Security Now! #939 - 09-12-23

LastMess

This week on Security Now!

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.

The plumber's contract didn't say anything about moving any rocks.

ValiDrive

I was getting ready to put this little side project to bed and make it available to everyone, when one of our testers found what very much appears to be a fake drive which ValiDrive shows as being 100% good to go. At the start of this work I had purchased a handful of sketchy looking "too good to be true" drives from Amazon and that drive was among those that I purchased. The drive's Amazon listing for $20 claims that it's 1TB, but it comes formatted as a 2TB and that's also the size that GRC's InitDisk utility obtains when it queries the drive at a low level to give it a shiny new format. What's troubling is that ValiDrive is passing this sketchy drive with an "all green" 100% healthy report but file level testing strongly suggests that it only has 64GB of actual storage. So, while all of the foundation work I've done so far remains valid and useful, I apparently have some more work to do. And I'm on it!

Security News

Who Blinked?

What can only be called wonderful and welcome news surfaced in the middle of last week from the UK. The short version is, the UK blinked and in the face of all secure messaging apps – and Apple – firmly stating that they will all pull their services out of the UK rather than compromise their user's privacy and security, the UK said... "Oh! We never said that we wanted you to do that!"

Uh huh. But, of course, nothing involving politicians and bureaucracies is ever clean and simple. And the details here are interesting. What first happened was that last Wednesday the Financial Times (FT) carried the breaking news which the tech press jumped on.

9to5Mac's headline was: "Future of iMessage safe in the UK, as government backs down on encryption". Wired's headline was: "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." So, you get the idea.

Unfortunately, much as those were attention commanding and welcome headlines, none of that was true. Or at least they were all probably deliberate oversimplification and exaggeration click-bait which, predictably, did not sit well with the UK. So the next 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." Here's what Reuters explained:

> *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.*
>
> *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 communications regulator* ***would only require them to scan content where “technically feasible”.***
>
> *Tech companies have said scanning messages and end-to-end encryption are fundamentally incompatible.*

Right. In other words “not technically feasible.” So essentially, by admitting and facing reality, this junior minister Stephen Parkinson set off a fire-storm. Was the fire set deliberately? Did the senior mister set this up to have junior drop this and then hide? Am I being too cynical? Reuters continues with their coverage:

> *[Senior technology minister Michelle] Donelan, however,* ***denied*** *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. If there was a situation where the mitigations that the social media providers are taking are not enough*

(And we already know they won’t be.)

> *... 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.”*

Huh???

> *She said further work to develop the technology was needed, but added that government-funded research had shown it was possible.*

Oh... kay.

So there’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 is desperately trying to obfuscate that fact. Here’s the way Apple Insider explained what happened:

> ***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,*** *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: "We haven't changed the bill at all."* [Wow. She continued:] *"If there was a situation where the mitigations that the social media providers are taking are 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."*
>
> [I don’t think that anyone’s ever going to listen to her or take her seriously again.]
>
> *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.”*
>
> *AppleInsider opines: There is no technology today that will allow only the good guys to break end-to-end encryption — and there never will be.*
>
> *Consequently 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 nonsensical and unenforceable plan.*

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 outspoken president, Meredith Whittaker. Here’s what Wired write of what Meredith had to say:

> *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.”*
>
> *Whittaker acknowledges that “it’s not enough” that the law simply won’t be aggressively enforced. “But it’s major. We can recognize a win without claiming that this is the final victory,” she says.*
>
> [Wired continues:] *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.”*

I believe this is authentically a big 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 this was the strongest legislation they were able to obtain.

We know that this will in no way pacify Sarah Gardner, whom we talked about last week, after she threatened Apple's Tim Cook with her forthcoming pressure campaign 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.

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 to see what the European Union does now.

How Microsoft lost the key to their kingdom

As we know, a China-based attacker, known as Storm-0558, somehow managed to acquire one of Microsoft's secret keys which allowed it to forge login tokens which it then used to access the private eMail of OWA and Outlook.com users. In the wake of these revelations the entire security world was left wondering exactly how Microsoft had managed to fumble the crucial protection of this important key. Last Wednesday, Microsoft provided what may be the conclusion of their investigation... and it did answer some of these questions. Here's what Microsoft shared:

> *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 multi-factor 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.*
>
> *Our corporate environment, which also requires secure authentication and secure devices, allows for email, conferencing, web research and other collaboration tools. 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.*
>
> *Our investigation found that a consumer signing system crash in April of 2021 resulted in a snapshot of the crashed process ("crash dump"). The crash dumps, which redact sensitive information, should not include the signing key. In this case, a race condition allowed the key to be present in the crash dump (**this issue** has been corrected). The key material's presence in the crash dump was not detected by our systems (**this issue** has been corrected).*
>
> *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. This is consistent with our standard debugging*

> *processes. Our credential scanning methods did not detect its presence (**this issue** has been corrected).*

So… three strikes and you're out!

> *After April 2021, 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?

> *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 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 (**this issue** 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 (**this issue** has been corrected using the updated libraries).*
>
> *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 https://aka.ms/storm-0558 blog. Items detailed in this blog are a subset of these overall investments. We are summarizing the improvements specific to these findings here for clarity:*
>
> - *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*
> - *Released enhanced libraries to automate key scope validation in authentication libraries, and clarified related documentation*

Okay. So as Microsoft has explained this mess-up, a series of five separate and previously undiscovered bugs, all which they have since found and fixed, were responsible for allowing a key, which should never have left Microsoft, to be exfiltrated by Chinese attackers and then used to remotely compromise what should have been highly secure enterprise eMail.

It’s obvious that the key should have never been allowed to leave Microsoft. But what they appear to have conveniently skipped over is why that key ever left the HSM – 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 cause **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 GRC’s code signing keys are sequestered in hardware, completely inaccessible to the outside world once installed and only able to be used to sign signature hashes.

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.

The car I drive is 19 years old – and I’m more glad than ever!

I titled this next 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. It’s a car. It does what a car is supposed to do. It’s 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 conducted and published last Wednesday by The Mozilla foundation which they titled: *“It’s Official: Cars Are the Worst Product Category We Have Ever Reviewed for Privacy.”* A subheading might be: “25 car brands tested and 25 car brands failed.” Here’s what their research uncovered. I’ve edited their posting for the podcast:

> *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.*
>
> *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. While we worried that our doorbells and watches that connect 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 one of them.*
>
> *We reviewed 25 car brands in our research and we handed out 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) 63% of the mental health apps (another product category that stinks at privacy) we reviewed this year received this “ding.” But it was 100% for automobiles!*

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 the 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" (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, abilities, and interests.

AND (get this) Most (84%) share or sell that data!

*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, and other businesses we know little or nothing about.*

*Worse, nineteen of the 25, (so 76%) say they can **sell** your personal data.*

*A surprising number **(56%)** also say they can share your information with the government or law enforcement in response to a "request."*

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. It has the potential to cause real harm and inspired our worst cars-and-privacy nightmares.

And keep in mind that we only know what companies do with 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 A to B, you're also funding your car's thriving side-hustle in the data business in more ways than one.

Next, most (92%) 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 other 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 (GDPR) privacy law.

> *In other words: car brands often do whatever they can legally get away with to your personal data.*
>
> ***We could not confirm whether any of them meet our Minimum Security Standards***
>
> *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 couldn't 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 are 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.*

They provide a car-by-car table of transgressions. But it's really not worth examining. They're all really bad. I think what we're seeing is the classic case of oversight. This is a recently emerged feature category that no one has yet really focused upon. So once again it's the Wild West.

Here are a few additional points:

> ***Tesla is only the second product we have ever reviewed to receive all*** *of our privacy "dings." (The first was an AI chatbot we reviewed earlier this year.) 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.*
>
> ***Nissan earned its second-to-last spot for collecting some of the creepiest categories of data we have ever seen.*** *It's worth reading the review in full, but you should know it includes your "sexual activity." Not to be out done, 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." Yes, reading car privacy policies is a scary endeavor.*

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 feel like overly broad policies that arise after some wingnut brings an unfounded lawsuit against an automaker. The firm's attorneys will then overreact by adding a clause stating that they cannot be held responsible for any anal probes being conducted by space aliens while operating their motor vehicle. This is not meant to suggest that anal probes will be performed, only that if they should be, don't go suing us since by driving this car you've already agreed that if it does happen it's not our fault.

With regard to references to sexual activity and sex life, that might refer to the car's GPS recording being used after the fact to infer something about the driver based upon when they went where. Again, an overly broad exclusion that is easily misinterpreted. 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 will comply with "lawful requests, whether formal or informal."* [Mozilla notes that's a serious red flag.]
>
> *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 US 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 that 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, 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.*
>
> *(Since many of the apps share a privacy policy with the vehicle, we can't always tell which data is taken from your phone so it's probably better to err on the side of caution by not using it.)*
>
> *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.*
>
> *We spent over 600 hours researching the car brands' privacy practices. That's three times as much time per product than we normally do. Even still, 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 can barely get to the bottom of what's going on with cars, how does the average time-pressed person stand a chance?*
>
> *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. We've talked before 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 says that even* ***the passengers*** *of a car that uses 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 doesn't mean what it should.*

As I said, I think this is just an area that has, until now, escaped oversight. Hopefully, research like this, which puts the problem squarely on the map, will help that to happen.

Closing the Loop

418: Tea Ready? / @ramriot

> *Hi Steve, This DOM issue is a tough but old nut, raised again in connection to Extensions. Would this be an opportunity for browser vendors to tighten Same-Origin rules for access to form fields? Perhaps make them write only, Immutable objects when accessed Cross Origin?*

I strongly suspect that the real problem at this point is breakage of the already existing quite rich browser extension ecosystem – not to mention the loss of 3rd-party password managers – 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 things remain extremely permissive.

Think of the degree to which we must trust today's browser. 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 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 there they sit... watching... because they provide some little browser doodle that we like and don't want to live without.

We started off without much concern for browser extension security and privacy. But now that we feel we need more security and privacy it's difficult to take it back without sacrificing the feature rich environment that extensions provide.

Both Firefox and Chrome are aware of this problem. This is why both of them allow their users to decide whether 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 in Chrome's Incognito mode or Firefox's private window modes. Another option, since we don't lack for browser choice, would be to reserve a secondary browser which is running perhaps only your password manager, for use during highly confidential private work.

In any event, we do currently have a problem that's going to require some eventual resolution.

person typing #22 / @pt22

> *Hey again Steve, In last week's SN 938 the tradeoff between security and convenience was mentioned with respect to websites and browser extensons 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 (1Password is an example, and I use KeePassium). 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.*

That’s a great solution to the dual-browser approach and 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 into the path. Very nice!

A husband who’s wife let him use her X account since he doesn’t use his often and X wouldn’t let him logon without lots of pain…

> *Hello Steve*
>
> *Thank you for the many great shows. I have been listening since episode 1 and was overjoyed to hear you are not stopping at 999.*
>
> *With regards to the web extension security research story, I help develop a web application used internally by the majority of banks in the US. A few years ago we implemented Content Security Policy (CSP) headers. CSP has a useful feature were 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 site 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 papers solution of adding a secure input element or alerting the user of nefarious actions 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 all 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, browser extensions have been disabled. It is annoying that many useful tools are blocked and yet I cannot argue with their decision.*

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 about 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. What a mess!

Anthony Bosio / @abosio

> *@SGgrc I think Topics might be DOA. People are interpreting it and spreading it basically as “shares your browser history with other sites”.*

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’ve said they are going to do next year – is going to succeed because there won’t be any alternative to using it. So 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 suspect that, while it might take some time for the less technical types to understand it, it’s where the entire industry is going to go.

Barbara / @Barbara130

> *Some downloadable software basically is a stub that phones home to download the rest while installing,. I don't think giving the stub to virus total would be helpful.*

Barbara makes a good point. As we know, not all software we download is the entire package. We’re now often seeing a small “installer” that connects back to home base to download the entire package which is often many modules.

SKYNET / @fairlane32

> *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 is 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? 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 clean one instead? I don’t see how that 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.* 🤔🤔

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 assigned to known security researchers and VirusTotal to cause different “clean” software to be delivered on demand. And we have seen other non-web server examples of malware being aware of the IPs of researchers it wishes to hide from.

E. Remington / @ericrdemington

> *Hi Steve, one email provider people overlook is iCloud, you can set up your own domain. While iCloud limits you to, if I remember right, 5 email addresses, by using a form like “something” + eMail address@yourdomain.com you can have an infinite number of email addresses. And RFC822 and all of its updates have supported <id>+<tag>@domain.*

I’m glad E.Remington mentioned iCloud since he’s correct that iCloud as an eMail provided is easily overlooked. Very nice. Thanks!

Magnify247 / @Magnify247

> *Steve, with Windows 12 being prepped for 2024, will INcontrol be updated - or will the current version allow for the VERsion/RELease application accordingly?*

You know that old expression about “fool me once.” As we know, the predecessor to “InControl” was “Never 10”. I would have named the next one “Never 11” 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 that 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 all it’s controlling is a few registry keys which Microsoft officially supports, I fully expect that InControl will **not** require any tweaking once Windows 12 becomes available. It should keep working as long as Microsoft honors those settings. And since their enterprise users depend upon them, I can’t see anything changing.

Christian. P. / @thewayeye_

> *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 source. You can remove any risk of testing the file directly, by getting virustotal to download the file, then re-check the hash before you execute it. Virus total reports the sha256, and the hash is also in the URL of the result.*
>
> *So perhaps a sensible process would be to get virustotal to download and check, if that looks largely ok then download directly and test again. The second test should take you to the same page. Virustotal doesn't automatically re-test files that have already been submitted, it just computes 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. Great podcast etc!*

I wanted to share this since the way the world is evolving, keeping VirusTotal in one’s back pocket makes a lot of sense. Christian is right about the way VirusTotal operates. Before it does anything else, it first calculates the SHA256 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 library of previously scanned files. And if so, it just returns the previous result.

And Christian’s suggestion of then uploading your own copy of the hopefully-identical file that VirusTotal first approved of to see whether they match makes sense. I wanted to also note that Windows now has a built-in “hashfile” command using it “certutil” which also allows a user to quickly and easily generate an SHA256 hash of any file they wish to check:

```
certutil -hashfile C:\Users\user1\Downloads\software.zip SHA256
```

InspClousseau / @InspClousseau

> *@SGgrc Steve what DNS services do you recommend for children (under 10) to avoid unsafe/unsuitable sites?*

I think the service that Leo uses and recommends is OpenDNS. They are definitely reputable, they've been acquired by Cisco, and they have a free family tier which they refer to as their "FamilyShield" service. Using it is as easy as configuring the family's router to use OpenDNS's servers at two specific IP addresses: 208.67.222.123 and 208.67.220.123. Once you've done that you can go to https://welcome.opendns.com which will confirm that you're now using their filtered DNS. (I suppose if you wanted a bit more proof you could also try going over to PornHub and see whether that works.) https://www.opendns.com/setupguide/#familyshield

LastMess

Any regular listener of this podcast can probably guess that today's title of "LastMess" will have something to do with "LastPass" – and they would not be wrong. There's growing 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 archives, is now being and has successfully been decrypted and 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". Here's some of what Brian reported. He wrote:

> *In November 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 crypto.*
>
> *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, run full nodes."*
>
> *Monahan has been documenting the crypto thefts via Twitter (now 'X') since March 2023, frequently expressing frustration in the search for a common cause among the victims. Then on Aug. 28, 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 included a screenshot 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. 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 have 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.*
>
> *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.*
>
> *Chainalysis confirmed that the employee had suffered a high-dollar cryptocurrency heist late last month, but otherwise declined to comment for this story.*
>
> *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 that this is a real problem that I've been urging my friends and family 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:

> *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.*

Yep. That's a standard dodge. 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"*

Yeah... perhaps some additional litigation now. They wrote:

> *"Since last year's attack on LastPass, we have remained in contact with law enforcement and continue to do so. We have shared various technical information, Indicators of Compromise (IOCs), and threat actor tactics, techniques, and procedures (TTPs) 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 the 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."*

Brian's reporting then covers for his readers everything that this podcast already covered for our listeners. 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's International Computer Science Institute (ICSI) and a lecturer at UC Davis said about brute force attacks that: "You just crunch and crunch and crunch with GPUs, with a priority list of vaults 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 4 million [password guesses] per second with 1000 iterations, but that would go down to 8 thousand per second with 500,000 iterations, which is why iteration count matters so much. So a combination of 'not THAT strong of a password' 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 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 is still trying to claw back his losses. For his reporting, Brian refers to this person as "Connor", which is not his real name. Brian writes:

> *Connor said he began using LastPass roughly a decade ago, and that he also stored the seed phrase for his primary cryptocurrency wallet inside of LastPass. Connor chose to protect his LastPass password vault with an eight character master password that included numbers and symbols (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 box 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, Aug. 27, 2023, 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, and 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.*
>
> *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. I had that ready to go two months ago, but have been procrastinating moving things to the new wallet."*

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' latest forced upgrades *"a stunning indictment of the negligence on the part of LastPass."*

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 8-character random passwords or 'correct horse battery staple' type passphrases — is outright negligence. 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 fiance about this last night. She's waiting for LastPass to tell her to* *change everything. Meanwhile, I’m telling her to do it now.”*

Based upon all of our 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, but decrypting it, while technically possible, is almost certainly a time-consuming process. 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 are fairly certain that a cash equivalent pot of gold awaits them if they are successful.

So, anyone whose iteration count was high – 100,100 or greater rather than 1, 500 or 5000 – and anyone who was using a high-entropy passphrase at the time, is almost certainly safe.

The powerful significance of this story is that it seems almost certain that the stolen LastPass encrypted data is actually being selectively decrypted and someone is turning it into cryptocurrency cash. And that’s a big deal.