LastPass Again

Description: This week we answer a few questions. What if an Australian company doesn't secure their own network? Has Ireland NOT levied fines against any major Internet property owned by Meta? What's in REvil's complete dump of Australia's Medibank data disclosure? We finally answer the question, is nothing sacred? (It turns out it's not rhetorical.) Also, whose root cert just got pulled from all of our browsers, and how did a handful of Android platform certs escape? What U.S. state has banned all use of TikTok? What country is prosecuting its own ex-IT staff after a breach? How has memory-safe language deployment actually fared in the wild? Are last August's Black Hat 2022 videos out yet? And which brand of IoT security camera do you probably NOT want to use or purchase? Which podcast had the most amazing guest last week? What happened when SpinRite was run on an SSD? And what does LastPass's announcement of another hacker intrusion mean for it and its users? Answers to those questions and more coming your way during this week's Security Now! podcast.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-900.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-900-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. You've got questions? Steve's got answers. A lot of questions. Australia's stiff fines for not disclosing a break-in. What country is prosecuting its own ex-IT staff for a memory breach? Which podcast had the most amazing guest last week? I wonder. And what happened when SpinRite was run on an SSD? Plus an analysis of LastPass's recent revelation of an attacker intrusion. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 900, recorded Tuesday, December 6th, 2022: LastPass Again.

It's time for Security Now!, the show where we cover your security, your privacy. We explain things. We talk about the world as it is with this fantastic fellow right here, Mr. Steve Gibson. Hello, Steve.

Steve Gibson: Hello, Leo. You noted, and I had previously seen, that we are now at Episode 900.

Leo: Wow.

Steve: Yeah. So for the beginning of December. And there was little doubt that I had to title today's podcast LastPass Again.

Leo: Yeah.

Steve: Yeah, boy.

Leo: Holy cow. Yeah, this was breaking news at the end of the episode last week. And now we've had - you had time to look at it.

Steve: Yup, yup. So...

Leo: And I still don't know if we know exactly what happened, but I'm looking forward to hearing you.

Steve: In fact, yes, we will get there. We're going to answer a few questions. And when I looked at the word "few," I thought, wait a minute, we're doing more than a few. What if an Australia company doesn't secure their own network? Also, has Ireland not levied fines against any major Internet property owned by Meta?

Leo: No.

Steve: What's in REvil's complete dump of Australia's Medibank data disclosure?

Leo: Oh.

Steve: We finally answer, Leo, the question...

Leo: Yes? Yes?

Steve: Is nothing sacred? It turns out it's not rhetorical.

Leo: Oh.

Steve: Also, whose root cert just got pulled from all of our browsers? And how did a handful of Android platform certs escape? What U.S. state, I kid you not, has banned all use of TikTok?

Leo: What?

Steve: Uh-huh. What country is prosecuting its own ex-IT staff after a breach? How has memory-safe language deployment actually fared in the wild? Are last August's Black Hat 2022 videos finally out yet? And which brand of IoT security camera do you probably not want to use or purchase?

Leo: Oh.

Steve: Which podcast had the most amazing guest last week?

Leo: Hmm.

Steve: Uh-huh. What happened when SpinRite was run on an SSD? And what does LastPass's announcement of another hacker intrusion mean for it and its users? Answers to those questions and more coming your way during this week's Security Now! podcast.

Leo: You've got questions. Steve has answers. Inquiring minds want to know. Well, that's why we tune in the show, right, every week.

Steve: And why is nothing sacred?

Leo: I can't wait to hear that one. Some of those I think I know the answer to, but we'll see. Go ahead, sorry.

Steve: Our Picture of the Week...

Leo: Oh, I haven't looked yet.

Steve: ...is a great one. So the caption that came with it, it was perfect. So it reads: "If you've ever messed up a dimension or a hole position on something you're building, don't be too hard on yourself. At least you're not the Cisco design engineer..."

Leo: Oh, dear.

Steve: Un-huh, "...who caused an entire product line recall by placing the mode button directly above an RJ45 port."

Leo: Oh, man.

Steve: "That button resets the switch to its factory default settings when it's held down."

Leo: What happens when you plug it in Ethernet with a lock?

Steve: Yes, well, yeah. So for those who can't see the picture, what we see is an RJ45 plug in the first port of the switch. Now, if it were just a minimal RJ45, that would be fine. That was what the engineer was thinking. But if you put one of those rubber boots on it, where you have that rubber flap that comes over and protects the...

Leo: The lock, yeah, yeah.

Steve: ...plastic lock; right? Then what happens is that thing is perfectly positioned over the factory reset button.

Leo: Oh, my god.

Steve: So if you were then to lift the cord up, that would rotate the plug so that the protective boot pushes the factory reset button down.

Leo: And holds it, holds it down.

Steve: And holds it down, returning the switch to its, like, where the login and password, or cisco/cisco...

Leo: I don't understand why every time I set up our Cisco router, it goes back to factory settings. I don't understand it.

Steve: Yeah, what could be wrong? Oh, goodness. Anyway, great Picture of the Week. Thank you, listener who sent that to me.

Okay. So continuing our recent Australia watch, recall from last week that a recent cybersecurity country ranking which was published by MIT gave Australia the number one slot for Cyber Defense, followed by, in decreasing order, the Netherlands, South Korea, U.S., and Canada. So Australia number one. We recently covered Australia's proactive declaration of cyberwar, which they'll be waging against the world's perpetrators of cybercrime, not waiting for them to commit another crime; but, you know, going after them. You know, it would seem that these high-profile attacks on Optus, Telstra, Medibank that we'll be talking about a little bit later, Woolworths, and EnergyAustralia really woke up the sleeping bear and galvanized Australia into action. They've decided to go active.

Last week saw another facet of this campaign, with the creation of new legislation to replace Australia's creaky 34-year-old Privacy Act of 1988. The new legislation ups the ante when Australia's own internal attack targets, like those companies just mentioned, turn out to be willfully negligent. The legislation bears the cumbersome name "Privacy Legislation Amendment (Enforcement and Other Measures) Bill 2022." It grants the Office of the Australian Information Commissioner, the OAIC, the power to levy hefty fines on companies and not only Australian companies, which is interesting, we'll get to that in a second which ignore security best practices to needlessly expose their customers' data through cybersecurity breaches.

Under this bill, which is expected to receive royal assent shortly to place it into force, companies that fail to safeguard their data face fines of up to the greater of 50 million AUS or 30% of the company's adjusted revenue, whichever is greater.

Leo: Wow.

Steve: Yeah. So if you have more than 30 - if 30% is greater than 50 million AUS, that's what you've got to pay. So, now, this is a huge change. The existing fines impose only 2.22 million AUS as a fine as a result of security breaches. So we're going from 2.22 million to 50 million, or 30% of revenue if 30% is more than 50 million.

Okay. So steep as this is, the updating of Australia's antiquated legislation has, not surprisingly, been greeted with positive feedback from Australian cybersecurity experts, who view it, and I think reasonably, as the incentive needed to get local companies to pay the attention that they must to the state of their IT systems. Right? I mean, it's like, oh, we're sorry, like when a breech happens. It's like, well, if you'd worked to fix this in the first place proactively, then you wouldn't have to be sorry. And Australian citizens wouldn't have had all of this data lost.

So given the historical reticence that we keep seeing to getting ahead of this problem - nobody does - I can't see any other way to bring about the changes we need. It does still feel a bit wrong for the suppliers of the too-often buggy systems which enterprises rely upon to continue to be held harmless. But that's a bridge we're not prepared to cross yet. So at this point we're going to say, well, were you patching? And that's the other problem, too, that is uncomfortable here, is how is this decision about negligence reached because it's a thorny problem.

I mentioned before that this law wasn't applicable only to Australian companies. According to the bill's text, its provision and fines will also apply to any non-Australian company who's doing business in Australia, even if they're headquartered outside of Australia. And that one promises to prove interesting. You know, Australia fining some other company who's doing business in Australia, who exposes the private information of Australian citizens, 50 million AUS. Wow.

Okay. While we're on the topic of fines, Ireland's data protection agency fined Meta 285 million euros due to Facebook's data breach a year and a half ago, in April of 2021. The Irish Data Protection Commission said that Meta failed to safeguard its Facebook platform from data scraping. Which, as we know, data scraping is just like bots, spiders, get in and scrape all of the pages of Facebook. Anyway, this Irish Data Protection Commission alleges that this allowed a threat actor to compile details on more than 530 million Facebook users. This data was later sold on an underground cybercrime forum, thus bad guys profited. Facebook told TechCrunch that following the incident they rolled out protections to detect scraping operations. To which I would ask, again, why not before this? I mean, it really does seem that we're having to, like, severely punish these tech companies in order to get their attention and get them to change their behavior, crazy as it seems after the fact.

Since Ireland had previously fined Meta's Instagram 405 million euros in September of also this year, and WhatsApp got a fine of 228 million euros in the previous September, this now rounds out the fines so that Ireland's data protection agency has now fined all or each of Meta's three main platforms. So that answers the question we posed at the beginning: Is there any major property of Meta that Ireland has not yet fined? The answer now, no. They've got them all.

And speaking of Medibank and Australia, okay, remember Medibank, as I mentioned, was one of the several ransomware embarrassments that Australia-based organizations recently suffered. Well, Medibank stood up to the REvil gang, as everyone thinks they should, if you can, refusing to pay or to buckle under to REvil's extortion threats. They threatened to release the entire contents of what had been illegally stolen from Medibank. So Medibank is Australia's largest private health insurer. And what's known is that the significant personal data for Medibank's 9.7 million current and past clients was stolen during REvil's original intrusion and data exfiltration.

So last Thursday Medibank released a lengthy statement. I won't bother with it all, but it began: "We're aware that stolen Medibank customer data has been released on the dark web overnight. We are in the process of analyzing the data, but the data released appears to be the data we believed the criminals stole. Unfortunately, we expected the criminals to continue to release files onto the dark web. While our investigation continues, there are currently no signs that financial or banking data has been taken. And the personal data stolen in itself is not sufficient to enable identity theft" - although when you hear what it is, you'll maybe question that - "and financial fraud. The raw data we have analyzed today so far is incomplete and hard to understand." Okay. We'll get back to that because I thought that was an interesting thing that they said.

A bit later in their released statement, they start quoting Medibank's CEO David Koczkar. And so they wrote his quote, saying: "Anyone who downloads this data from the dark web, which is more complicated than searching for information on a public Internet forum, and attempts to profit from it, is committing a crime. The Australian Federal Police have said law enforcement will take swift action against anyone attempting to benefit, exploit, or commit criminal offenses using stolen Medibank customer data. We continue to work closely with the Australian Federal Police, who are focused, as part of Operation Guardian, on preventing the criminal misuse of this data. Again," he says, "I unreservedly apologize to our customers. We remain committed to fully and transparently communicating with customers, and we will continue to contact customers whose data has been released on the dark web."

Okay. At the end of the statement, Medibank then - and the statement goes on and on and on, like at great length - they enumerate the sobering details of what they believe the REvil gang both obtained, and a little bit of what they did not obtain. So on the daunting bit that they did get, they said, the name, date of birth, address, phone number and email address for around 9.7 million current and former customers, and some of their authorized representatives. This figure represents around 5.1 million Medibank customers, 2.8 million ahm customers, and around 1.8 million international customers.

Also, Medicare numbers for ahm customers, though not the expiration dates; passport numbers; and the travel visa details for international student customers, though again not expiration dates. Health claims data for around 160,000 Medibank customers, around 300,000 ahm customers, and around 20,000 international customers. This includes service provider name and location, where customers received certain medical services, and codes associated with diagnosis and procedures administered.

Additionally, around 5,200 My Home Hospital (MHH) patients have had some personal and health claims data accessed, and around 2,900 next of kin of these patients have had some contact details accessed. Also health provider details, including names, provider numbers, and addresses. So, you know, wow. A huge breach.

On the flipside they said that REvil did not access primary identity documents such as drivers' licenses for Medibank and ahm resident customers. Medibank did not collect primary identity documents for resident customers except in exceptional circumstances. REvil also did not access health claims data for extras services such as dental, physio, optical and psychology. And they did not access credit card and banking details.

Okay. But, wow, they still got a lot; right? Lots of stuff on 9.7 million individuals. What most caught my attention was their statement that "The raw data that we have analyzed today so far is incomplete and hard to understand." What occurs to me is that a raw unorganized dump of data concerning 9.7 million current and past customers is far less actionable than an organized searchable online database containing the same information. In other words, it's almost entirely the structure of the data, I mean, that much data, that gives it meaning and makes it useful.

If REvil grabbed raw data files without formatting templates and indexes into the data, if the database is highly relational in nature and deeply depends upon the interrelationships of pieces of it and the indexes into it in order to pull it together into something coherent, then the release of a massive blob of raw and disorganized data, where there's nothing to make clear what pieces goes with which, might be much less damaging than it at first appears. Which is interesting. It sort of says that, you know, while the data is in place, it's actually useful to the organization that owns it, and that knows how to interpret it. But if you just grab a static blob, a static file, then there may not be, I mean, it's all there. But if you can't, like, find all the little bits and pieces which get pulled together by having this data understood by the database and databases that contain it, then it's probably not so useful to anyone. Anyway, just interesting to have this actually happen. Okay. Leo?

Leo: Yes?

Steve: I couldn't resist giving this story, this short bit of news, the heading "Is Nothing Sacred?"

Leo: Okay.

Steve: Because the official website of the Vatican was pushed offline last Wednesday.

Leo: Oh.

Steve: By a DDoS attack.

Leo: Is nothing sacred?

Steve: Exactly.

Leo: I'm sure Father Robert had a late night that night, I would bet.

Steve: Yeah. It was pro-Russian hacktivists. And I had to look this up. CNA, all right, the Catholic News Agency.

Leo: Oh, well, there you go. They're serious.

Steve: CNA, the Catholic News Agency, pointed out the attack came a day after Moscow criticized Pope Francis's latest condemnation of Russia's invasion of Ukraine. So, no, nothing is sacred. And if the Vatican says, "Putin, you are bad," well, get ready to be blasted off the Internet for a while.

Okay. Mozilla yanks a no longer trusted root. As we know - actually, not only Mozilla, but everybody else - because we've covered this through the years, web browsers are extremely reticent. Okay, it's not the web browsers, right, it's the people who manage

them. But web browser teams, I should have said, are extremely reticent to remove root certificates - actually, Leo, maybe one day web browsers themselves will be sentient and will then decide whether they should trust a root certificate.

Leo: I'm sorry, Dave. I can't go there.

Steve: I don't trust that website any longer, and you shouldn't either. Okay. So they are reticent to remove root certificates from their trusted root stores because doing so immediately renders invalid and not trusted any and all outstanding certificates which have been previously signed by that certificate's authority using their matching private key. Pulling the trust from the root effectively puts the certificate authority out of business overnight. And as we know, this has happened a few times since the start of the podcast, and it's always interesting.

In this case, the certificate authority in question is an apparently shady Panamanian firm called TrustCor. Nearly a month ago, long simmering questions about TrustCor were brought to a boil by a piece in the Washington Post whose headline was "Mysterious company with government ties plays key Internet role. TrustCor Systems vouches for the legitimacy of websites, but its physical address is a UPS Store in Toronto."

Leo: Oh, that's not good.

Steve: Whoa.

Leo: Oh, boy.

Steve: No. That'll get your attention. Here's just a sampling of the juicy bits from The Washington Post's reporting. And for what it's worth I've got the link in the show notes. I really recommend any of our listeners who enjoy gossip based in facts, this thing is just - this reads like you couldn't make this up.

Okay. The Washington Post wrote: "The company's Panamanian registration records show that it has the identical slate of officers, agents, and partners as a spyware maker identified this year as an affiliate of Arizona-based Packet Forensics, which public contracting records and company documents show has sold communications interception services to U.S. government agencies for more than a decade.

"One of those TrustCor partners has the same name as a holding company managed by Raymond Saulino, who was quoted in a 2010 Wired article as a spokesman for Packet Forensics. Saulino also surfaced in 2021 as a contact for another company, Global Resource Systems, that caused speculation in the tech world when it briefly activated and ran more than 100 million previously dormant IP addresses assigned decades earlier to the Pentagon. The Pentagon reclaimed the digital territory months later, and it remains unclear," wrote the Post, "what the brief transfer was about, but researchers said the activation of those IP addresses could have given the military access to a huge amount of Internet traffic without revealing that the government was receiving it."

And our listeners may recall that we talked about this weird event at the time, noting how odd it was that this previously dormant and DoD-reserved block of IPv4 address space was suddenly being routed and tied to some random private company no one had ever heard of. Anyway, the Post continues: "TrustCor's products include an email service

that claims to be end-to-end encrypted, though experts consulted by The Washington Post said they found evidence to undermine that claim. Researchers said that a test version of the email service also included spyware developed by a Panamanian company related to Packet Forensics. Google later banned all software containing that spyware code from its app store.

"A person familiar with Packet Forensics' work confirmed that it had used TrustCor's certificate process and its email service, MsgSafe, to intercept communications and help the U.S. government catch suspected terrorists. Speaking on the condition of anonymity to discuss confidential practices, the person said, "Yes, Packet Forensics does that." And come on. The name "Packet Forensics" should be an obvious enough tell.

Leo: It should tell you everything you need to know; right? Yes.

Steve: Yes, about the company's intentions. Remember, any device that's holding a certificate which is able to sign other end certificates is thereby able to intercept any and all TLS-secured traffic bound for any remote web server. It accepts the connection, the TLS connection; examines the domain being requested; creates and signs a TLS certificate on the fly; and returns it to the browser. In this case, so long as all web browsers contained the TrustCor CA root certificate, they would happily accept that on-the-fly signed certificate. So the connection to the intercept middlebox, as they're called, would be encrypted where the middlebox would decrypt the TLS data for completely in-the-clear analysis.

Leo: It's a man-in-the-middlebox attack.

Steve: That's exactly what it is. The middlebox would then initiate its own connection to the actual destination server so that its interception was invisible, while it continued to surveil all of the intercepted browser's traffic.

Leo: Oh, my god. Oh, dear.

Steve: So what is really, I mean, it's heartwarming to see how long the thread was in the Google group's back and forth, while they explained very patiently to the TrustCor representative, who kept trying to rebut all of their evidence that, you know, we're going to pull the plug on you.

Leo: Yeah. It's pretty obvious, yeah.

Steve: Because this is not okay, yeah.

Leo: How did he get in in the first place? That's my question.

Steve: Well, I mean, apparently if you have enough money - there was a quote somewhere later in this long Washington Post article. If you had enough money, basically you could buy yourself certificate authority privileges. And, you know, essentially they don't want to deny a company just out of hand who, like, passes all the requirements

and certifications and looks like they're going to be a reputable and reliable certificate authority. It's like, well, you know, why should we say no to them? So...

Leo: I would have hoped there was a better way to vet certificate authorities given the power that it gives them.

Steve: I know. And, you know, Mozilla currently recognizes 166 root certificates, but no longer the three from TrustCor.

Leo: Oh, good.

Steve: Our really long-time listeners may recall that episode...

Leo: The Hong Kong Post Office Episode. You don't have to - say no more. I know exactly where you're going.

Steve: That followed my chance discovery of what was then the explosion in certificate authorities in Firefox's root store. The time before when I had looked, there were like seven or eight certificate authorities. Now there appear to be hundreds. And as I said at the time this is inherently not good. All web browsers are trusting any certificate signed by the owners of any of these root certs. It makes it an inherently unstable system. But in fairness, you'd have to say that things have gone much better, certainly than I expected. The industry has been amazingly effective at policing itself. And the events, you know, of these trusted root store abuses have been very few and far between. You know, it's an obvious privilege to be granted certificate authority rights. It's a license, essentially, to print money, but only so long as the certificate authority's signature means something.

Okay. So the Washington Post's story is not behind any paywall, and it reads, as I said, like someone's imaginative fiction, being well researched and backed by facts. I've got a link, as I said, in the show notes for anyone who's curious to know more. And also I have a link to the Google group's discussion, which...

Leo: It sounds pretty good. I've got to read that.

Steve: Oh, yeah. All of the industry's participants, including TrustCor's, you know, TrustCor's representative was there, trying to rebut this. But ultimately they lost the argument because the evidence was just there.

Leo: Yeah.

Steve: You know, it just - it went on and on and on. And they said, look, sorry, but we can no longer trust you to sign anything.

Leo: This initial - the kickoff post has 34 references, footnoted references. That's a pretty good start.

Steve: Yeah. Yeah.

Leo: That's not just something off the cuff. I think this is a problem.

Steve: Again, you have to take your hat off to these guys. They do not yank this privilege casually.

Leo: Yeah.

Steve: I mean, you really, you know, they fret and worry and make sure that this, you know, that it wasn't a one-off sort of event.

Leo: In fact, there's almost a smoking gun connecting TrustCor to spyware.

Steve: Yes, yes.

Leo: I mean, it's pretty bad. I mean...

Steve: Yes.

Leo: A UPS post box in Toronto.

Steve: Oh, my god.

Leo: Holy cow. Holy cow. Oh, here's the TrustCor response, yeah, yeah. Oh, I'm going to read this. That's some good late-night reading. Thank you.

Steve: Yeah, it really is good.

Leo: Yeah, yeah.

Steve: Okay. So one more, and then we'll take our next break. While we're on the subject of crucial certificates and certificate management, last Wednesday the 30th, an internal Google report which was originally created on the 11th of November was made public. And two days later, on Friday, the security firm Rapid7 pulled the pieces together. Google's report is titled "Platform certificates used to sign malware."

And under Technical Details of Google's report they said: "A platform certificate is the application-signing certificate used to sign the 'android' application on the system image. The 'android' application runs with a highly privileged user ID - android.uid.system, basically it's like the root - and holds system permissions, including permissions to access user data. Any other application signed with the same certificate can declare that it wants

to run with the same user ID, giving it the same level of access to the Android operating system." In other words, it's a full penetration of Android security.

Digging a bit deeper, we find that the Android Security Team discovered several malware samples in the wild that were signed by platform certificates issued by major vendors including Samsung, LG, MediaTek, and Revoview. After discovering the incident, the Android Security Team worked with the affected companies to revoke and rotate the leaked platform certs.

I liked what Rapid7 had to say about this because what they said made a lot of sense about what didn't make a lot of sense about the whole escapade. Here's what they wrote. They said: "On November 30, 2022, a Google report initially filed on November 11th was made public. The report contained 10 different platform certificates and malware sample SHA-256 hashes, where the malware sample had been signed by a platform certificate, the application signing certificate used to sign the 'Android' application of the system image. Applications," they wrote, "signed with platform certificates can therefore run with the same level of privileges as the 'Android' application, yielding system privileges on the operating system without user input. Google has recommended that affected parties should" - I be they've recommended. "Affected parties should rotate their platform certificate. However, platform certificates are considered very sensitive," Rapid7 wrote, "and the source of these certificates is unknown at this time."

They said: "This use of platform certificates to sign malware indicates that a sophisticated adversary has gained privileged access to very sensitive code-signing certificates. Any application signed by these certificates could gain complete control over the victim device. Rapid7 does not have any information that would indicate a particular threat actor group as being responsible; but historically, these types of techniques have been preferred by state-sponsored actors." Meaning, right, like those like at the top of the food chain. "That said," they wrote, "a triage-level analysis of the malicious applications reported shows that the signed applications are adware, a malware type generally considered less sophisticated. This finding suggests that these platform certificates may have been widely available, as state-sponsored actors tend to be more subtle in their approach to highly privileged malware."

Okay. So some low-end malware adware was somehow signed by like the most closely guarded private keys belonging to some of Android's largest and most reputable vendors. Either those closely signed private keys escaped, or somehow those still-resident keys were used to sign the malware. Either way, the fact that malware was signed means that something went wrong. What's weird is that any agency that somehow obtains the ability to get any malware signed by major platform keys is not going to waste that awesome privilege on easily discovered adware. They would treasure that capability and hold it close, choosing to reserve its use for only highly targeted infiltration specifically so that it never was discovered because, as soon as it is, it's going to be rendered, you know, it's going to be neutered by having the keys rotated.

Now, thanks to the casual misuse of a collection of certificates that somehow escaped from something, whoever or whatever gained the ability to sign those certs has almost certainly lost those rights. None of the signatures of those certificates will be trusted going forward. Given what we know, none of this makes any sense. So we have a mystery. But, you know, it's been dealt with, thank goodness. Wow.

Leo: Interesting.

Steve: And after we tell our listeners why we're here, Leo, we're going to find out what state has banned the use of TikTok.

Leo: I think I know the answer. But I'll wait and hear.

Steve: Oh, my goodness.

Leo: I think it's only for government employees, not for you and me.

Steve: That is true. Well, yeah, you can't ban it statewide, of course.

Leo: You couldn't. You couldn't.

Steve: No.

Leo: Yeah.

Steve: Thank god.

Leo: Yeah. What would my son do without all his millions of fans on TikTok? I ask you.

Steve: So here's one for you. Last week South Dakota's Governor Kristi Noem...

Leo: Oh, good old Kristi. Ah, yes.

Steve: ...signed Executive Order 2022-10, which bans all use of the Chinese social media platform TikTok by state government agencies, employees, and contractors. The Executive Order's news release stated that the order is in response to the growing national security threat posed by TikTok due to its data-gathering operations on behalf of the Chinese Communist Party. You know, Leo, you've got to keep your eye on those commies.

The press release quoted Governor Noem saying - wow. "South Dakota will have no part in the intelligence-gathering operations of nations who hate us. The Chinese Communist Party," she says, "uses information that it gathers on TikTok" - apparently from watching Hank make things - "to manipulate the American people."

Leo: Yeah.

Steve: They're being manipulated by Hank's salt.

Leo: Yes.

Steve: And they gather data off the devices across the platform. She says: "Because of our serious duty to protect the private data of South Dakota citizens, we must take this action immediately. I hope that other states will follow South Dakota's lead; and Congress should take broader action, as well." The order took effect immediately and applies to all employees and agencies of the State of South Dakota - no more TikTok for you.

Leo: No.

Steve: Including persons and entities who contract with the state, Leo, commissions and authorities or agents thereof. And thinking about that, you know, I thought, I really do wish that I would still be alive in another hundred years to see what the Internet has become by then. You know, will it have succeeded in pulling the world together? Or will the world's fearful leaders have established borders and regional controls just as they have everywhere else?

Leo: Wouldn't that be a terrible thing for the Internet? I mean, that's just what we don't want; right?

Steve: Yes. And it seems to be happening. It's getting chopped up and fragmented and regulated. And now unfortunately everybody is like suing everyone because they're not happy with the outcome of using it. Speaking of which, Albania has blamed its IT staff. Remember the drama that we covered back in July, where Iran retaliated against Albania by melting down their government networks. Then Albania retaliated back, which I guess is redundant, by severing diplomatic ties with Iran and sweeping into the just-closed Iranian Embassy looking for anything that Iran might have not sufficiently destroyed before leaving.

Also recall that it turned out that Iran had been rummaging around in Albania's networks since April of 2012, so for more than 15 months without ever being detected.

Leo: They need a Thinkst Canary.

Steve: I thought that as I was pulling this together.

Leo: They really could use that, yeah.

Steve: I thought, you know, someone needs to give those Albanians a clue. Well, who is to blame for all this? It must be someone's fault; right? And we can't blame the vendors of the buggy systems. After all, they provided patches. For some of the problems. Usually. Eventually. Again, we've got to blame someone. So Albania has decided that it was all the fault of the IT staff, and so now they're in trouble.

Leo: Oh. Oh, boy.

Steve: Albanian prosecutors have charged and asked for the house arrest of five government employees. The prosecutors say the five accused failed to apply security

updates to government systems and also failed to detect the hackers that had been wandering around inside their network as far back as April 2021.

Okay. So maybe the IT guys were seriously negligent. But we know that's not necessarily the case. If I may segue for a moment, a perfect example of Albanian-scale negligence not being necessary is the news that the U.S. Department of Homeland Security's Cyber Safety Review Board recently said that it intends to review attacks carried out by the Lapsus$ extortion group and will publish a report detailing how Lapsus$ managed to bypass a broad range of security measures without the use of advanced malware and managed to breach a large number of high-profile targets including Cisco, Microsoft, Nvidia, Samsung, Uber, Rockstar Games, and others. These companies are not firing their IT department staff because they recognize that it's possible to do nothing wrong and still be breached.

Okay. In Albania's case, it could just as easily have been His Excellency the President of the Republic of Albania who clicked a link in a phishing email to invite those crafty Iranian cyberwarriors to come for a visit. And who knows what managerial opposition or budgetary constraints the intrepid five might have faced in their department? IT departments are notoriously understaffed, overworked, and unappreciated. And IT people are just like everyone else; right? There are good ones, and there are bad ones. Which are they? We don't know. What I wonder, though, is who they're going to get to fill those vacated jobs? With the risk of prosecution...

Leo: Oh, good point.

Steve: With the risk of prosecution for attempting to do a job that might be impossible, and knowing what happened to the last five guys, I would not be surprised to learn that those IT staff positions are difficult to fill. So, you know, I would be careful, you know, how you deal with problems like this in the future. Wow.

We do have some good news on the memory-safe languages front. Since the August release of Android 13, which was the first Android where a majority of the new code added to the project was written in memory-safe languages including Rust, Java, and Kotlin, Google noted that since shifting its focus to memory-safe languages, the number of memory safety vulnerabilities reported in the Android OS has dropped to less than half of comparable counts. So that's good news for memory-safe languages. You know, I've always been saying that we're never going to get our systems fixed if we keep messing with them. You know, this is of course the big problem with Windows is Microsoft refuses to stop. And they just keep doing stuff.

Well, when you do new stuff, you're going to have new problems. And so the only thing you could do would be to start using, I mean, high-quality, memory-safe languages for all the new stuff you do. And that's what Google's been doing with Android, and they're seeing a precipitous drop. I have a chart in the show notes that shows successive years of Android releases - 2018, '19, '20, '21, and 2022. And, I mean, it is really looking good. So something has to change in order for these problems to change. And empowering programmers with languages that help them makes all kinds of sense.

Leo: I'm kind of surprised Kotlin is so low because, boy, everybody's so excited about Kotlin.

Steve: Yeah.

Leo: It is still a tiny fraction of the overall development.

Steve: I think it's because it's just the very start.

Leo: It's new, yeah, yeah.

Steve: Yeah. And isn't Kotlin the one that runs on top of the Java VM? I think it...

Leo: Ultimately, at some point, Java's got to be - I think has to be in there, although I see Rust and C and C++.

Steve: Yeah.

Leo: But the Google, the last time I looked at Android development, the underneath, underlying stuff is Java. So I don't know, yeah, Kotlin would make sense to a virtual machine for Java VM, or a frontend for a Java VM.

Steve: And I think it is a different language on top of the Java VM.

Leo: Oh, it's a wonderful language. I mean, it's a very, yeah, it's for the JVM, that's right.

Steve: Right.

Leo: It's a great language, and I would guess probably more memory-safe, has null safety and stuff like that. So, yeah.

Steve: Right. It's pretty-looking.

Leo: Yeah.

Steve: For me, the prettiest-looking language I ever did any serious work in was Pascal. It was just - it was just - it was pretty. And you could come back later, and it made sense to you.

Leo: Yes. It was very concrete.

Steve: Yeah. And the least pretty was Forth. Forth is a write-only language.

Leo: Forth is fun.

Steve: You could stare at that and have no idea what the hell is going on.

Leo: Huyen Tue Dao is one of the hosts on All About Android, an Android developer, loves Kotlin. And I remember when they announced that they were going to support Kotlin first-class, and this is four years ago or five years ago at Google I/O, the developers cheered. So I have high hopes.

Steve: Good.

Leo: Don't want to see all that C. Look how big the C slice is.

Steve: I know.

Leo: Talk about not memory-safe.

Steve: And C++ is, like, together they're almost half.

Leo: Yeah. Yeah, the other half is Java. Big slice for Rust, though. That's also good, good news.

Steve: Yeah, yeah. Okay. We're going to answer another question. Have those Black Hat USA 2022 talk videos, which were recorded back in August, finally been published? Why am I asking, you might wonder?

Leo: Probably because...

Steve: Because the answer is yes.

Leo: Oh, they have.

Steve: They have been, yes.

Leo: Woohoo.

Steve: I have a link in the show notes for anyone who's interested, right below that graphic you were just showing. It brings you up to a playlist of all the Black Hat 2022 videos. And those are always interesting for hackers.

Leo: And they're all on YouTube, which is great.

Steve: Yup. We do have another Chrome zero-day biting the dust, which brings the total up to nine for Google. It was a type confusion bug in Chrome's JavaScript V8 engine. It was discovered internally by one of Google's tag researchers. But being a zero-day it was found because somebody was using it. So, wow. There's a lot of pressure to get into Chrome, it being the majority browser now. And this was another way that's now been foreclosed. At some point over the weekend I restarted Chrome, and it came up with an announcement of, yay, you've got a new version. It's like, oh, okay, good.

Okay. The Verge's coverage of Anker's Eufy, spelled E-U-F-Y, IoT cameras did not pull any punches. Their headline read "Anker's Eufy lied to us about the security of its security cameras." And then the subhead said "Despite claims of only using local storage with its security cameras, Eufy has been caught uploading identifiable footage to the cloud. And it's even possible to view the camera streams using VLC." Okay. Since I can't improve on The Verge's coverage and reporting, here's what they wrote, the beginning of it. It was long, but this will give you the idea.

They wrote: "Anker" - a company we all like - "has built a remarkable reputation for quality over the past decade," said The Verge, "building its phone charger business into an empire spanning all sorts of portable electronics, including the Eufy home security cameras we've recommended over the years," said The Verge. "Eufy's commitment to privacy is remarkable. It promises your data will be stored locally, that it 'never leaves the safety of your home,' that its footage only gets transmitted with end-to-end military-grade encryption" - okay, at this point you start to have to worry, right, when someone says "military-grade encryption."

Leo: I never want to see that phrase again.

Steve: No.

Leo: Every advertiser puts it in there. I just take it out and say AES-256 or something.

Steve: It is frightening, yeah. And they said that "It will only send the footage 'straight to your phone.' So," The Verge wrote, "you can imagine our surprise to learn you can stream video from a Eufy camera from the other side of the country with no encryption."

Okay, now, The Verge's coverage of this might seem somewhat harsh. But they then show a snapshot of the marketing for the Eufy camera which makes all of these claims quite clear, which then makes the reality of what Anker is doing somewhat stunning. So in the show notes I have this snapshot, which is right off of the Eufy marketing page. They said: "Our Technology Keeps Your Privacy Safe." Okay, now, I'm not sure that privacy can be kept safe. So the wording of the headline...

Leo: Yeah, you're not safe, but your privacy is.

Steve: Yeah. So, but they got two, they got the words in there that they wanted; you know.

Leo: That's the main thing, yeah.

Steve: You know, "your," "privacy," and "safe." So, okay. So maybe actually whoever wrote this didn't even understand what they were saying. I don't know. But that might be their way out of this corner. So they've got three big icons. Local Storage, and we've got kind of your house, and there's a server with a power symbol on it. And it says: "For Your Eyes Only. Home is where your data belongs. With secure local storage, your private data never leaves the safety of your home and is accessibly by you alone." Okay, now, consider that in the context of the fact that you can stream it from the other side of the world with VLC.

Okay, second icon: "End-to-End Encryption. Peeking Prohibited. All recorded footage is encrypted on-device and sent straight to your phone, and only you have the key to decrypt and watch the footage. Data during transmission is encrypted." None of that is true. "On-Device AI." Oh. "Everything In-House. Our super smart AI" - apparently much better than their super dumb crypto - "is built into every Eufy device. It analyzes your recorded footage without the need to risk your privacy by sending it to the cloud." Okay? Like all of this is untrue, stunningly.

Leo: Wow. Wow.

Steve: I mean, it's just, it's, like, incredible. None of it's true. So get ready for the lawsuits. And how. The Verge continued, okay, The Verge said: "Worse, it's not yet clear how widespread this might be because, instead of addressing it head-on, the company falsely claimed to The Verge that it wasn't even possible. On Thanksgiving Day, infosec consultant Paul Moore and a hacker who goes by Wasabi both alleged that Anker's Eufy cameras can stream encryption-free through the cloud, just by connecting to a unique address at Eufy's cloud servers with the free VLC Media Player." So, I mean, there shouldn't even be cloud servers; right? What? It never leaves your house. It goes straight to your phone. What do you need the cloud for? But apparently there's a cloud, and all your video is there. And you don't even need an app. You just use VLC and give it the URL.

"When we asked Anker point-blank to confirm or deny that, the company categorically denied it. 'I can confirm that it is not possible to start a stream and watch live footage using a third-party player such as VLC,' said Brett White, a senior PR manager at Anker."

Leo: Oh, well, he knows, yeah.

Steve: And of course that's exactly whose opinion you want regarding anything potentially damaging.

Leo: I couldn't do it. I tried, but I couldn't do it. What is this VLC?

Steve: Yeah, I clicked the link, and it just said hello. Wow. They wrote: "But The Verge can now confirm that's not true." That is, what Brett said. "This week we repeatedly watched live footage from two of our own Eufy cameras" - which of course they had been recommending in the past, so they probably had some - "using that very same VLC media player from across the United States, proving that Anker has a way to bypass encryption and access these supposedly secure cameras through the cloud."

They said: "There is some good news. There's no proof yet that this has been exploited in the wild." Oh, great. Now everyone's going to jump on that.

Leo: I don't know how you would know. I mean, you couldn't - there's no way to prove or disprove.

Steve: Right, right. You know? Suddenly the cameras are getting hot. I wonder why.

Leo: I do think you need to know the serial number of the camera.

Steve: That is true.

Leo: So that's some protection; right?

Steve: They said: "The way we initially obtained the address required logging in with a username and password before Eufy's website will cough up the encryption-free stream." Again, none of this, none, I mean, like what they're doing completely belies what they said they were doing. I mean, the fact that you have, like you log into the cloud, well, then the video must be there. It's not in your house. I mean, this, like, this should make your head explode. If you think about this, nothing they're claiming matches the services that they're offering.

So they said: "But it also gets worse." They said: "Eufy's best practices appear to be so shoddy that bad actors might be able to figure out the address of a camera's feed because that address largely consists of your camera's serial number encoded in Base64, something you can easily reverse with a simple online calculator. The address also includes a Unix timestamp you can easily create, plus a token that Eufy's servers don't actually seem to be validating - we changed our token to 'arbitrarypotato' and it still worked [thank you, The Verge] - and a four-digit random hex whose 65,536 combinations could easily be brute forced." And I'll note that other people have already done this, and they did it.

So a Mandiant vulnerability engineer, Jacob Thompson, tells The Verge: "This is definitely not how it should be designed." Yeah, no kidding. For one thing, serial numbers don't change, so a bad actor could give or sell or donate a camera to Goodwill and quietly keep watching the feeds. But also he points out that cameras don't tend to keep their serial numbers secret. Some stick them on the box and sell them at Best Buy. Yes, including Eufy.

On the plus side, Eufy's serial numbers are long at 16 characters and aren't just an increasing number. We've seen that done before. Not here. So: "You're not going to be able to just guess at IDs and begin hitting them," says Mandiant Red Team consultant Dillon Franke, calling it a possible "saving grace" of this disclosure. "It doesn't sound quite as bad as UserID 1000, then you try 1001, 1002, 1003," and so forth.

Anyway, I'm reminded of the fact that I don't have a single connected video camera anywhere within my environment, and that your wife Lisa, Leo, early on intuited the inherent dangers of having unknowable video capture technology, which is what all of this is, lurking around the house. In a TNO (Trust No One) world, the simple though impractical truth is: "Unless you designed it yourself, you don't know what it does." And I

should add that due to the crazy complexity of the things we design today, even if you did design it yourself, you may still not know that it does what you think it does.

Leo: Right. What you said, man.

Steve: Yeah.

Leo: Right on. The point is, maybe it isn't an easy to exploit, but they completely misrepresented what was going on.

Steve: Oh, my god, yes. Nothing that they said about it was the truth.

Leo: And they must have known better. I mean, it's not an accidental mistake disclosure.

Steve: The one out I had was when you have a situation like Anker was apparently in, where everybody loved their power supplies, which they probably themselves actually did create. There's a tendency to go buy other companies.

Leo: Right.

Steve: In order to expand yourself.

Leo: That may well have happened, yeah, yeah.

Steve: Yeah. And so it probably is like, well, we've got all this money from power supplies, who looks good? Oh, let's get Eufy.

Leo: Right. They make everything, too. They make headphones. They make all kinds of stuff. Anker has a sound core division.

Steve: Yeah.

Leo: They, exactly as you say, they found success, and they then expanded.

Steve: Yup. And so they, the Anker people, unfortunately are tying their good name to products that they can't actually vouch for, but which are making them money now, and they've got Brett out there on the front lines saying, what? What link? I don't have a link. Where did you get that link? That's illegal for you to have that link.

Leo: I tried, but I couldn't do it.

Steve: Okay. So this is moderately random, but not too far afield for this podcast. Everyone knows of my passion for coding. But I predate electronic computers. And before computers was electronics. Although coding has taken over, electronics will always be my first love. So in addition to coding, I occasionally do a bit of tinkering, hacking and designing with electronics.

At some point in the past, some googling must have taken me to a place called Seeed Studio. That's S-E-E-E-D, spelled with three E's, SeeedStudio.com. I purchased something from them, I don't now remember what, and as a consequence was promptly added to their periodic mailing list. In this case, I don't mind the spam because the mail contains photos of the stuff they're promoting, and my jaw spends most of its time hanging down with my mouth open over the insanely low cost of the technology that's currently available from China. It is truly astonishing.

For example, a recent mailing showed the "Seeed Studio XIAO ESP32C3." It's a tiny module about the size of a quarter with 14 electrical connections, seven on either side, and what appears to be two tiny buttons and an LED. It also has a tiny USB-C connector, presumably for programming this little thing. And all of the software for doing so is open source. Its description says: "Seeed Studio XIAO ESP32C3 adopts new RISC-V architecture..."

Leo: Ooh, ooh.

Steve: "...supporting both WiFi and BLE wireless connections." On this little thing.

Leo: On that thing?

Steve: WiFi and Bluetooth.

Leo: What?

Steve: "For Internet of Things applications, you will find it is flexible and suitable for all kinds of IoT scenarios." Okay. I was curious, so I looked into the chip this uses. The ESP32C3 is a 32-bit RISC-V microprocessor, which includes a whole host of I/O peripherals in addition to WiFi and Bluetooth 5. It has cryptographic hardware accelerators that support AES-128/256.

Leo: What?

Steve: SHA hashes, RSA, HMAC, digital signature, secure boot...

Leo: What?

Steve: And has a hardware random number generator. And how much is it if you purchase just one? $4.99.

Leo: Oh, my god. Oh, my god.

Steve: Five dollars for that. And that is just typical of what this Seeed Studio has for sale. Anyway...

Leo: Just don't try to bring it into South Dakota, that's all I'm saying.

Steve: No, no, no, no. That's outlawed. It might have TikTok embedded on it. Anyway...

Leo: That is so cool. And RISC, everybody's very interested in this RISC-V. This is the newest kind of open source digital architecture.

Steve: Well, yes. And license-free; right? The reason there's no ARM on this is you have to pay ARM for that. And you're not going sell something for $5 that has this and everything else it has if you have to pay some ARM licensing fee.

Leo: And it shows you what the ARM tax is, really, if you think about it.

Steve: Yeah. And so, I mean, RISC-V is - it's a beautiful architecture. It's been, like, moving along for years, and it's evolving. And it has an absolutely mature open source free tool chain for doing stuff. But none of that is why I've brought this up today.

Leo: Although, Steve, you could put, you know, it could be a SpinRite hardware device. Does it have room for software? You could put SpinRite on it. You wouldn't have to use DOS or anything. You just plug it in and boot to it.

Steve: Wonderful.

Leo: Five dollars, I'm telling you. I'm just telling you. All right. All yours.

Steve: Okay. That's not why I'm telling anyone about this. I'm telling anyone about this because a month or two ago, maybe three, something in one of those mailings brought me up short because it was similarly stunning, and I thought you guys, our listeners, all needed to at least know about it. It got away from me when I went back to try to find it. I didn't know where it went. But when their most recent mailing mentioned it again, I thought, okay, this time it's not getting away from me.

Okay. Get a load of this. It's called the LinkStar-H68K-1432 multimedia router. It has WiFi 6, 4GB of RAM, 32GB of eMMC flash storage onboard, with an SD card slot for more. It's powered by a quad-core 64-bit Cortex-A55 chip, an ARM G52 2EE GPU. There's a GPU because it can output HDMI 4K video at 60 frames per second.

Leo: What?

Steve: By the way, Leo, it's 2.5 by 3.5 inches. That's the size of that little thing that you're looking at. It has a USB 3 port, two USB 2 ports, a USB Type-C that can be attached to a SATA 3 drive. On the router side, aside from its dual-band 1200 Mbps WiFi 6, it also has four Ethernet ports, individual interfaces, two running at up to 2.5GB, another two at 1GB. It comes with Android 11 preinstalled...

Leo: What?

Steve: ...but also supports Ubuntu, Debian, Armbian, OpenWRT, and Buildroot, which is used to build embedded Linux systems.

Leo: I just thought of a new geek game we could play: Geek Price Is Right. Try and guess.

Steve: So what will this little pocket-sized fanless WiFi 6, four Ethernet interface router set you back? How about $119?

Leo: Unbelievable. Wow.

Steve: That's what got my attention. That little Netgate SG-1100 router that I love and use and have recommended, it's 189, and it only has three Ethernet interfaces and no WiFi. This thing has four separate interfaces and WiFi 6, and a ton more. The fact that you can drop OpenWRT onto it and have an operating, state-of-the-art router with four ports, all isolated individual subnets, and WiFi 6? For $119?

Leo: Could you put pfSense on it, do you think?

Steve: That I don't know. That's a question I have, and maybe one of our listeners will be interested to try. Again, it's hard to imagine this thing from the picture. It's 2.5 inches by 3.5 inches. And it's fanless. It's got a little heat sink on the bottom, multiple USB ports, 4K HDMI, and an SD card slot. It's just incredible. For 119 bucks.

Leo: In the palm of your hands.

Steve: I want to be clear, I don't own one. I don't have time to own one, and I'm not vouching for it in any way.

Leo: Do not buy it. We're going to get mad at you.

Steve: That's right. So I'm not vouching for it in any way. Unlike the ZimaBoard, which I was happy to vouch for since I had several, and I loved them, you're on your own with this thing if you should decide to take the plunge. For the right hardware tinkerer, this could be so much fun, and it's not very expensive. I have the link in the show notes, and it is Episode 900, this episode's GRC shortcut of the week, so grc.sc/900. That will take you to this thing's web page, where you can see for yourself. Anyway, I just - it was so

cool, so inexpensive, it could be the perfect home router. Four ports, and WiFi 6 for $119.

Leo: That's really cool. Boy, we live in amazing times, Steve. Can you imagine if you were a young guy, you know, a teenager, at the time building the Portable Dog Killer, if you'd had something like Seeed Studio available to you? You might have...

Steve: Unfortunately, I'd probably be bringing Elon's satellites down if I had this.

Leo: It's a good thing. It's a good thing you didn't have it. Wow. Wow.

Steve: I've figured out how to fire the retrorockets. Cool. Okay. Speaking of Elon, one last piece of lunacy. When asked during a scheduled Twitter Space chat this past Sunday why he bought Twitter, Elon explained his decision as follows, and I'm not making this up: "I can't exactly say why because it's one of those things where it's like my biological neural nets said, 'It is important to buy Twitter.' And just like with a digital neural net, you can't really exactly explain why the neural net is able to understand an image or text. The collective result of the neural net says this is an important decision, or this is the right action. And my biological neural net concluded that it was important to buy Twitter; and that if Twitter was not bought and steered in a good direction, it could be a danger for the future of civilization. And so that's why I bought it."

Leo: Wow. Wow. Clear as mud, Elon.

Steve: Yeah, Elon. So okay, you're passing the responsibility off to your brain, whose operation you don't understand.

Leo: I don't know what I'm doing or why I'm thinking it, but I'm going to do it anyway.

Steve: Yeah, because, you know, I'm a biological neural net just like those image recognizers, and we don't know how they work either.

Leo: It's a puzzle whether he knows what he's saying is moronic, and he's saying it to confuse and distract you, or if he actually believes it, which I do not know which is worse. It's amazing.

Steve: So I did not have time to make a comprehensive scan of my DMs this week. Frankly, the DMs channel is becoming quite popular, and there's a lot to go through. I'm only going to share my own tweet from last Wednesday for those of our listeners here who don't follow me on Twitter, and there are many. I tweeted on Wednesday: "To all Security Now! Listeners: I'm currently listening to Alex Stamos on Wednesday's 'This Week in Google.' Alex has not let anyone get a word in edgewise because he has so much amazing information to share. Without reservation, I RECOMMEND, all caps, listening to this. It's FANTASTIC! in all caps, exclamation point." That tweet received about three times as many likes as any of my weekly Security Now! notes posting tweets do, as well as 13 replies and 12 retweets. Alex was amazingly wonderful.

Leo: Thank you, yeah.

Steve: And I just wanted to make sure, I wanted, like, this is a listening assignment for all of our listeners. Last week's, so what was that, it'd be November 30th, This Week in Google. It was amazing.

Leo: If you just go to TWiT.tv/twig will take you to the TWiG page. It's Episode 692. So it'll be the first episode if you go right now. But even in a few weeks it'll still be TWiT.tv/twig, Episode 692.

Steve: And you don't even have to wade through a bunch of crap in the beginning. I mean, like...

Leo: No, there's no crap.

Steve: Alex said: "Is this microphone on?" And that's all it took.

Leo: It was very - rich density of information. He was great. He's a wonderful guy. And I really enjoyed him, and I'm hoping we can get him back because he had a lot to say.

Steve: So on my end, to say that things are going well with SpinRite's alpha release testing would be an understatement, considering how poorly things could have easily gone. I'm still somewhat in shock that we're very close to having a final release. I have things to fix, but nothing major so far. Mostly people who are now really engaged and involved are, like, keep asking for new features. And so it's like, oh, oh, oh, you know, that's really not what we should be doing now. So it's really looking good.

There was one posting to the newsgroup last week that I wanted to share because it makes a point that I need to drive home, not only for everyone's safety, but it's part of the reason why I am so fired up about SpinRite's potential long-term future. This person posted this. He said: "I have a ThinkPad Helix, and the SSD is a Samsung EVO 1TB mSATA." So a high-end Samsung EVO. That's all, that's the only brand I buy now. He said: "When the SpinRite pre-release starts, it estimates 31.7 minutes for processing. However, a Level 2 pass, with no errors detected, takes two hours and 56 minutes." Okay, so just shy of three hours. He says: "So that's more than five times longer" - it's actually almost six times longer; right? It estimated 31 minutes, it actually took three hours, so six times - "than the estimated time." He asks: "Is that normal?"

Okay. So I replied to him in the newsgroup: "We've found that whereas the fronts of spinning drives tend to be the fastest regions because they contain more sectors around their longer outer tracks, the fronts of many SSDs are conspicuously slow. We posit that this is due to the presence of much more on-the-fly error correction and data recovery. We first saw this using the ReadSpeed benchmark tool. One of my future plans is to locate these slow-to-read spots and selectively rewrite them to restore their speed by eliminating this unseen error correction and data recovery which results in a significant reduction in SSD performance."

I said to him: "If you do discover that the front of that drive is quite slow, you could identify the slow region and run Level 3, which does a rewrite, over just that region, and it might very well speed it back up." I said: "And that would likely also increase its reliability by solidifying those sectors which might be on the verge of transitioning from very slow to unrecoverable."

He replied: "I just did a Level 3 on the entire drive with Alpha 4, and now SpinRite estimates the 1TB drive will require 29.7 minutes, and a full Level 2 scan completed in 29:50, which is almost six times faster than it scanned a couple of days ago." He said: "Thank you so much for creating 6.1."

So what that means is, to summarize this, he first did a simple read pass on that SSD in his ThinkPad. It took him three hours just to read all the sectors. He ran Level 3, a SpinRite Level 3 across the entire drive. Now that same process of simply doing a read scan takes 30 minutes, just shy of 30 minutes. So running a Level 3 SpinRite on an SSD that had nothing technically wrong with it increased its speed by a factor of six, on average, across the entire SSD. That's what we're seeing.

So this is the reason I am very excited. What's happening is we are - SSDs are having trouble reading their contents, but still able to; yet it's revealed by a significant slowdown which is going unnoticed. But how many times have we heard that an SSD-based machine, which being solid-state you would think, right, it's solid-state, that an SSD-based machine doesn't seem to be as fast as it was when it was new. That wouldn't seem logical, but this might be what's going on.

So as I said, one of the things I have in store once SpinRite 6.1 is launched, and we start working immediately on 7, is to profile the performance of mass storage media to locate and selectively repair sluggish spots. But for what it's worth, what anyone can do, or will be able to do as soon as 6.1 is out, is to give such drives a single Level 3 pass, as this person who just posted did, which very well could significantly improve both the system's performance and its reliability by rewriting the drive's sectors to recharge those leaky storage capacity cells. So anyway, a cool outcome from the 6.1 work.

Leo: And now the continuing LastPass saga. And, you know, we still use LastPass for our enterprise access...

Steve: I do.

Leo: ...and [indiscernible] management. So let me know if we should stop. I mean, I use Bitwarden at home, but TWiT uses LastPass.

Steve: Like many other LastPass users, last week I received another note from LastPass. The note is short, so I'll read it into the record. We got a note from Team LastPass, is how it was signed. "Dear Valued Customer: In keeping with our commitment to transparency, we wanted to inform you of a security incident that our team is currently investigating. We recently detected unusual activity within a third-party cloud storage service, which is currently shared by both LastPass and its affiliate, GoTo. We immediately launched an investigation, engaged Mandiant, a leading security firm, and alerted law enforcement.

"We've determined that an unauthorized party, using information obtained in the August 2022 incident, was able to gain access to certain elements of our customers' information. Our customers' passwords remain safely encrypted due to LastPass's Zero Knowledge

architecture. We're working to diligently understand the scope of the incident and identify what specific information has been accessed. As part of our efforts, we continue to deploy enhanced security measures and monitoring capabilities across our infrastructure to help detect and prevent further threat actor activity.

In the meantime, we can confirm that LastPass products and services remain fully functional. As always, we recommend that you follow our best practices around the setup and configuration of LastPass, which can be found here." And then they provided a link. "As is our practice, we will continue to provide updates as we learn more. Please visit the LastPass blog for the latest information related to the incident. We thank you for your patience while we work through our investigation."

Okay. So there's no follow-up yet. Last time we waited precisely three weeks from the first announcement. The first announcement came on August 25th, and three weeks later we got the update on September 15th.

For me, the two most relevant pieces of information in this first of presumably two disclosures are where the person, the team, wrote: "We've determined that an unauthorized party, using information obtained in the August 2022 incident, was able to gain access to certain elements of our customers' information." So that. And "Our customers' passwords remain safely encrypted due to LastPass's Zero Knowledge architecture."

Okay. So as usual from all such partial disclosures, we're left wanting more. Of course. We don't know, at the moment, which "certain elements" of their customer information was inadvertently made available. But they apparently know. They're not saying yet. We know from the follow-up note from the first intrusion that the bad guys were rummaging around in some developer network that was not connected to their production systems. But now it appears that those who were doing the rummaging managed to get sufficient information...

Leo: They found something.

Steve: Right.

Leo: They found some keys in there.

Steve: Right, whether in the form of source code that disclosed, like, some other way to get into something else, or maybe in the form of credentials that were bound into whatever it was they got, which were used to perpetrate this latest breach and information exfiltration. If LastPass lost control of their customers' billing data names, credit cards, street addresses and so on that would not be good. But at this point we're just speculating. Presumably in another two weeks, or three, whenever, we'll be told more. You know, let's hope.

Last week, after this happened, I popped on for the first 15 minutes of Tech News Weekly with Jason and Mikah to talk about this latest breach. I made the same point that I always make, which is that none of the passwords and other secret data that's being stored by any of the many competing password managers, you know, LastPass and all of them, should ever be vulnerable to any breach of the data that's being stored on our behalf in the cloud. You know? That's thanks to what we once called - we coined the abbreviation or an acronym PIE, which stood for Pre-Internet Encryption, which the industry now calls end-to-end encryption, though that term is becoming less useful as

non-end-to-end systems abuse it. I mean, you know, even Apple says that iMessage is end-to-end encrypted, except that, well, it's not in the cloud because they have the keys to that.

So again, unfortunately the definition of "end-to-end" is being abused. But that's going to happen when any popular phrase gets into the marketing department. But the point is, if it's done correctly, we're really just using a password manager's cloud service to keep our various devices synchronized. It's simply a synchronization mechanism. And in fact my annoyance with the passkeys system which is becoming, hopefully will become popular, is that there isn't an all-in-one cloud sync for them. This is an opportunity for Microsoft to keep their devices separate from Apple, to keep their devices separate from Google, to keep their devices separate, unfortunately. But anyway, maybe that'll change when password managers start supporting passkeys. Maybe we'll get sync back. That's a possibility.

The threats we face to our stored secrets, and this is the point, are only on our end. In order to do its work, at some point any password manager must have at least the user's username and password decrypted for the site that they're visiting and want to fill in the form for. I don't know whether the entire password archive is decrypted as a whole, or whether sites can be decrypted individually, which would seem safer to me. But either way, at some moment in time, the data must exist in the clear in the user's browser. Way back at the start of this podcast we noted the inherent impossibility of protecting encrypted DVD video content because the player itself needed to be able to decrypt the DVD in order to play it for its owner, and the DVD's publisher ultimately had no control over the DVD player and what it did.

So all they could do was the best they could do. If the password manager's browser add-on were to be adulterated in some way to break its security design, or if something was able to somehow intercept its operation in the client, that would prove devastating. But it's difficult to see how any breach of LastPass or any other password manager's cloud synching facility could ever endanger a user's always-encrypted secrets.

Now, of course, none of this prevents reputational harm to LastPass. You've got to know they're not happy about this. But they're also, like, the biggest target in town. So that's what comes with being the big target, like Chrome is the big browser target. And most users will have no idea what it means for all of their data to always be encrypted before it leaves their browser. They don't see anything leaving their web browser. They have no concept of the cloud or of client-side encryption or what any of that means. They just know that they're using this or that password manager, and this or that password manager suffered a breach, and that the press is now able to say that this is, in the case of LastPass, the second breach in a little over four months.

If we assume that the decision to change password managers is unwarranted and I'm not suggesting whether it is or it not, it's a personal decision then one huge advantage any password manager has is inertia. It's much easier to change search engines than it is to change password managers. You know, certainly it can be done. And the password managers have provided means to lower the bar to doing so, offering various importers of other password managers' content, you know, and archives. But it seems most likely that, until users learn that someone's passwords were actually stolen, inertia will reign. The statement "Our customers' passwords remain safely encrypted" I think matters a lot, especially when changing password managers is a pain to do.

The listeners of this podcast, as knowledgeable and sophisticated as any, anywhere, you know, they're able to make an informed decision. I don't know what it means for them. I'm continuing to remain with LastPass because I have no interest in punishing them for making a mistake, and there is no indication that the security I actually require from them has ever been endangered. You know? That said, I'll be interested to learn more in

the next couple weeks when they're able to tell us more. I think their fiduciary obligation is to immediately engage law enforcement and an outside firm and tell us that this has happened. So it's reasonable for us to give them a few weeks for them to tell us more. And I'm sure we'll get something in a week or two.

Leo: Good. Nothing to fear yet, anyway.

Steve: Yeah. I mean, I just - I do think the fact that something that happened, some information that escaped from the first breach got used in some way, it's like, well, okay, that makes sense.

Leo: What it wasn't is a leak of - they don't know what your master password is. They can't unencrypt your vault.

Steve: Correct.

Leo: So there was nothing that could be leaked in that regard.

Steve: Right.

Leo: Another thing maybe to worry about, and this was something that Tavis Ormandy brought up, is the way that that JavaScript Chrome extension or Firefox extension works, the code in there, because at some point it does have to see it in the clear.

Steve: Yes, that is the concern is what happens on the client side.

Leo: Right.

Steve: But not on the cloud side.

Leo: Right, right.

Steve: Yeah, and in fact, you know, one of my favorite expressions back when I was spending time doing SQRL stuff was that I used to say "SQRL gives websites no secrets to keep."

Leo: Yes. That's the key.

Steve: And that's the key.

Leo: That's Trust No One.

Steve: Exactly. It's Trust No One. Not even - you don't even have to trust your password manager.

Leo: Right, right. And somebody's saying, well, didn't Steve review the code? Well, you reviewed the code, but like 10 years ago. I mean, I don't think any of that code survives.

Steve: And it wasn't their code in their cloud. What it was was it was the algorithm that Joe was using, which is this concept of client-side blob encryption.

Leo: Right, right.

Steve: That's what it was.

Leo: Right.

Steve: And in fact he even provided - he provided me a web page where you could go and, like, see a simplified JavaScript and understand it and see what it was doing.

Leo: Yeah. My guess is that all the major password companies do basically the same thing.

Steve: They'd be insane not to.

Leo: Yeah, E2EE encrypted block.

Steve: You have to do this. The moment a password company actually compromises their users' passwords, they might as well just declare Chapter 11. It's over.

Leo: Right, right.

Steve: Roll up the carpet. Nobody will ever use them again.

Leo: The other good side of this, LastPass has been spun off Log Me, which it's a complicated long corporate chain. But the equity capital company that bought Log Me is spinning LastPass off as a standalone company, which if it isn't already, it will be any day now. And in a way that's good because of course equity capital...

Steve: It helps them keep focus.

Leo: Exactly. And equity capital is always looking to make their investment back. They're usually highly leveraged. So just getting that off in its own corner where nobody's going to tell them what to do means that, I mean, they're good people. They know what to do, and they will keep doing the right thing as long as they're left to their own devices. And I think that's sort of what this is.

Steve: And, you know, this hurt them. There's no doubt. This hurt them. You know? No one wants to say whoops, we were breached. But, you know, nothing to see, move along. There's going to be a bunch of people that go, okay, that's two in one year, you know, I'm not waiting for three strikes.

Leo: Yeah, no, that's true.

Steve: I'm going with two strikes.

Leo: But you've reassured us, I hope. You know, I think everybody says...

Steve: I'm not changing. I'm...

Leo: And if you're in Australia, just, you know, don't you [laughter]. No, they had - they did responsibly disclose it; didn't they. They were the letter of the law in Australia, and GDPR has a responsible disclosure clause. I don't know if we do in the U.S. Feels like we ought to. But since...

Steve: CISA is definitely, certainly among government agencies, you know, thou shalt disclose.

Leo: Yeah.

Steve: I think it was within, was it 48 hours or 24 hours? I mean, it was...

Leo: It was a short time, yeah.

Steve: Yeah, there's a short window, at least. But I don't know if...

Leo: But they don't have the $50 million AUS penalty to back it up. That's the problem.

Steve: Ooh, boy.

Leo: That does get the job done, I must say. I must say. Steve Gibson, you get the job done each and every Tuesday with Security Now!. Thank you so much for Episode 900. What a milestone, huh?

Steve: Wow.

Leo: We would never have thought, way back in what was it, 2007, 2006...

Steve: We were still using cranks to start our cars, Leo.

Leo: It's a lifetime ago. But we're doing it, and we're going to do it again next Tuesday, and we're going to keep doing it for at least two more years. Go to GRC.com. That's Steve's website, the Gibson Research Corporation. Couple of things you can do there. Of course, first thing you should do is get SpinRite, the world's best mass storage maintenance and recovery utility.

Steve: Looking better every day.

Leo: You're getting 6.0, but we're minutes away from the release of 6.1. And anybody who buys now will get 6.1 for free as an automatic upgrade. You'll also get to participate in the development, which is kind of a fun thing to do. There's wonderful forums going on all the time. Steve's got his own newsgroup.

Steve: Yeah, don't ask for any new features. No new features.

Leo: No, no, too late. It's locked.

Steve: Please, yeah.

Leo: Solid. Locked in. While you're there, you also should check out all the other freebies that Steve offers. ShieldsUP is world famous. Things to test your network, router and so forth. And you can get a copy of the show while you're there, too, save yourself a step. He has 16Kb audio for the bandwidth-impaired; 64Kb audio, that's the standard version. And transcripts, written by Elaine Farris, so they're very well done. So you can read along as you listen, or you can search to find a part of it. All that's at GRC.com. Feedback forms are open at GRC.com/feedback. Steve's also on Twitter, although it sounds like it's a little busy now. Your DMs might be a little full.

Leo: Don't DM him there. But you can follow him.

Steve: I have some catching up to do.