restricting access to systems and data using Just in Time and Just Enough Access policies."

Okay, so that paragraph went to explaining, like, all the things they did and designed on purpose to prevent anything like this from ever happening. And they go on: "Our corporate environment, which also requires secure authentication and secure devices, allows for email, conferencing, web research, and other collaboration tools." So in other words, there's the production environment, and then the corporate environment. They said: "While these tools are important, they also make users vulnerable to spear phishing, token stealing malware, and other account compromise vectors. For this reason, by policy and as part of our Zero-Trust and 'assume breach' mindset, key material should not leave our production environment." That is, the first one that's highly protected.

"Our investigation found that a consumer signing system, okay, a consumer signing system crash in April of 2021" - right, so 2.5 years ago - "a consumer signing system crash in April of 2021 resulted in a snapshot of the crashed process, a crash dump. The crash dumps, which redact sensitive information, should not include the signing key." And I'll just put a pin in this here and add that, okay, the crash dumps should not need to redact sensitive information since a signing key should never be in RAM to be dumped

after a crash. They should be in a hardware security module. What the heck is a signing key ever doing in RAM? But we'll get back to that.

So Microsoft says: "In this case, a race condition allowed the key..." - which obviously was present in RAM to be put out in a dump. They said: "...to be present in the crash dump." Okay. "A race condition allowed the key to be present in the crash dump. This issue has been corrected." So if we had a bell, we would ring it. Ring, you know, ding. There's the first bug fixed. "The key material's presence in the crash dump was not detected by our system." Ding. They said: "This issue has been corrected." Bug number two.

"We found that this crash dump, believed at the time not to contain key material, was subsequently moved from the isolated production network into our debugging environment on the Internet-connected corporate network." Which we've already said is not as secure as a production environment. They said: "This is consistent with our standard debugging processes." Then, "Our credential scanning methods did not detect its presence." Ding. Bug number three. "This issue has been corrected." Okay. So, so far we've got three strikes.

They continue: "After April 2021" - that's when this crash occurred - "when the key was leaked to the corporate environment" - and remember the key was leaked because it passed through three bugs, none of which should have existed. They've been fixed now. Okay. "When the key was leaked to the corporate environment in the crash dump, the Storm-0558 actor was able to successfully compromise a Microsoft engineer's corporate account. This account had access to the debugging environment containing the crash dump which incorrectly contained the key. Due to log retention policies, we don't have logs with specific evidence of this exfiltration by this actor, but this was the most probable mechanism by which the actor acquired the key."

Okay. So why was a consumer key able to access enterprise email? Right? Microsoft explains: "To meet growing customer demand to support applications which work with both consumer and enterprise applications, Microsoft introduced a common key metadata publishing endpoint in September of 2018. As part of this converged offering, Microsoft updated documentation to clarify the requirements for key scope validation, which key to use for enterprise accounts and which to use for consumer accounts.

"As part of a pre-existing library of documentation and helper APIs, Microsoft provided an API to help validate the signatures cryptographically, but did not update these libraries to perform this scope validation automatically." Ding. Number four. "This issue," they say, "has been corrected. The mail systems were updated to use the common metadata endpoint in 2022. Developers in the mail system incorrectly assumed libraries performed complete validation and did not add the required issuer scope validation. Thus, the mail system would accept a request for enterprise email using a security token signed with the consumer key." Ding. Number five. "This issue has been corrected using the updated libraries."

Finally: "Microsoft is continuously hardening systems as part of our defense in depth strategy. Investments which have been made related to MSA key management are covered in the" - and they have a blog, the storm-0558 blog. "Items detailed in this blog are a subset of these overall investments." In other words, as a consequence of this, we are now doing better than we were before. They said: "We are summarizing the improvements specific to these findings here for clarity."

We have four bullet points: "Identified and resolved race condition that allowed the signing key to be present in crash dumps. Enhanced prevention, detection, and response for key material erroneously included in crash dumps. Enhanced credential scanning to better detect presence of signing key in the debugging environment. And finally, released

enhanced libraries to automate key scope validation in authentication libraries, and clarified related documentation."

And really, those last two things, the library kind of stuff, synchronization and so forth, that just kind of feels like huge corporation stuff; like yeah, it's understandable that something changed over in Department A, and Department B was using it, but they didn't get theirs refreshed and updated. I mean, yeah.

Okay. So as Microsoft has explained this mess-up, a series of five separate and previously undiscovered bugs, all of which they have since found and fixed, were responsible for allowing a key, which should have never left Microsoft, to be exfiltrated by Chinese attackers and then used to remotely compromise what should have been high-security enterprise email. It's obvious that the key should have never been allowed to leave Microsoft.

But as I said before, what they appear to have conveniently skipped over is why that key ever left the HSM, you know, the Hardware Security Module, which was the only place it should have ever existed. It's really worth having all of us note that not one of those five flaws would have caused any trouble if that secret key had not been in the system's RAM at the time of that fateful crash. This is precisely why HSMs exist. It's why, for example, GRC's code signing keys do not ever exist in any RAM. They are sequestered in hardware, completely inaccessible to the outside world once installed there, and only able to be used to sign signature hashes. You cannot query the hardware for the key. It won't give it to you. It will only agree to use it until it expires. You know?

I've always said that anyone can make a mistake. In this instance Microsoft made five big ones. But policy is a different matter. And Microsoft completely dodged the question of how they could have ever had a policy to allow a crucial signing key to be present in RAM. That's just never okay, and it got them in trouble here. So anyway, at least now we understand how it happened. It was a bunch of mistakes. I'm actually, I'm impressed that they have all of those buggy things in the pipeline that were meant to work and meant to catch this problem. I mean, there was an intention to prevent this from happening. Again, they're doing an awful lot of work which could have been resolved by having this in a hardware security module in the first place. So I don't get that. But the fact that they have those things demonstrated noble intent.

Unfortunately, they were buggy. And so the key slipped right past all three of them in that chain. But kind of cool that it was there. But anyway, at least we know how it happened. It was a miss. And they're going to do better. But really the number one takeaway from this entire debacle is don't have important keys in RAM, ever. No matter how tightly you believe you have protected them from ever being divulged. That's why we have HSMs, and they're not even expensive. I'm sure Microsoft can afford one. I've got several.

Jason, let's tell our listeners why we're here, and then I'm going to explain why I'm glad my car is as old as it is.

JASON: I can't wait for you to explain to me why I shouldn't be driving these new cars, or newer cars. At least newer than 19 years old.

Steve: Sorry to tell you, my friend, that Tesla is the worst.

JASON: I'm not that surprised to hear that. We're going to talk all about that coming up next. Here's where the world comes crashing down on us if we have a new car or a newer than 19-year-old car like Steve has, apparently. We might not be so happy about that after hearing this story.

Steve: Well, to be, as we often say, to be forewarned is to be forearmed.

JASON: Sure.

Steve: I titled this piece "The car I drive is 19 years old, and I'm more glad than ever." The reason is, it's a car. It's not a continuously connected mobile entertainment system on wheels.

JASON: Right.

Steve: It's a car. It does what a car is supposed to do. It moves my butt from one place to another. The reason I'm more glad than ever that all my car does is move me, is that I read the research that was just conducted and published last Wednesday by The Mozilla Foundation. They titled this "It's Official: Cars Are the Worst Product Category We Have Ever Reviewed for Privacy."

JASON: Oh, no.

Steve: And a subhead might be "25 car brands tested and 25 car brands failed." So here's what their research uncovered. I've edited their posting a little bit for the podcast. They said: "Ah, the wind in your hair, the open road ahead, and not a care in the world except for all the trackers, cameras, microphones, and sensors capturing your every move. Ugh. Modern cars are a privacy nightmare," they wrote.

"Car makers have been bragging about their cars being 'computers on wheels' for years to promote their advanced features. However, the conversation about what driving a computer means for its occupants' privacy hasn't really caught up." As we'll see. I make this point later. I think that's exactly the case. It's happened quickly, and we haven't caught up.

They said: "While we worried that our doorbells and watches that connect us to the Internet might be spying on us, car brands quietly entered the data business by turning their vehicles into powerful data-gobbling machines. Machines that, because of all those brag-worthy bells and whistles, have an unmatched power to watch, listen, and collect information about what you do and where you go in your car.

"All 25 car brands we researched earned our 'Privacy Not Included' warning label, making cars the official worst category of products for privacy that we have ever reviewed. The car brands we researched are terrible at privacy and security. For one thing, they collect too much personal data - every single one of them. We reviewed 25 car brands in our research, and we handed out 25 'dings'" - you know, as in a dent in your car - "25 'dings' for how those companies collect and use data and personal information." That's right. "Every car brand we looked at collects more personal data than necessary and uses that information for a reason other than to operate your vehicle and manage their relationship with you.

"For context, only 23% of the, by comparison, only 23" - I'm sorry. "Only 63% by comparison of the mental health apps, and they said "another product category that stinks at privacy we reviewed this year received this 'ding.' But it was 100% for automobiles." So cars are worse than mental health apps at managing privacy is their point.

They said: "And car companies have so many more data-collecting opportunities than other products and apps we use, more than even smart devices in our homes or cell phones we take wherever we go. They can collect personal information from how you interact with your car, the connected services you use in your car, the car's app - which provides a gateway to information on your phone - and can gather even more

information about you from third-party sources like Sirius XM or Google Maps. It's a mess.

"The ways car companies collect and share your data are so vast and complicated that we wrote an entire piece on how that works. The gist is they can collect super intimate information about you, from your medical information, your genetic information, to your 'sex life.'" And they put in parens "(seriously), to how fast you drive, where you drive, and what songs you play in your car - in huge quantities. They then use it to invent more data about you through inferences about things like your intelligence, your abilities, and your interests." And get this. "Most - 84% - share or sell that data." Okay. And just to stop here, that was the surprise to me. It's like, what? They are data retailers. They are retailing the data that they're collecting about their drivers.

Mozilla said: "It's bad enough for the behemoth corporations that own the car brands to have all that personal information in their possession to use for their own research, marketing, or the ultra-vague 'business purposes.' But then, most (84%) of the car brands we researched say they can share your personal data with service providers, data brokers, or other businesses we know little to nothing about. And worse, 19 of the 25, (76%) say they can sell your personal data. A surprising number" - 56% of the total 25 - "also say they can share your information with the government or law enforcement in response to a 'request.'" And, they write, "Not merely a high bar court order, but something as easy as an 'informal request.' A very low bar. Car companies' willingness to share your data is beyond creepy," writes Mozilla. "It has the potential to cause real harm and inspired our worst cars-and-privacy nightmares.

"And keep in mind," they say, "that we only know what companies do with your personal data because of the privacy laws that make it illegal not to disclose that information, such as California's Consumer Privacy Act. So-called anonymized and aggregated data can and probably is shared, too, with vehicle data hubs - who are the data brokers of the auto industry - and others. So while you're getting from point A to point B, you're also funding your car's thriving side-hustle in the data business in more ways than one.

"Next, most" - 92% in their study - "give drivers little to no control over their personal data. All but two of the 25 car brands we reviewed earned our 'ding' for data control, meaning only two car brands, Renault and Dacia, both owned by the same parent company, say that all drivers have the right to have their personal data deleted. None of the others do. While we would like to think this deviation from the norm is one car company taking a stand for drivers' privacy, it's probably no coincidence that these cars are only available in Europe, which is protected by the robust General Data Protection Regulation, the GDPR privacy law. In other words, car brands often do whatever they can legally get away with to your personal data."

They wrote: "We could not confirm whether any of them meet our Minimum Security Standards." They said: "It's so strange to us that dating apps and sex toys publish more detailed security information than cars. Even though the car brands we researched each had several long-winded privacy policies - Toyota wins with 12 - we could not find confirmation that any of the brands meet our Minimum Security Standards. Our main concern is that we can't tell whether any of the cars encrypt the personal information that sits on the car. And that's the bare minimum. We don't call them our 'state-of-the-art security standards,' after all. They're our minimum security standards. We reached out, as we always do, by email to ask for clarity; but most of the car companies completely ignored us. Those who at least responded (Mercedes-Benz, Honda, and technically Ford) still didn't completely answer our basic security questions.

"A failure to properly address cybersecurity might explain their frankly embarrassing security and privacy track records. We only looked at the last three years, but still found

plenty to go on with 17 (68%) of the car brands earning the 'bad track record' ding for leaks, hacks, and breaches that threatened their drivers' privacy."

Okay. So then in the article they provide a car-by-car table of these transgressions, several columns of classification of problems by 25 rows for each of the car brands. But frankly the table's not worth examining. They're all really bad. I think that, as I said, I think we're seeing a classic case of oversight. This is a recently emerged feature category that no one has yet really focused on. And, boy, I really do hope that the privacy people take a look at this and say, whoa, what? It's still the Wild West out there.

So they had then a few additional points. They said: "Tesla is only the second product we have ever reviewed to receive all of our privacy 'dings.' The first was an AI chatbot that we reviewed earlier this year." They said: "What set them apart was earning the 'untrustworthy AI' ding. Tesla's AI-powered autopilot was reportedly involved in 17 deaths and 736 crashes and is currently the subject of multiple government investigations." So, ouch.

"Nissan earned its second-to-last spot for collecting some of the creepiest categories of data we have ever seen." They wrote: "It's worth reading the review in full, but you should know it includes your 'sexual activity.' Not to be outdone, Kia also mentions they can collect information about your 'sex life' in their privacy policy. Oh, and six car companies say they can collect your 'genetic information' or 'genetic characteristics.'" They said: "Yes, reading car privacy policies is a scary endeavor."

Okay. Now, I'll just interject here to suggest that the fact that it can be done doesn't mean that it is being done or has ever been done. These sorts of statements in privacy policies feel like overly broad policies that arise after some wingnut brings an unfounded lawsuit against an automaker. You know, the firm's attorneys will then overreact by adding a clause stating for example they cannot be held responsible for anything that happens if you're picked up by space aliens while operating their motor vehicle. This is not meant to suggest that those things will happen if you're abducted. They're just saying, if they should, then don't go suing us because the policy you already agreed to by driving our car says it's not our fault if something happens. So, you know.

And with regard to references to sexual activity and sex life, that could refer to the car's GPS recording, that GPS is recording where you're going. And so if it's used after the fact to infer something about the driver based upon when they went where, well, then again they've included a broad exclusion because of some past lawsuit that they suffered. So that's probably what that is about. I hope.

Mozilla also said: "None of the car brands use language that meets Mozilla's privacy standard about sharing information with the government or law enforcement, but Hyundai goes above and beyond. In their privacy policy, it says they'll comply with 'lawful requests, whether formal or informal.' All of the car brands on this list except for Tesla, Renault, and Dacia signed on to a list of Consumer Protection Principles from the U.S. automotive industry group Alliance for Automotive Innovation, Inc. The list includes great privacy-preserving principles such as 'data minimization,' 'transparency,' and 'choice.'

"But how many of the car brands actually follow these principles? Zero. It's interesting if only because it means the car companies do clearly know what they should be doing to respect your privacy, even though they absolutely don't do it. This is usually where we'd encourage you to read our reviews, and to choose the products you can trust when you can. But unfortunately cars aren't really like that. Sure, there are some steps you can take to protect more of your privacy, and we've listed them all in each of our reviews under 'Tips to protect yourself.' They're definitely worth doing. You can also avoid using your car's app or limit its permissions on your phone.

"But compared to all the data collection you can't control, these steps feel like tiny drops in a massive bucket. Plus, you deserve to benefit from all the features you pay for without also having to give up your privacy." And they finish: "We spent over 600 hours researching the car brands' privacy practices. That's three times as much time per product than we normally spend. Even so, we were left with so many questions. None of the privacy policies promise a full picture of how your data is used and shared. If three privacy researchers" - that's how many they had on this project - "can barely get to the bottom of what's going on with cars, how does the average time-pressed person stand a chance?"

JASON: No kidding.

Steve: "Many people have lifestyles that require driving. So unlike a smart faucet or voice assistant, you don't have the same freedom to opt out of the whole thing and not drive a car at all. We've talked about the murky ways that companies can manipulate your consent. And car companies are no exception. Often they ignore your consent. Sometimes they assume it. Car companies do that by assuming that you have read and agreed to their policies before you step foot in their cars. Subaru's privacy policy even says that the passengers of a car that used connected services have 'consented' to allow them to use - and maybe even sell - their personal information just by being in the car. So when car companies say they have your 'consent' or won't do something 'without your consent,' it often means what it should."

So as I said, I think this is an area that has until now escaped oversight. Hopefully, research like this which puts the problem squarely on the map will eventually help that to happen. And, you know, wow. We are driving around inside of connected computers, Jason; and they do indeed have sensors galore. They're connected. They know where we are. They know what time it is. They know everything we do with our computerized entertainment systems, what stations we're listening to. And you could imagine.

JASON: Filled with cameras.

Steve: Are they not going to monetize that? Why would we imagine they would not monetize that?

JASON: Oh, of course. Yeah, I mean, none of this is that, I mean, it's shocking, but it's also unsurprising; right?

Steve: I could say it's saddening. I don't know that it's shock - you know, just it's saddening.

JASON: Yeah, yeah, yeah. It is saddening. But at the same time like we live in the data economy. And man, at this point I just feel so beaten down about, like, how my data is used all over the place. Like it would be easy for me to be, like, well, but I carry around a smartphone, and it does a lot of those things. And apparently I've said that that's okay because I still have a smartphone and have for years. But, I mean, you're right, vehicles, you know, have the potential of having cars inside that can be monitored potentially. And that's just one example. And it's always following you wherever you go. So that data has value. That's the unfortunate reality.

Steve: Yeah. And I guess it seems to me that the minimum we could request is the ability for transparency, to know exactly what data is being collected, and then the ability to ask for it to be deleted. And, you know, true. Most drivers never will. They're not listening to this podcast. They're just not concerned. They go, oh, yeah, well, whatever. But the good news is there is legislation which is moving in this direction, which gives

consumers control if they want it. And at this point the automobile industry is way behind on that score, obviously.

Okay. So we've got some neat feedback from our listeners. Someone whose name, I know him from years of transacting with him, his handle is @ramriot. But his name, he's named himself in Twitter "418," which of course is one of the error codes that can be returned by HTTP. It's about, I think it's Are You a Teapot? Anyway, so his is "418: Tea Ready?"

Anyway, he said: "Hi, Steve. This DOM" - meaning the document object model for our web browsers. "This DOM issue is a tough but old nut, raised again in connection to extensions," which we talked about last week. "Would this be an opportunity for browser vendors to tighten up the Same-Origin rules for access to form fields? You know, perhaps make them write-only, immutable objects when accessed cross-origin?"

And so great point and question. I strongly suspect that the real problem at this point that we face from a practical standpoint is breakage of the already existing quite rich browser extension ecosystem, not to mention the loss of third-party password managers, which do have to poke into every website's forms in order to do their work. But even more broadly, there would be breakage that would result from any further tightening of access by extensions. We already saw the uproar that Chrome's rather modest move to the V3 Manifest caused. And the trouble we were talking about last week was after this move. So these are things that you can still do, even under the V3 Manifest. So things remain extremely permissive today.

Just think of the degree to which we must trust today's browser. And that's - this is really sobering, when you think about it. The degree to which we must trust today's browser, the browser itself. Through it passes everything. Nearly our entire interaction with the world today is through our chosen browser. Interactive applications have moved or are moving from desktop applications to browser-hosted apps.

All of our usernames and our passwords and the private information we fill out as we interact with anything - the IRS or credit bureaus or loan applications or our doctors' offices or dating sites or any retailer - everything we do today passes through our browser. And now we're reminded that, if our password managers are able to see everything we do, then so are other extensions which we might trust far less. Yet here they sit, watching, because they provide some little browser doodle that we like, and we don't want to now live without.

We started off without much concern for browser extension security and privacy many years ago. But now that we feel we need more security and privacy, it's difficult to take it back without sacrificing the rich feature set and environment that these extensions provide to us. Both Firefox and Chrome are aware of this problem, which is why both of them allow their users to decide which extensions should be allowed to follow them into the browser's private viewing mode.

And depending upon how extension-laden any user's normal browsing is, it might be worthwhile to consider trimming back on the extensions that are allowed to run, well, I would argue first normally, but also specifically, in Chrome's Incognito mode or browsers' private window modes. So you're able to then switch into there and have many fewer things watching what you do. It's not just the browser that's watching. And as I've said, we have to utterly and absolutely trust the browser because it sees everything we do. Turns out extensions, which we may trust far less than the browser, who knows where they came from, are able to see what's going on, too.

The alternative, and I've heard that some of our listeners are doing this, would be to reserve a secondary browser, because we certainly don't lack for choice of browsers

today, Everyone can use Firefox or Chrome. So reserve a secondary browser which is running, you know, maybe only your password manager and nothing else. Obviously we're trusting our password manager a lot also. So have a browser that only does that, so it's able to log you in places, but then you don't have any other perhaps sketchy extensions that are doing things. You may not want to live there without all your extensions, but that's where you might want to go when, you know, you're doing something that is much more confidential.

In any event, we do currently have a problem that's going to require some eventual resolution because right now, as we noted last week, the need to trust every extension with everything we do, just like we do with our password manager and our browser, that's a problem.

Someone posting as "person typing #22," I guess he feels he's rather generic, he said: "Hey again, Steve. In last week's SN-938 the tradeoff between security and convenience was mentioned with respect to websites and browser extensions like password managers. I figured it was worth mentioning that on the Mac and on iOS, I use Apple's universal AutoFill with a compatible password manager." And he says: "1Password is an example, and I use KeePassium." He said: "For most browsing, I use Firefox. But to log into banking and similar sites I use Safari on the Mac without ANY extensions. The OS itself recognizes websites' password fields and allows me to choose a password to autofill from my password manager. I feel like this provides lots of both security and convenience."

And I completely agree, that is a great solution. It's sort of a solution using the dual-browser approach, but it also strengthens the isolation from even the password manager by interposing an OS that's as secure as Apple has been able to design, any OS, into the path. I think that's very nice.

And I had a great comment from a guy whose name I don't know because he uses what is now we're calling X, you know, what used to be known as Twitter, so infrequently that it wouldn't let him log in. So he posted from his wife's account. Anyway, whoever he is, he said: "Hello, Steve. Thank you for the many great shows. I've been listening since Episode 1 and was overjoyed to hear you're not stopping at 999.

"With regard to the web extension security research story, I helped develop a web application used internally by the majority of banks in the U.S. A few years ago we implemented Content Security Policy (CSP) headers. CSP has a wonderful feature where all violations can be reported to the website to help fix bad rules. During the initial rollout we reviewed the violations and found that javascript was being injected into our sites by browser extensions. A few of these extensions seemed to have questionable intentions and were likely installed by adware. I do not think the research paper's solution of adding a secure input element or alerting the user of nefarious activities is adequate since an extension can alter the source before the DOM is rendered, and therefore could strip out these protections.

"A website can try their best to obfuscate input and output; but at the end of the day, a browser extension can access or modify headers (including cookies), requests, and responses. It is an ideal position for a man-in-the middle attack while the user thinks their connection is secure and private. Maybe something similar to CSP or HSTS where a site with sensitive information could request the browser to disable all browser extensions could help protect users. Of course, sites with advertisements would quickly abuse this power to block good extensions like uBlock Origin, so maybe this would just add complexity to an already impossible problem. Where I work," he said, "browser extensions have been disabled. It is annoying that many useful tools are blocked, and yet I cannot argue with their decision."

So wow. You know, this guy's workplace said nope, sorry, there's just too much danger there. All browser extensions are disabled. And, you know, it would be a pain not to have the benefit of a built-in password manager. But on the other hand, you know, browsers are now universally offering their own built-in password managers, so he's not without autofill.

What was interesting was that the moment I heard this listener talk about some means to allow websites to force-disable extensions, I was reminded that exactly such a proposal was floated through the industry, I'm not exactly sure when, like a month ago. I don't recall the details, and I don't think we discussed it here. But I believe I remember Leo, Stacy, and Jeff discussing it on This Week in Google. And I also recall that many naysayers were suggesting that this was a just a slimy way of disabling ad blockers. So, you know, right. To this listener's point, this is all a mess.

Anthony Bosio, he said: "I think Topics might" - meaning, you know, Google's Topics solution for profiling, providing some information about what people are currently interested in. He said: "I think Topics might be DOA. People are interpreting it and spreading it basically as 'shares your browser history with other sites.'"

Okay. So let's hope that this is just the initial uninformed reaction to anything that's new. Given Chrome's massive market power, any technology Google creates and enforces by virtue of foreclosing on all alternatives - which they have said they're going to do next year - is going to succeed because there won't be any alternative to using it. So I believe Topics cannot be DOA if Google doesn't want it to be. Then to that we add that Topics is also an extremely benign, non-tracking, privacy-enforcing system. And I expect that, while it may take some time for the less technical types to catch up and to understand it, it's where the entire industry is going to go. And to that I say yay.

Barbara said: "Some downloadable software basically is a stub that phones home to download the rest while installing. I don't think giving the stub to VirusTotal would be helpful." And of course she's referring to our previous conversation about sending things to VirusTotal and having it, you know, sending things you download to VirusTotal and having it check on them before you trust them. And I think Barbara makes a very good point. As we know, not all software we download is in the entire package. We're now often seeing a much smaller "installer" that immediately connects back to home base to download the entire package, which is often many modules deep. The promise is that you select the things you want, and it only downloads those things that you have said you want and intend to use. But again, it doesn't give you a chance to check all of them against VirusTotal.

Someone tweeting from the handle Skynet said: "Hi, Steve. Regarding Martin's 'duh' about VirusTotal being served ostensibly 'clean' files from a malicious source, how would such a site even know who or when they would be doing this to know to send a clean executable? Do websites even actively monitor who's downloading their content? And, if so, wouldn't they have to time it so as to know when to give VirusTotal a clean one instead of a malicious one?" He says: "I don't get where Martin is getting this idea from. You'd have to be checking logs of IP addresses; wouldn't you? And by the time they'd discovered that, 'Oh, look, VirusTotal is trying to get one of our most malicious executables, quick, give them this one instead?' I don't see how it works. Even with some redirect link I'd think it would be too late to detect that it was VirusTotal asking for the file; no? Am I being a doofus for missing a big duh? Please explain."

Okay. So no one's being a doofus here. When I created ShieldsUP! 24 years ago, back in 1999, it was because I knew that the IP address of anyone connecting to my web server was immediately known to the server. So I was able to return custom webpage results based upon the security I had detected at their connecting IP. So it would, in fact, be simple for the IP address blocks which have been assigned to known security researchers

and VirusTotal to cause different 'clean' software to be delivered on demand. And we've seen other non-web server examples of this where malware is actually aware of the IP addresses of researchers and acts differently and, like, changes its behavior to be non-malicious when it realizes that known researchers are downloading it or examining it. So, yeah, this actually does work and can happen.

E. Remington says: "Hi, Steve. One email provider people overlook is iCloud. You can set up your own domain." Which, by the way, I had no idea of. He said: "While iCloud limits you to," he says, "if I remember right, five email addresses, by using a form like 'something plus emailaddress@yourdomain.com," you can have an infinite number of email addresses." And, he says, RFC 822, which is like the original email, and all of its updates have supported the idea of using the + symbol added to a tag in order to differentiate it, you know, basically that the tag is ignored, and it goes into your main ID for the domain. So anyway, I'm glad that Remington mentioned iCloud since he's correct that iCloud as an email provider is easily overlooked. And, you know, certainly they are reputable.

Magnify247 said: "Steve, with Windows 12 being prepped for 2024, will InControl be updated, or will the current version allow for the version and release application to be locked accordingly?"

Well, you know the old expression about "fool me once." As we know, the predecessor to "InControl" was "Never10." I would have named the next one "Never11" except, after Microsoft changed what was clearly their original intention for Windows 10 to be the last Windows ever - which everyone recalls, even if now Microsoft claims it was never what they said - I decided that I had to drop any major version numbering from the utility. So "InControl" gets to live on without any further name changes.

And since it's all about controlling just a few registry keys which Microsoft officially supports, it should keep working, as long as Microsoft honors those settings. And since their enterprise users depend upon those, I can't see anything changes. So I think when Windows 12 occurs next year, InControl will probably continue to work. And of course, if not, if something changes, I'll update it. But I don't expect that I'm going to have to.

Christian P. said: "Hi, Steve. Just an observation about the concern from the user on the last Security Now! podcast about testing a file directly from VirusTotal, and the risk of the file being swapped based on the source of the test. You can remove any risk of testing the file directly by getting VirusTotal to download the file, then recheck the hash before you execute it. VirusTotal reports the SHA-256, and the hash is also in the URL of the result.

"So perhaps a sensible process," he says, "would be to get VirusTotal to download and check. Then, if that looks largely okay, then download directly to your system and test again with VirusTotal. The second test should take you to the same page. VirusTotal doesn't automatically re-test files that have already been submitted. It just recomputes the hash and looks up the last submitted report. You do have the option to reanalyze, but there is little point if the hash is the same. Perhaps if a new scanning has been added since the last test. Anyway," he says, "great podcast, et cetera."

Okay. So I wanted to share this since the way the world is evolving, keeping VirusTotal in one's back pocket I think makes a lot of sense. Christian is right about the way VirusTotal operates. Before it does anything else, it first calculates the SHA-256 hash of the file that it either downloaded or that the user submitted. It then checks to see whether that file's hash already exists in its known library of previously scanned files. And if so, it just returns the previous result. No need to rescan since it already did. And the matching SHA-256 signature is absolute proof that the file has not changed from the one that it previously scanned.

And Christian's suggestion of then uploading your own copy of the hopefully identical file that VirusTotal first approved of to see whether they indeed match makes sense. I wanted to also note that Windows has for a while now had a built-in "hashfile" command which is an actual subcommand of the "certutil" which also allows a user to quickly and easily generate an SHA-256 hash of any file they may wish to check on their own. So you open a command prompt and say certutil, C-E-R-T-U-T-I-L, space hyphen hashfile space, then the path to the file, then space SHA256 and hit ENTER. And you'll immediately get the SHA-256 hash of the file which you can then manually check if you wish.

And lastly, InspClousseau says: "@SGgrc Steve, what DNS services do you recommend for children under 10 to avoid unsafe and unsuitable sites?" I think that Leo uses and recommends OpenDNS. They are definitely reputable, they've been acquired by Cisco, and they have a free family use tier which they refer to as their "Family Shield" service. Using it is as easy as configuring the family's router to use OpenDNS's servers at two very specific IP addresses: 208.67.222.123 and 208.67.220.123. Once you've done that, you can go to welcome.opendns.com, which will confirm that you're now using their filtered DNS.

And, you know, I suppose if you wanted a bit more proof, you know, that it was working, you could also try going over to Pornhub and see whether that works. And I would expect you will not be able to go there when you are using the family-safe Family Shield service. So anyway, a little quickie, and I'm glad that InspClousseau thought to ask.

So Jason, let's share our last sponsor piece and then, wow, spill the beans.

JASON: No.

Steve: Speaking of, yeah, when LastPass finally really screwed this, and we had the podcast "Leaving LastPass," which was where I finally said, okay, there's just no more excuse for this. You know, a lot of people were nervous about what that meant. And their nervousness was based on or should have been on their password and the number of iterations that LastPass had been using for them. And we know that unfortunately not everybody was set to the iteration count of 100,100, which was where we last left it when we last visited this issue for LastPass. Many people were set to one. Some were set to 500. Some were set to 5,000. It turns out that there is pretty convincing evidence now, which is what we're going to share, that the data that was stolen is being decrypted and is being used to hurt hopefully previous LastPass users.

JASON: Yeah. Wow. All right. Well, we are going to, like I said, spill the beans on all of this. And, you know, a couple of spit takes. Oh, my goodness. That's coming up next. The LastPass saga, it continues. And I feel like every time we check in on it, it's worse than it was before. Really happy I'm not with LastPass anymore, but I'm super curious to hear all about this.

Steve: Nothing I'm going to say is going to disabuse you of that concern, Jason.

JASON: Indeed.

Steve: Any regular listener of this podcast can probably guess that today's title of, as I said, LastMess, will have something to do with LastPass. So there's growing significant circumstantial evidence which, under the circumstances, is probably the only sort of evidence anyone would ever be able to obtain, which suggests that the encrypted LastPass Vault data, which LastPass had been storing for its many users, and which they famously had exfiltrated from their backup location, is now being and has successfully been decrypted. I don't mean all of it, I mean incrementally, but that's bad enough; right? And it's being used by those unknown cyber assailants.

Brian Krebs reported the news of this last Tuesday on his KrebsOnSecurity site under the title "Experts Fear Crooks Are Cracking Keys Stolen in LastPass Breach." So here is some of what Brian reported. He wrote: "In November of 2022, the password manager service LastPass disclosed a breach in which hackers stole password vaults containing both encrypted and plaintext data for more than 25 million users. Since then, a steady trickle of six-figure cryptocurrency heists targeting security-conscious people throughout the tech industry has led some security experts to conclude that crooks likely have succeeded at cracking open some of the stolen LastPass vaults.

"Taylor Monahan is lead product manager of MetaMask, a popular software cryptocurrency wallet used to interact with the Ethereum blockchain. Since late December 2022, Monahan and other researchers have identified a highly reliable set of clues that they say connect recent thefts targeting more than 150 people. Collectively, these individuals have been robbed of more than $35 million worth of their cryptocurrency.

"Monahan said virtually all of the victims she has assisted were longtime cryptocurrency investors and security-minded individuals. Importantly, none appeared to have suffered the sorts of attacks that typically preface a high-dollar crypto heist, such as the compromise of one's email and/or mobile phone accounts.

"Monahan wrote: 'The victim profile remains the most striking thing. They truly all are reasonably secure. They are also deeply integrated into this ecosystem, including employees of reputable crypto orgs, VCs (venture capitalists), people who built DeFi protocols, deploy contracts, and run full nodes.'

"Monahan has been documenting the crypto thefts via Twitter (now 'X') since March of 2023, frequently expressing frustration in the search for a common cause among these victims. Then on August 28th Monahan said she'd concluded that the common thread among nearly every victim was that they'd previously used LastPass to store their 'seed phrase,' the private key needed to unlock access to their cryptocurrency investments."

Brian Krebs included a screenshot in his coverage of Taylor's tweets. On August 28th she tweeted: "The diversity of key types drained is remarkable: 12 and 24 word seeds generated via all types of hardware and software wallets. Ethereum Presale wallet JSONs. Wallet.dats. Private key generated via MEW and others."

And she also noted that the diversity of the chains and coins which had been drained was striking. So it wasn't as if there was a fault in any specific chain or crypto contract that had been exploited, or service. There was no common denominator - well, except for LastPass.

Again, Brian writes: "Armed with your secret seed phrase, anyone can instantly access all of the cryptocurrency holdings tied to that cryptographic key and move the funds anywhere they like. This is why the best practice for many cybersecurity enthusiasts has long been to store their seed phrases either in some type of encrypted container, such as a password manager; or else inside an offline, special-purpose hardware encryption device such as a Trezor or Ledger wallet.

"Nick Bax, director of analytics at Unciphered, a cryptocurrency wallet recovery company, said: 'The seed phrase is literally the money. If you have my seed phrase, you can copy and paste that into your wallet. Then you can see all my accounts, and you can transfer my funds.'

"Bax said he closely reviewed the massive trove of cryptocurrency theft data that Taylor Monahan and others had collected and linked together. He said: 'It's one of the broadest and most complex cryptocurrency investigations I've ever seen. I ran my own analysis on

top of their data and reached the same conclusion that Taylor reported. The threat actor moved stolen funds from multiple victims to the same blockchain addresses, making it possible to strongly link those victims.'

"Bax, Monahan and others interviewed for this story say they've identified a unique signature that links the theft of more than $35 million in crypto from more than 150 confirmed victims, with between two and five high-dollar heists happening each month since December 2022." So in other words, it's not all at once. It's even following the pattern of brute-forcing something, and the common link of somethings is LastPass.

Anyway, Brian says: "The researchers have published findings about the dramatic similarities in the ways that victim funds were stolen and laundered through specific cryptocurrency exchanges. They also learned the attackers frequently grouped together victims by sending their cryptocurrencies to the same destination crypto wallet. By identifying points of overlap in these destination addresses, the researchers were then able to track down and interview new victims. For example, the researchers said their methodology identified a recent multi-million dollar crypto heist victim as an employee at Chainalysis, a blockchain analysis firm that works closely with law enforcement agencies to help track down cybercriminals and money launderers."

Okay. Just to make sure everyone is following this, based on what they were seeing, they followed victims to the bad guys. Then they looked at all of the transactions on the bad guys' wallets that identified new victims. Then they looked up the known victims' wallets and were able to find, for example, in this case essentially new victims, in this case an employee at Chainalysis. Then they went to Chainalysis and said, hey, has anybody had a problem?

Brian writes: "Chainalysis confirmed that the employee had indeed suffered a high-dollar cryptocurrency heist late last month, but otherwise declined to comment further. Bax said the only obvious commonality between the victims who agreed to be interviewed was that they had all stored the seed phrases for their cryptocurrency wallets in LastPass."

Bax told Brian Krebs: "On top of the overlapping indicators of compromise, there are more circumstantial behavioral patterns and tradecraft which are also consistent between different thefts and support this conclusion. I'm confident enough," he said, "that this is a real problem that I've been urging my family and friends who use LastPass to change all of their passwords and migrate any crypto that may have been exposed, despite knowing full well how tedious that is."

Brian Krebs asked LastPass for any comment, about which Brian wrote. He said: "LastPass declined to answer questions about the research highlighted in this story, citing an ongoing law enforcement investigation and pending litigation against the company in response to its 2022 data breach."

Yup. That's the standard dodge, of course. LastPass said in a written statement to Brian: "Last year's incident remains the subject of an ongoing investigation by law enforcement and is also the subject of pending litigation." Uh-huh. Perhaps some additional litigation now.

Anyway, they continued: "Since last year's attack on LastPass, we have remained in contact with law enforcement and continue to do so. We've shared various technical information, Indicators of Compromise, and threat actor tactics, techniques, and procedures with our law enforcement contacts, as well as our internal and external threat intelligence and forensic partners in an effort to try and help identify parties responsible. In the meantime, we encourage any security researchers to share any useful information

they believe they may have with our Threat Intelligence team by contacting securitydisclosure@lastpass.com."

So Brian's reporting then covers for his readers everything that this podcast already covered for our listeners back at the time, you know, things like how crucial the PBKDF iteration count is for increasing the difficulty of cracking the user's password by brute force. How the early LastPass users may have originally had iteration counts of 1 or 500, and how despite LastPass moving the defaults upward over time as necessary to keep ahead of brute force cracking capabilities, for reasons that no one has ever explained, many of the original much-too-low original defaults remained in place.

Nicholas Weaver, a researcher at University of California Berkeley, their International Science Institute, and he also lectures at UC Davis, said about brute force attacks that: "You just crunch and crunch and crunch with GPUs, with a priority list of targets that you target." He said that a password or passphrase with average complexity, such as "Correct Horse Battery Staple," is only secure against online attacks, and that its roughly 40 bits of entropy means that a graphics card can blow through it in no time.

An Nvidia 3090 can do roughly four million password guesses per second with an iteration count of 1,000. But that would go down to 8,000 per second with 500,000 iterations, which is why he says iteration counts matter so much. So a combination of 'not that strong of a password' and an 'old vault with a low iteration count' would make it theoretically crackable. It would take real work, but the work is worth it given the high value of the targets.

And here's something else that Brian reported which is very interesting. Brian interviewed one of the victims tracked down by Monahan. This person is a software engineer and a startup founder who was recently robbed of approximately $3.4 million worth of different cryptocurrencies. This engineer agreed to tell his story in exchange for anonymity because he's still trying to claw back his losses. Good luck. For his reporting, Brian refers to this person as "Connor," which is not his real name.

So Brian writes: "Connor said he began using LastPass roughly a decade ago, and that he also stored the seed passphrase for his primary cryptocurrency wallet inside LastPass. Connor chose to protect his LastPass vault with an eight-character master password that included numbers and symbols." Okay, so, maybe around 50 bits of entropy. Connor said: "I thought at the time that the bigger risk was losing a piece of paper with my seed phrase on it. I had it in a bank security deposit vault before that, but then I started thinking, 'Hey, the bank might close or burn down, and I could lose my seed phrase.'

"Those seed phrases sat in his LastPass vault for years. Then, early on the morning of Sunday, August 17th, 2023, Connor was awoken" - meaning just recently, right, a couple weeks ago - "Connor was awoken by a service he'd set up to monitor his cryptocurrency addresses for any unusual activity. Someone was draining funds from his accounts, fast.

"Like other victims interviewed for this story, Connor didn't suffer the usual indignities that typically presage a cryptocurrency robbery, such as account takeovers of his email inbox or mobile phone number. Connor said he doesn't know the number of iterations his master password was given originally, or what it was set at when the LastPass user vault data was stolen last year. But he said he recently logged into his LastPass account, and the system forced him to upgrade to the new 600,000 iterations setting." Which we know as too little, too late.

"He said: 'Because I set up my LastPass account so early, I'm pretty sure I had whatever weak settings or iterations it originally had.' Connor said he's kicking himself because he recently started the process of migrating his cryptocurrency to a new wallet protected by a new seed phrase. But he never finished that migration process, and then he got

hacked. He said: 'I had set up a brand new wallet with new keys. And I had that wallet ready to go two months ago, but had been procrastinating moving things to the new wallet.'" And I thank Connor for his honesty.

JASON: No kidding.

Steve: Wow. "Nicholas Weaver, the UC Berkeley researcher, said what we all know, which is that LastPass deserves blame for not having upgraded iteration counts for all users a long time ago, and called LastPass's latest forced updates 'a stunning indictment of the negligence on the part of LastPass.'" Meaning they could and should have done this years and years ago.

"He said: 'That they never even notified all those with iteration counts of less than 100,000, who are really vulnerable to brute force even with eight-character random passwords or "correct horse battery staple" type passphrases, is outright negligence.' He said: 'I would personally advocate that nobody ever uses LastPass again, not because they were hacked, not because they had an architecture that makes such hacking a problem, but because of their consistent refusal to address how they screwed up and take proactive efforts to protect their customers.'

"Bax and Monahan both acknowledged that their research alone can probably never conclusively tie dozens of high-dollar crypto heists over the past year to the LastPass breach. But Bax says at this point he doesn't see any other possible explanation. He said: 'Some might say it's dangerous to assert a strong connection here, but I'd say it's dangerous to assert there isn't one. I was arguing with my fianc about this last night,' he said. 'She's waiting for LastPass to tell her to change everything. Meanwhile, I'm telling her to do it now.'"

So, yeah. Based upon our own experience with LastPass, anyone who is waiting for them to do anything like take responsibility, which might open them to additional litigation, is probably going to be waiting for quite a while. As for all of our listeners whose LastPass Vaults were exfiltrated back at the time of the heist, I'm fairly certain that most of us likely have little to fear. The bad guys obtained a massive treasure trove of encrypted information for more than 25 million individual users. But decrypting it, while, yes, technically feasible on a one-by-one instance, depending upon passphrase length and iteration count, is still a time-consuming and massive undertaking. So attacks on the vault repository are going to be highly targeted. They want money, plain and simple. They're not going to waste the cost of decrypting someone's vault unless they're fairly certain that a cash equivalent pot of gold awaits them if they are successful.

Anyone whose iteration count was high - 100,100, which was what we last told all of our listeners to switch to when we talked about it years prior to all of this, or even greater than that, rather than the 1 that it started at, the 500 that it moved to, or the 5,000 that it moved to, all without it ever being done proactively, retroactively, you know, and anybody whose iteration count was high, you know, we set it to 100,100 the last time we talked about it, years before all this happened, you're almost certainly safe. On the other hand, somebody whose iteration count was one, 500, or 5,000, that's a problem.

Also a high-entropy passphrase, which all of our listeners hopefully had, that would be protection. But now we know that the bad guys didn't just grab these 25 million vaults and say, oh darn, they're all encrypted. No, they're attacking them. They're going after pots of gold. So if by chance you did own cryptocurrency, you do own cryptocurrency, and your key, your passphrase or wallet key was in LastPass, absolutely don't hesitate. Create a new wallet, move the cryptocurrency to the new wallet. It really only looks like these attacks are going to be targeted, and I know that all of our listeners have moved over to Bitwarden now anyway. But remember, it was the contents at the time of the theft that matters.

JASON: Right.

Steve: Not what we did afterwards.

JASON: Right.

Steve: And, you know, and this guy Bax said he's telling his friends, you know, you really need to go change all the passwords that LastPass was storing for you at the time of the theft. A pain in the butt, yes. But, you know. And again, to me, it seemed really unlikely. There's 25 million of us. And if the iteration count was high, and you're just, you know, some random guy, then I don't think you're ever going to get it decrypted. And, you know, there's no money; right?

JASON: Right.

Steve: They want money.

JASON: Right. And the longer it goes not decrypted, the staler that data gets for them. So probably I'm guessing the lower likelihood that they have to target something that they don't know has some sort of a major payload. Don't be like Connor. Don't set half of this stuff up if you think there is really something to save on your end.

Steve: And then don't follow through.

JASON: And then don't follow through. That's just so heartbreaking.

Steve: The other thing is we were just talking about the power of somebody getting your email address because email is how all the password recovery occurs.

JASON: Yeah, oh, yeah.

Steve: So if you do nothing else, change your email password.

JASON: Yeah. Yeah.

Steve: Just because that's what they will go for. They'll go for your email password, get into your email, then start clicking on sites that they know you have access to and start doing the whole account takeover routine.

JASON: Yeah. Indeed. All right. Well, I am happy that I'm no longer with LastPass. It is disappointing that we're still not getting any sort of like real resolute confirmation, slash, you know, just stating that they totally and completely messed up in so many ways. I mean, this is just...

Steve: They're owned by a hedge fund. We're never going to find out.

JASON: Never going to get it.

Steve: They no longer care. You know? The guys we knew and cared about are long gone. They cashed out. It's all about private equity and suck as much money out of their corporate install base as they can.

JASON: Ooh, man, you put it that way. Well, Steve, thank you for diving deep into that. That was fascinating, if not scary. If you're still with LastPass, I think you have yet another reason to get out. Thank you, everybody. We'll see you next time on Security Now!. Bye-bye.

Steve: Bye.