The Inside Tracks

Description: I want to start off this week by following up on last week's podcast about the hardware backdoor discovered in Apple's silicon, to support the conclusion I've reached since then, that this was deliberate on Apple's part, that they always knew about this, and why. Then we're going to wonder whether everyone is as cyber-vulnerable as Ukraine appears to be. And if so, why and just how serious could cyberattacks become? What's the latest on the mess over at 23andMe? How's cryptocurrency been faring; and are things getting better, staying the same, or getting worse? What Google Mandiant account got hacked? Just how seriously, and legally, do we take the term "war" in "cyberwar," and what are the implications of that?

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-956.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-956-lq.mp3

LastPass recently announced some policy changes; even if they are about two years late, what lessons should the rest of the 'Net take away? During 2023, how did Windows 11 fare against Windows 10? What happens when users discover that Chrome's Incognito mode is still tracking them? And then, after exploring some questions from our terrific listeners, I want to share the result of some interesting research I conducted last week during the final days of the work on SpinRite 6.1 for this week's podcast, titled "The Inside Tracks."

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. An update on that CVE that provided a backdoor into Apple hardware. He's got some interesting afterthoughts. He also talks a lot about cyberwarfare and the vulnerabilities in Ukraine, plus 23andMe's latest disclosure. Is it really enough? All that and more coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 956, recorded Tuesday, January 9th, 2024: The Inside Tracks.

It's time, I know you've been waiting all week for it, Security Now! on the air. Steve Gibson, our man about town. He's the expert on privacy, security, how things work, and really I think one of the most trusted people in this business. He's right here. Hi, Steve.

Steve Gibson: Hey, Leo. Great to be with you again, as we're plowing through year 19 on the podcast.

Leo: Yikes.

Steve: Yeah. Okay. So this was a little bit sparse week in terms of security news. And I have to say it's difficult to follow last week's blockbuster with something similar, although I do think I have something I'm lining up for next week. So, but there's still lots to talk about. I want to start off this week by following up on last week's podcast, you know, which was obviously about the hardware backdoor that had been discovered in Apple's silicon, to support the conclusion I've reached since then, as I've continued to think about it, that this was deliberate on Apple's part, that they always knew about this, and why.

Then we're going to wonder whether everyone is as cyber-vulnerable as Ukraine appears to be; and, if so, why, and just how serious could cyberattacks become? What's the latest on the mess over at 23andMe? How's cryptocurrency been faring, and are things getting better, staying the same, or getting worse? What Google Mandiant, you know, their big security firm, account got hacked? Just how seriously, and legally, do we take the term "war" in "cyberwar," and what are the implications of that? LastPass recently announced some policy changes. Even if they're two years late, what lessons should the rest of the 'Net still take away from this?

During 2023, how did Windows 11 fare against Windows 10? What happens when users discover that Chrome's Incognito mode is still tracking them? And then, after exploring some questions from our terrific listeners, I want to share the result of some interesting research I just conducted last week during the final days of the work on SpinRite 6.1. Thus today's podcast is titled "The Inside Tracks."

Leo: Oh, of the inside tracks of your hard drive.

Steve: Uh-huh.

Leo: Yeah. Oh, I like that. Someday hard drives will no longer have inside tracks, but...

Steve: Some interesting news.

Leo: Yeah. Oh, good. Steve, I'm ready for the Picture of the Week. I've got it all cued up here. Haven't looked at it yet, though. I like to be surprised. Oh, I get it. That's great. Here, let me show everybody.

Steve: It's perfect for the podcast.

Leo: Yeah.

Steve: Perfect for the podcast. So this is the, today actually, on January 9th is the 17th anniversary or birthday of the introduction of the Apple iPhone.

Leo: Yes.

Steve: And anyway, so this is just an apropos birthday cake. And really, Leo, when I thought about it, until I realized, okay, most people wouldn't get it.

Leo: Wait a minute. There's only eight candles, Steve.

Steve: That's true.

Leo: And only two of them are it.

Steve: That's true.

Leo: Well, is that binary?

Steve: Uh-huh. Eight candles is a byte; right? An 8-bit byte. And, yes, binary. So the candle representing 16 is lit.

Leo: That's hysterical.

Steve: As is the candle representing one. Now, what occurred to me though is, first of all, you could probably get away with only seven, right, because that would bring you up to 127 years old, which that's, you know, humans don't live that long.

Leo: Plus you want a significant bit so you can have negative numbers, too.

Steve: Uh, well, that would be interesting, yes.

Leo: Yes.

Steve: Now, the danger is, though, that you can only approach this cake from the front.

Leo: From the right side, yeah, yeah.

Steve: Yes. Because if it's backwards, then 17 turns into 136.

Leo: Well, save it because in 129 years, no, 119 years, you'd be able to use this.

Steve: Leo, I'm afraid this podcast will not last that long.

Leo: Okay.

Steve: We're going past 999, but not to 9999999.

Leo: We can flip it. I love it. All right. Very nice.

Steve: Okay. So I want to begin, as I said, with a bit of a follow-up to last week's news.

Leo: Yeah, we've been talking about it. We talked about it on MacBreak Weekly today. We talked about it on TWiT on Sunday. I think it's really a big deal story, which I'm not seeing anywhere but here.

Steve: That's exactly right. And in fact one of the - I have some, we have some Q&A later. And one of our listeners said, how is this not getting more attention?

Leo: Yeah.

Steve: Anyway, we will talk about that when we get there. But the way we left things last week in the wake of this revelation was with a large array of possibilities. Since then, I've settled upon exactly one, which I believe is the best fit with every fact we have. You know, again, no speculation here. Although, again, we're never going to have a lot of answers to these questions. Many people sent notes following up on last week's podcast. Many doubted the NSA conspiracy theory because of those other, not easy to effect steps, involving other clearly inadvertent mistakes in Apple's code, that were needed by this particular malware.

I don't know why it didn't occur to me last week, but it has now. As we know and have covered here in great detail in the past, Apple has truly locked their iPhone down every way from Sunday. I believe from all the evidence and focus that Apple has put into it, that Apple's iPhones are truly secure. But would Apple actually produce a smartphone handset that they, and I mean they, absolutely positively truly could not get into, even if it meant the end of the world?

Leo: Oh, that's a very good point.

Steve: Right. If Apple believed that they could design and field a truly and totally secure last resort backdoor means of accessing their devices in the event that the world depended upon it, I believe that they would have designed such a backdoor. And I believe that they did, deliberately and purposefully.

Leo: For their own use.

Steve: Yes. And I do not think less of them for it. In fact, I think that the case could be made that it would be irresponsible for Apple not to have provided such a backdoor.

Leo: Yeah. What if Dr. Evil had an iPhone with the launch codes on it; right?

Steve: Well, that's where I'm going here. We'll likely never know whether any external agency may have made them do it. And yes, doing so could hardly be more controversial. But I can imagine a conversation among just a very few at the very top of Apple, Tim Cook and his head of engineering and of security. They had to have had a conversation about whether it should be possible, under some set of truly dire circumstances, for them to get into somebody else's locked phone. Obviously, the security of such a system would be more critical than anything. But their head of engineering security would have explained, as I did last week, that as long as the secret access details never escaped, because it's impossible to probe anything that must be accompanied by a signature hash, there would truly be no way for this backdoor to be discovered. As I said last week, from everything we've seen, it was designed to be released in the field where it would be active, yet totally safe.

So if Tim Cook were told that Apple could design and build-in an utterly secure, emergency, prevent the end of world escape hatch into their otherwise utterly and brutally secure devices, and this escape hatch could never possibly be opened by anyone else ever, I imagine Tim would have said, under those conditions, yes. I think that most CEOs who are in the position to understand that with great power comes great responsibility, when assured that it could not possibly be maliciously discovered and abused to damage their users, would say, yes, build it in.

I trust Apple as much as it's possible to trust any commercial entity operating within the United States. I believe that they absolutely will protect the privacy of their users to the true and absolute limit of their ability. If the FBI were to obtain a locked iPhone known to contain, exactly as you said, Leo, the location and relevant disarming details of a nuclear weapon set on a timer and hidden somewhere in Manhattan, I would be thanking Apple for having the foresight to create a super-secure means for them, and them alone, to gain entry to their device.

And I'd argue that in doing so they did have the best interests of their customers in mind. In this scenario, a great many iPhone users' lives would be saved. There are all manner of conspiracy theories possible here. Yeah, obviously. And this one of mine is only one of many. But of all the possible theories, I believe this one fits the facts best and makes the most sense.

Of course, the first thing everyone will say is, "Yeah, but Gibson, they did lose control of it, and it was being used by malware to hurt some of Apple's users." And that's true. In fact, that's the only reason the world learned of it. If this scenario is correct, Apple never divulged this to any entity, and never would. This would never have been meant for the FBI, CIA, or NSA to have for their own use. If an impossibly high bar were reached, Tim Cook would say "Have an agent bring us the phone, and we'll see what we can do."

But somewhere within Apple were people who did know. Perhaps someone inside was set up and compromised by a foreign group. Perhaps Apple had a longstanding mole. Perhaps it was gambling debt, or the threat of some extremely embarrassing personal information being disclosed. One thing we've learned and seen over and over on this podcast is that when all of the security technology is truly bulletproof, the human factor becomes the lowest hanging fruit. Just ask LastPass how they feel about the human factor, which bit them badly.

Okay, so where does this leave us today? We know that all Apple iPhones containing the A12 through A16 chips contain this backdoor and always will. We don't know that it's the only backdoor those chipsets contain, as we touched on last week. But Apple doesn't need another backdoor since they still have access to this one. They locked a door in front of this one, but they can always unlock it again. After being contacted by Kaspersky, Apple's iOS updates blocked the memory-mapped I/O access that was discovered being taken advantage of by malware. But Apple is able to run any software

they choose on their own phones. Which means that Apple still has access to this backdoor, should they ever need to use it.

And this means that they've lost plausible deniability. They have the ability to open any iPhone that they absolutely must. So this poses a new problem for Apple when law enforcement now comes knocking with a court order, as it's almost certain to, with way-below-that-bar requests for random iPhone unlocking to "assist" in this or that case. So this is a new mess for Apple. I'm sure they're facing that.

Apple's most recent silicon is the A17, yet Kaspersky told us that this facility had only been seen in the A12 through A16. If the malware did not contain that initial unique per-chip-generation unlock code for the A17 silicon, and we know that it didn't, then this backdoor might still be present in today's iPhone 15 and other A17-based devices. That's the most reasonable assumption, since it was there for the first, you know, for the previous five generations. Apple obviously likes to have it. But what about the next iteration of silicon for the A18?

Another thing we don't know is what policy change this disclosure may have for the future. We don't know how committed Apple was to having this capability, but I think I've made a strong argument for the idea that it has to have it. Have they been scared off? Well, maybe. We'll see what happens now, you know, as I said, with law enforcement asking them to unlock everybody's iPhone. Will they move it to a different location within the ARM's 64-bit address space, yet keep it around?

As we were after last week, we're left with a handful of unanswered and unanswerable questions. But my intention here was to hone this down to explore what appears to be the single most likely scenario. Apple designed, and is quite proud of, the GPU section of those chips which contains the backdoor hardware. There's no chance they were not aware of the little hash-algorithm-protected DMA engine built into one corner of the chip. People within Apple knew.

Listeners of this podcast know that I always separate policy decisions from mistakes, which unfortunately happen. So I sincerely hope that Apple's policy was to guard this as perhaps their most closely held secret, and specifically that it was never their policy to disclose it to any outside agency, of any kind, for any reason. Somewhere, however, a mistake happened. And I'd wager that, by now, Apple knows where and how that mistake occurred.

Leo: What about the assertion that it may be as an ECC, that that SBox is an ECC table?

Steve: We've seen the SBox. That's not error correction.

Leo: It's not.

Steve: In fact, when I first saw that posting, I was excited that there might be somebody who actually had some information. But they know nothing more than anybody else. The idea that that hash is ECC is complete nonsense. And so for me, I just thought, okay, well, these guys don't know anything more than the rest of us do.

Leo: It's easy to speculate, yeah.

Steve: Yeah. Well, we know what that is. That's not ECC. Anybody who knows about error correction, and I have to because of SpinRite, knows that's got nothing to do with error correction.

Leo: Aha.

Steve: So that was just - that was complete nonsense.

Leo: It was B.S., okay.

Steve: Yes, yes.

Leo: All right.

Steve: Error correction, you feed something in. It does a lot of work. You get a syndrome, which is an XOR mask against the data, and it also presides a position of where that mask is slid to provide the XORing which flips the bad bits into good bits. That's not what that SBox is. That SBox is a simple hash.

Leo: Cool.

Steve: So that's not ECC. But I'm glad you connected with this thought of mine, Leo, that, you know, of course they have...

Leo: There had to be a backdoor. Yes, of course.

Steve: Yes. Of course they have, they and they alone, to prevent the end of the world, can open up any of their devices.

Leo: That makes sense. And then they would super-strong secure it.

Steve: Yup.

Leo: And of course they didn't need to go through the other three exploits to get to that part.

Steve: Right.

Leo: Bad guy would, right.

Steve: Right. And they didn't even know it was being abused until Kaspersky said, uh, look what we found over here. And they're like, oh, fit, you know, [crosstalk].

Leo: But this is the problem with a backdoor of any kind is it's impossible to secure.

Steve: And that's what everyone has said. Yes, I mean, that argument that, you know, and this of course has been the argument that everyone has always used against, you know, the EU saying, oh, just give us a way in. It's like, no. It will be abused. This was. And this is really unfortunate. But I don't think Apple was wrong to do it.

Leo: No, when you put it that way it makes sense. It seems like a good thing to do. And now I'm hoping that Dr. Evil doesn't put the nuclear codes on his iPhone, unless they have another backdoor, which is possible.

Steve: Well, they haven't lost this one, Leo. This one is still available.

Leo: Oh.

Steve: All they have to do, I mean, they...

Leo: They can reenable it.

Steve: They locked a door in front of this one with an update. They can un-update any phone that they want to.

Leo: Oh, of course, duh.

Steve: Yeah. So they still have full access to A12 through A16 get out of jail free card. And this is the problem, too, because now the FBI, well, if the FBI didn't know, the NSA will now tell them. Maybe someone is listening to this podcast. Now Apple can open any, you know, now that cat is out of the bag.

Leo: Yeah.

Steve: They can no longer say what they have been saying to the courts, we have no way in.

Leo: Oh, interesting.

Steve: Yes, you do.

Leo: Oh, boy. That's interesting, too.

Steve: You always have had a way in. You just didn't tell anyone. Well, now we've found out because it leaked. It leaked from inside Apple, unfortunately.

Leo: Okay.

Steve: So we talked earlier about the interesting abuse of Internet-connected cameras by Russia to obtain attack and improved targeting of their weapons strikes inside Ukraine. I saw a bit of an update on that last week after Russian hackers successfully hijacked at least two security cameras and used those cameras' live video feeds to adjust their missile strikes targeting the city of Kyiv last week. Once Ukraine's SBU, which is their Security Service, detected the attacks, they took down the cameras to prevent their further abuse. In part of the coverage of that was the information that, since the start of Russia's invasion, more than 10,000 security cameras have been taken down across Ukraine. Ten thousand.

So hearing this, you know, I kind of paused to wonder and also to worry about the situation here in the U.S., in Europe, and with any of our allies. Back when we began this podcast, Leo, you know, when loincloths were in fashion...

Leo: We're even older than the iPhone.

Steve: That's right, baby. The idea of "cyberwar" was still squarely the stuff of fiction. From everything we hear, there are constant low-level "cyber skirmishes" going on all the time. We know that the usual suspects of China, North Korea, and Russia have talented hackers who are more or less continually poking around inside our computer networks. And for our own sake, I hope we're giving at least as well as we're getting, since a cyber standoff is in everyone's best long-term interest.

Everyone in the world, however, is pulling from the same common pool of technology. So we're probably all about equally vulnerable to each other. There's no reason to believe that the cameras we have everywhere here in the states and that Europe has through their countries are any more secure than those that Ukraine was using. And the same applies to all of our other interconnected technology. I've often wondered what I would do if I were starting out in the world today. I've always had a strong intellectual curiosity about whether I could hack other people's stuff. But doing so is both unethical and illegal. So I never have.

But participating in my country's defense and, if necessary, its offensive operations, I have to admit, that has some appeal, while also resolving the ethical and legal roadblocks. One thing is very clear. This is no longer the stuff of sci-fi. It's very real. And it appears that our countries need us. So I've said it before, I'll say it again to our younger listeners who maybe haven't chosen a career path. Get really good at this stuff; and, you know, your country needs you. Seriously. We know that our country is hiring, and this is real. And, boy, it would be a lot of fun. I don't know about the camo, though, Leo.

Leo: You don't have to wear camo. You can wear your BDUs at home and wear civvies to work. It's okay. I give you permission. General Leo says it's okay.

Steve: We have heard, however, we've heard from our listeners, oh, no, I've got to put my camo on every morning. But they did assure me, though, that they're comfortable, that they're not like stiff and starchy.

Leo: Oh, yeah. Oh, yeah, yeah.

Steve: You know, they've got some sweat rings under their armpits, so that's good. Okay. So I think this topic is important enough for me to spend a bit more time on a specific example. Last Thursday, Reuters news service published an article titled "Russian hackers were inside Ukraine telecoms giant for months." So here's some new information that was not public before that Reuters just published. They said: "January 4th. Russian hackers were inside Ukrainian telecoms giant Kyivstar's system from at least May last year in a cyberattack that should serve as a 'big warning' to the West, Ukraine's cyber spy chief told Reuters.

"The hack was one of the most dramatic since Russia's full-scale invasion nearly two years ago knocked out services provided by Ukraine's biggest telecoms operator for some 24 million users for days starting December 12th. In an interview, Illia Vitiuk, head of the Security Service of Ukraine's (SBU) cybersecurity department, disclosed exclusive details about the hack" - which I would call more than a hack - "which he said caused because, he said, caused 'disastrous' destruction and aimed to land a psychological blow while gathering intelligence." He said: "This attack is a big message, a big warning, not only to Ukraine, but for the whole Western world to understand that no one is actually untouchable."

He noted that Kyivstar was a wealthy private company with heavy investments in cybersecurity. The attack wiped "almost +everything," including thousands of virtual servers and PCs, he said, describing it as probably the first example of a destructive cyberattack that "completely destroyed the core of a telecoms operator." Later, following some investigation, on December 27th he said that they found that the hackers probably attempted to penetrate Kyivstar originally in March or earlier, much earlier last year. He said: "For now, we can say securely that they were in the system at least since May of 2023. I cannot say right now when they had full access, probably at least since November." The SBU assessed the hackers would have been able to steal personal information, understand the location of phones, intercept SMS messages, and perhaps steal Telegram accounts with the level of access they had gained.

A Kyivstar spokesperson said the company was working closely with the SBU to investigate the attack and would take all necessary steps to eliminate future risks, blah blah blah. Of course that's, you know, the PR guy at the company that got blasted. And following the major break there were a number of additional attempts aimed at dealing more damage to the operator.

Kyivstar is the biggest of Ukraine's three main telecoms operators, and there are some 1.1 million Ukrainians who live in small towns and villages where there are no other choices, no other providers. People rushed to buy other SIM cards after the attack, which created large lines. ATMs using Kyivstar SIM cards for the Internet all ceased to work; and the air-raid sirens, which are used during missile and drone attacks, also did not function properly in some regions.

Post attack forensics are made more difficult because of the wiping of Kyivstar's entire infrastructure. But Vitiuk said he was "pretty sure" it was carried out by Sandworm, a Russian military intelligence cyberwarfare unit that has been linked to cyberattacks in Ukraine and elsewhere. A year ago, Sandworm penetrated a Ukrainian telecoms operator, but was detected by Kyiv because the SBU had itself been inside Russian

systems. Vitiuk said the pattern of behavior suggested telecoms operators could remain a target of Russian hackers, and during 2023 the SBU said it had thwarted over 4,500 major cyberattacks on Ukrainian government bodies and critical infrastructure.

So, okay, again, sadly, there's no actual reason to believe that things are any different anywhere else. Ukraine is using the same technology as everyone else. As I've said, all of the evidence we have suggests that our actual security is far more soft than we would like. What's generally and thankfully missing is the motivation to abuse it. But the rise of cryptocurrency created the motivation to extort enterprises that smugly believed until then that their IT security budget was sufficient, and that the threats were being overblown. No one thinks that any longer. The last thing we need is an escalation.

Leo: Wow.

Steve: Uh-huh. I know. It is sobering. You know, we've got little poking around the edges. But as I said, I hope that we're able to give as good as we get because, you know, we only hear about attacks against us. We don't get any information about, you know, what we're doing, how we're in other people's networks. But as I said, there's a career there.

Leo: I also think there's some reluctance to go full bore on this because you really, I mean, the obvious end game is attacking infrastructure, which could be horrific to civilians.

Steve: Well, and when you do so, you're no longer covert.

Leo: Right.

Steve: So it is a use it and then lose it. So nobody wants...

Leo: Right. And the threat of retaliation is so strong because you don't need - it's not like building a nuclear weapon. You know, you need a few good hackers. And probably almost any nation-state could muster up enough hackers to be a threat.

Steve: Leo, North Korea.

Leo: Right.

Steve: Like, where did they get their education? They didn't come over here and get taught at MIT.

Leo: Actually, apparently some did. But anyway - covertly.

Steve: Whoops.

Leo: But that's the point is that this information is out there. And it doesn't cost a lot to create a Fancy Bear. So it doesn't, you know, it's really an interesting issue. It's not the end of the line, I just want to say.

Steve: You do not need huge rooms of spinning centrifuges...

Leo: No, exactly.

Steve: ...and years of time.

Leo: Right.

Steve: You know, you need literally some guy in his mother's basement.

Leo: Literally.

Steve: And Leo, let's take a break, and we're going to talk about 23andYou.

Leo: And Me. And I'm excited, or not excited, but I'm interesting being a longtime 23andMe user. And I was really kind of mad at the company because they said it was our fault.

Steve: You should be. You should be, Leo. I know, you should be.

Leo: Now, tell me how much trouble I'm in with my DNA; okay?

Steve: Okay. So things are still a mess at 23andMe. They've been hit with 30, three zero, lawsuits.

Leo: Oh, gosh.

Steve: Since last - uh-huh. Yeah, I mean, and people take their DNA as like a privacy issue. Who would have thunk?

Leo: Mm-hmm.

Steve: So anyway, this was, you know, last December was the revelation of the breach which disclosed the personal information - and this is the number that stuns me - of 6.9 million of their users. Okay, now, just to remind everyone, the story is that 14,000-some accounts were first directly compromised using simple credential stuffing, you know, reusing known, previously used passwords of each victim. I would argue that this had to

be detectable right there, but you won't see what you're not looking for, and nothing else these guys did seems particularly impressive on the IT side, so...

Leo: They weren't looking.

Steve: They apparently weren't looking.

Leo: They weren't looking. They had their heads turned.

Steve: So, okay. From there, we're told that simply using the API that was available to any logged-on user - since that's all these bad guys apparently were - the attackers were then able to siphon off the personal data of an additional, expanded, 6.9 million unbreached users.

Now, I suppose I'm still skeptical about this explanation because in my gut I find it difficult to believe that the designers of 23andMe's architecture could deliberately have set things up so that anyone logged into their system could have direct access to, on average, the personal data of 493 other members. 6.9 million divided by 14,000 is 492.857, which is the average "disclosure reach" of each of 23andMe's 14,000 logged-on users. In order to believe that this is what actually happened, 23andMe's system had to be horribly designed from the start, which is quite dispiriting.

And then, in the wake of this catastrophe of their own making, adding insult to injury, 23andMe attempted, and I'm sure you saw this, Leo, to change the Terms of Service for their users retroactively, if you can believe that, to require them to agree to settle any disputes through arbitration in lieu of other legal action. That didn't pass notice, and you can imagine that it didn't go over very well.

As we know, anyone can make a mistake. But they're directly responsible for their apparently incredibly crappy system design which, again, if they're to be believed, allowed any legitimate user to log on with their own credentials, and then have access to the personal details of, on average, nearly 500 other users who they don't even know. So, wow; you know? It looked like a good deal. I'm also a member, for what it's worth. I spit in the tube years ago because it seemed like a curiosity. Wow.

Okay. And speaking of the incentives created by cryptocurrency, the Estonian cryptocurrency platform CoinsPaid - don't think they knew who their coins were going to be paid to - was the victim of another cyberattack, losing an estimated $7.5 million worth of crypto assets. I said "another" since this is the second time this company has been hacked. The first time it lost $37.3 million in July. CoinsPaid blamed last year's incident on, guess who, North Korea. I wonder who you call in North Korea to negotiate a settlement? "Hey, how 'bout we'll give you a 10% bounty and no hard feelings if you'll return the rest?" Right. I wouldn't hold my breath on that happening.

Meanwhile, the Gamma cryptocurrency platform says it lost $6.1 million worth of assets after a threat actor abused the infrastructure of one of its providers to manipulate exchange prices, and another threat actor has stolen nearly $4.5 million worth of crypto assets from the Radiant Capital cryptocurrency platform. The technique used there was a so-called "flash loan attack." So I don't know where all this money is, you know, where it's coming from and going to. But I am sure glad that none of it's mine.

Leo: Somebody's getting a yacht.

Steve: My god, Leo.

Leo: Geez.

Steve: Oh, yeah, we lost 37 million. Well, you know, that's, you know, stuff happens.

Leo: We can make more. It's okay.

Steve: Call North Korea and ask them if they'll accept a bounty and give us 90% back. No. It's the North Koreans who are having a party and got a yacht.

Leo: Yeah.

Steve: Okay. But, you know, let's take a larger view. Stepping back, overall, during all of last year, hacking attackers made off with more than $1.8 billion U.S. worth of crypto assets, and that was across 751 individual security incidents. There is some good news here, though, since that number is way down, as in by half, from the $3.7 billion U.S. that were lost the year before, during 2022. So that's good. Either people are like pulling their money out of this crazy business, or security is beginning to get better. I mean, it certainly was the Wild West there in the beginning, you know, when I don't remember the details, Leo, but Kevin was like buying icons of monkeys or robots or something?

Leo: Yeah, he was. He was buying bored apes.

Steve: Like what the heck?

Leo: And then crypto punks. And then, because he saw the writing on the wall, he offered his own icons. They were owls. And I think he and the consortium that did this, including a number of well-known NFT people, like people, made $50 million.

Steve: And where were we?

Leo: It's unbelievable.

Steve: Ah, those youngsters.

Leo: Those kids.

Steve: Anyway, according to the blockchain security firm CertiK, whom we've quoted from time to time, last year's top 10 most costly incidents accounted for more than $1.1 of that total $1.8 billion stolen in total. So it's not like it's, you know, all thefts are equal; right? The top 10 got 1.1 of the total 1.8. And speaking of behind the curtain, the most

costly incidents were linked to leaks or compromises of private keys. So that's how these attacks, the biggest ones, happened was that people's private keys got loose. More than $880 million was stolen just that way last year.

According to TRM Labs, North Korean hackers were linked to $600 million of those total stolen assets. In other words, one-third of all the cryptocurrency lost last year, the $1.8 billion, went into North Korea. So, yeah, maybe those are some of the MIT grads that they've got over there at work. Anyway, they're not slouches.

Leo: Nope.

Steve: Oh. And lest we believe that these things only happen to people with low security awareness, an unknown threat actor recently hijacked the Twitter account of Google's Mandiant division. Right? Like the high-end cybersecurity gurus.

Leo: That doesn't bode well.

Steve: No, it doesn't.

Leo: Not the best on your rsum.

Steve: No. The account takeover was used to promote a - guess what? - cryptocurrency scam. The attack was just one of a number of similar incidents that hit many high-profile Twitter gold badge accounts at the start of the year.

Leo: I don't know if you saw this morning, the SEC, as in United States Securities and Exchange Commission, was hacked on Twitter to put an announcement they have just approved Bitcoin ATF spot purchases, which they were quick to point out they hadn't, and that was a hack. So...

Steve: Nice. Nice.

Leo: I don't, you know, at this point I don't know if I want to blame the accountholders. I don't know if I blame Mandiant. Who knows what the Twitter security status is these days.

Steve: Good point. Good point. It could easily be someone inside.

Leo: Yeah, yeah. It could be [crosstalk].

Steve: I saw a little blurb...

Leo: Who knows?

Steve: Well, I saw a little blurb that - it just passed by on my phone, saying that Elon had used illegal drugs and that executives at Tesla and SpaceX were concerned.

Leo: I brought this up, this was a big...

Steve: Hey, you know.

Leo: This was a big story in The Wall Street Journal on Sunday for some reason. I brought this up on TWiT on Sunday. And the panel said, "Everybody knows that. We've known that for years." Oh, my god, he's using illegal drugs.

Steve: Yeah. And I would just note that there wouldn't be executives at Tesla or SpaceX, neither would exist were it not for Elon.

Leo: Yeah.

Steve: And who knows? Maybe as a result of some of those illegal drugs.

Leo: Maybe there's a secret. Maybe that's his secret, yeah.

Steve: That's right. Okay. So just how seriously, and legally seriously, do we take the term "war" when it's used in the phrase "cyberwar"? Remember back nearly seven years ago - and Leo, you're going to get a kick out of this because you said, "Whaaaaaat?" at the time. In 2017, the monster American pharmaceutical company Merck suffered a serious ransomware cyber breach by the NotPetya group. What stunned us at the time, and as I said, I remember your, like, "Whaaaaaat?" was Merck, who was carrying significant cyberattack insurance, was claiming that the attack, which they said affected 40,000 of their PCs, would cost them $1.4 billion, with a "B," to clean up.

Leo: And I still say "Whaaaaaat?"

Steve: And it's like...

Leo: You're kidding.

Steve: Oh, wow. So naturally, their cyber insurance carriers, of which there were three, were none too pleased by the prospect of having to fork over $1.4 billion to, what, finance Merck's physical replacement of their entire PC inventory? I mean, it's not as if the machines melted. So, what? I would love to see the justification for this. Anyway, we covered this of course at the time. And recall that the three Merck insurers who were on the hook for this were attempting to get out of their policy obligations by claiming that an exemption applied in the case of "Hostile/Warlike Action." That's literally in quotes. That's what it says. It's a commonly present policy exclusion.

So the question that has ever since then been working its way through our U.S. legal system was whether or not a "cyberattack" could and should be considered to fall under this standard "Hostile/Warlike Action" policy exclusion. And of course this would be precedent-setting since devastating cyberattacks are no longer theoretical, and insurance to make enterprises whole in the wake of one have become crucial and ever more costly, both in premium and in reimbursement.

Leo: Yeah, I think that that clause, by the way, is pretty common. Acts of god and acts of war are often exempt, you know, because...

Steve: Yup, yup, exactly. And so the question is...

Leo: We can't be expected to insure against that.

Steve: ...is this an act of war? So until last week, when New Jersey's state Supreme Court was set to hear oral arguments from both sides, a lower New Jersey appeals court had ruled that Merck was entitled to half of what they were seeking under their policy coverage. In other words, $700 million. The insurers still wanted to pay less, and Merck wanted more. But just hours before oral arguments were set to begin, the parties announced that they'd reached a settlement, though the terms of that settlement have not yet been disclosed. Given that Merck is a publicly traded company, owned by its stockholders, I would imagine that the terms of the settlement will eventually become known.

Interestingly, in amicus briefs filed before the scrapped oral argument, national associations for big business, manufacturers, and corporate insurance litigators had all argued the court to uphold the ruling that cybercrime did not fall under an insurer's "Hostile/Warlike Action" policy exclusion and plant a national flag, as they put it, on this issue to benefit insured businesses. But the decision, it turns out, was not cut and dried. Dueling briefs from international law scholars debated whether foreign-linked hacking against corporations is warlike action.

The takeaway for insurers should probably be that they are going to need to stand behind their cyberattack policies, and those paying for coverage by those policies should probably demand some explicit clarification from any policy that contains such potential wiggle room language because we haven't seen the last of cyberattacks. And unfortunately, we didn't really get the hard precedent set that many people were hoping one way or the other that this case would create. So, you know, last minute settlement, and it's like, okay, fine. Move along.

Last Tuesday, LastPass posted a blog titled "LastPass Is Making Account Updates. Here's Why." So I'm just going to share the opening paragraph because, you know, we're all well able to read between the lines. But LastPass said: "You may have noticed that lately we've been asking our customers to make some changes to their LastPass accounts. These changes include requiring customers to update their master password length and complexity to meet recommended best practices and prompting customers to re-enroll their multifactor authentication, among other changes.

All of these changes are intended to make our customers more secure, and we want to share additional context about the evolving cyberthreat environment that's driving these requests so customers can better understand why these changes are important. To do this, we'll address some of these recent changes, and explain what threats are driving them, and how these updates are designed to help."

Okay. So anyway. And it goes on. My only complaint, of course, is that it's closing the barn doors after the horses have all run off; you know? This would have been very nice to see several years ago, and history would have been written differently had that been the case. This effort is clearly an attempt to respond to the theft of the master data vault and to mitigate future disasters. Requiring everyone to "reenroll" their multifactor authentication basically means get a new private key at each end so that if that has also somehow been compromised, nobody, you know...

Leo: Okay, I was wondering about this. So it's the theory being that the secret, which is key to a time-based one-time password, has also been leaked.

Steve: Right.

Leo: Okay.

Steve: Or could be. So they're just saying, let's, you know...

Leo: Start over. A new secret.

Steve: Yeah, no reason not to.

Leo: Yeah.

Steve: It's sort of the equivalent of like people saying, oh, change your password just because we think it's time. It's like, uh, okay.

Leo: Yeah, I've always felt like, wait a minute, if my password's good, why am I changing it? But, yeah.

Steve: And as we know, that advice has been reversed now. It is no longer thought...

Leo: Yes, NIST took that back, yeah.

Steve: Yes.

Leo: But it isn't a bad idea to occasionally redo your two-factors, it sounds like.

Steve: I would agree. There's, I mean, [crosstalk] been some leakage.

Leo: That secret is kept in the clear, generally.

Steve: Well, that secret is at their end and in your authenticator.

Leo: Right.

Steve: So you don't want, you know, if there's been a breach, I mean, this sort of says that maybe they also lost their two-factor authentication data. And, you know, instead of just their user vault data.

Leo: It's hard not to read between the lines and say, what, you had another problem here?

Steve: Or this was a little more extensive than you said it was.

Leo: Maybe. Wow. Holy cow.

Steve: Yeah. Now, as for this new 12-character minimum password complexity requirement, that only makes sense. And I want to talk about that a little bit. What should really be happening at this point across the Internet, and I mean everywhere, is that users should begin to be forced to increase the security of their logons. It should not just be happening at scattered sites in the wake of devastating attacks. Any service that supports logons where a breach could have devastating consequences for its users should start doing the same. Users really want to reuse "their," and I have that in air quotes, personal password everywhere.

Leo: Yeah.

Steve: You know? Monkey123 forever. That's still today, 2024, that's the typical behavior. Obviously, not among this podcast's listeners, but pretty much everywhere. Never underestimate the strength of inertia. Users do not want to change, and they will not change unless and until they are forced to. We now have the technology to enforce password complexity rules on the user in their browser thanks to client-side JavaScripting. Users hate password requirements. Why? Because those requirements prevent them from using their favorite universal pet password everywhere. And those requirements mean that they may need to deal with unique passwords per site, at least to some degree. The question is whether the Internet should continue to let them.

If the Internet continues to allow this past behavior, it will never change. We all know that. Why would it change? Users will need to be forced. But every site is understandably terrified of doing that because they don't want to alienate their users. The rational solution is for sites not to pretend that their users have security that does not exist. If a site is not going to enforce a sufficiently high level of password complexity, then it should not assume that its users have any actual logon protection, and it should act accordingly.

Or perhaps the client-side JavaScript, which can see the user's plaintext password for itself before it is locally hashed and then sent to the server, should examine, the JavaScript should examine the password's complexity and send along a complexity ranking of the hashed password's strength. Then a site that does offer some sensitive services could explain to its logged-on user that the password they are using is fine for logging on; but, for their protection, a better password will be required before they're

allowed to do anything sensitive that they would not want hackers to be able to do in their name.

So I suppose I'm saying that the industry has clearly been dragging its heels because it has not been forced to change, and this has allowed users to, in turn, drag their heels and continue with habits that no longer serve their best interests. Web portal designers would be well served to keep this in mind. So good thing that LastPass said, okay, we're going to make some changes. But gee, you know, had they been keeping up with current practice and recommendations, that would have been happening all along. And talk about a base of users who would understand. I mean, it's one thing, you know, to ask logons at Granny's Cookies site to, you know, do complex passwords. But LastPass? Obviously people would be willing to do this.

Leo: But also talk about inertia. I mean, who's still using LastPass except somebody who's said I'm not going to change? I'm still - I'm not going to change.

Steve: True.

Leo: I mean, I would think many people would have done what you and I did and switched.

Steve: Yup. Remember the podcast title was "Leaving LastPass."

Leo: "Leaving LastPass," yeah. I mean, I guess the theory is, well, now LastPass will be more secure than anyone because they got bit. And so they're going to do everything they can not to get bit twice. I guess.

Steve: Except that we know that they thought they were secure; right?

Leo: Right.

Steve: I mean, that's always the conundrum. They thought they had this covered.

Leo: Right.

Steve: And whoops.

Leo: Yeah, whoops.

Steve: So as of - just a little quickie. As of the beginning of 2024, because this fascinates me, Windows 10 is holding onto two-thirds of the desktop, while Windows 11 has been gradually creeping upward from about 16% to now 26% of desktops across 2023. People generally like what they have, and again, inertia.

Leo: Yeah.

Steve: We should really rename the podcast.

Leo: The Inertia Show.

Steve: Yeah. I love what you and Jeff do, This Week in General. That's a good title.

Leo: Yeah, yeah.

Steve: And I think, you know, This Week in Inertia would be...

Leo: I'm not changing. I'm happy.

Steve: Yeah, just pry it out of my...

Leo: This from the guy who's been on Windows 7 since, you know, before the Stone Age. But okay.

Steve: Sitting in front of it right now.

Leo: Okay. Okay, Steve. I notice Windows 7 is holding strong, by the way, has not gone down.

Steve: I think that's the yellow line on that chart; right?

Leo: Yeah, yeah.

Steve: Yup. Yeah, that's me. I'm there. In fact, there's a little uptick there, no, I can install another one. No. I'll be moving to Windows 10.

Leo: Yeah, there's Steve, he's installed another one.

Steve: I'm Windows 10 in the evening. I'm Windows 7 during the day because it works great.

Leo: It's like a mullet, you know, Windows 7 in front, Windows 10 in back.

Steve: Is that there then a mulligan?

Leo: Yeah, no, a mullet, you know, that's the haircut.

Steve: Okay. So again, the following bit of Google tracking news made a lot of headlines recently, so I thought I would just mention it, too. Remember back in 2020 when Google was found to be tracking users in "incognito mode," and this resulted in a ridiculously large class action lawsuit. And just for the record, everybody, I know you already know it, but Leo and I are not generally fans of ridiculously large class action lawsuits.

Leo: No.

Steve: Because it's just attorney-enriching. So the news is that the lawyers on each side of this dispute have reached an agreement, as happens more often than not on the eve of such cases moving forward to trial. When you're big, you tend to be a target of attack since the presumption, at least among scummy attorneys, is that it's worth some money from the big guy just to make the nuisance lawsuit go away because they're going to spend more money defending this nonsense than they are just saying, fine, here, buzz off. At the same time, unfortunately, being big also increases the tendency of companies to throw their weight around, bully others, and imagine that they can get away with whatever behavior they want.

Thursday before last, U.S. District Judge Yvonne Gonzalez Rogers put the trial that had been scheduled for this case on hold in California after attorneys said they had reached a preliminary settlement. Judge Rogers had previously rejected Google's bid to simply have the case dismissed, saying she could not agree that users consented to allowing Google to collect information on their browsing activity when in incognito mode. The class action, which was filed in 2020 by the law firm Boies Schiller Flexner, and that's "Boies" as in David Boies...

Leo: Oh, yeah.

Steve: Uh-huh. Don't mess with David.

Leo: Don't mess with David, nn-nnn.

Steve: Unh-unh, claimed that Google had tracked users' activity even when they set the Google Chrome browser in incognito mode. It said this had turned Google into an "unaccountable trove of information" on the user preferences and "potentially embarrassing things." It added that Google could not "continue to engage in the covert and unauthorized data collection from virtually every American with a computer or phone." Oh, and I forgot to mention, the class-action lawsuit, $5 billion with a "B."

Leo: Oh, boy.

Steve: Dollars.

Leo: That's actually low for class-action.

Steve: More than the cryptocurrency, all the cryptocurrency lost in the last two years combined.

Leo: Oh, my.

Steve: I'm sorry, you were saying...

Leo: I think that's actually low for a class-action lawsuit. But maybe I'm...

Steve: Against Google.

Leo: Yeah. Apple just decided to settle its half billion dollar lawsuit. People are getting $92 each. But you had to - so I guess half a billion compared to five billion. Wow.

Steve: Yeah, 10 times. Anyway, so I think my take on this is that it's a case of the fine print coming back to bite you. Google claims that the users of their incognito mode were duly informed and knew that tracking was still occurring even though the post-incognito mode residues from their browsing such as history and cookies were not retained. Apparently some of their users disagreed and felt betrayed. So anyway, just, you know, another lawsuit settled. The industry moves on. Maybe this creates some pressure on Google to change this aspect of their behavior. I don't imagine most people spend much time in incognito mode. They only jump in to do something that they don't want to have...

Leo: Well, it's really hide it from your spouse mode.

Steve: Yeah.

Leo: And that's what people probably try to explain to people. But they didn't do a very good job of it.

Steve: So as I said, this was a rather thin news week. I think we made the best of it, talked about a lot of interesting stuff. I think I may be onto an interesting independent analysis of the privacy protections created by Google's Topics API and other components of Chrome's privacy sandbox. If it pans out, I'll have that for next week.

Leo: Steve, you're up.

Steve: So as I mentioned at the top, we had a listener, Carl Smith, who sent a tweet. He said: "@SGgrc, how has Operation Triangulation not received more press coverage?" He says: "This is huge," and then four exclamation points. And of course as I mentioned, I agree with Carl. I suspect that Apple is benefiting hugely from the fact that while what's really going on here, which obviously everyone who listens to this podcast understands, is truly monumental, it was also "patched" with yet another iOS update, and the public at

large has no way of discerning that this one is any different from any of the others that preceded it through the years.

And really, you know, some Russian security analysts found something they presented during the Chaos Communications Conference in Hamburg? What's that? You know, that's not going to make the nightly news. The popular news media cannot begin to explain this to the average consumer. So I bet the news producer just says "Talk about the weather," while Apple breathes a huge sigh of relief in the knowledge that they didn't take any PR hit from what might have been a disaster for them.

Leo: And let me put a plug in for you, Steve. That's why people listen to this show, that's why they listen to TWiT, because we can cover technical stuff in a way that's intelligent, and so that you get that information. If you're not a member of Club Twit, support what Steve's doing, twit.tv/clubtwit. That's all I'm going to say. But this is why you need us, and we need Steve. Sorry, didn't mean to throw you off there. I just, I couldn't resist. It's like, this is why we do what we do, because the mainstream media's not going to cover this stuff.

Steve: No, never. And even the tech press, it waters it down. You know, I mean, they're in a hurry. They've got lots of other stuff to do. You know, this is just one of a gazillion stories that they're trying to cover. For us it's like, whoa, hold on, stop the presses. This is a podcast today.

So Vjirasek said: "Hi, Steve. Great work on SN-955." That was last week. He says: "I am wondering why Apple has not implemented ROP attack protection similar to what Intel has done. Would this break the chain of this sophisticated attack? Also concerning to see that Apple has left the back door in the SoC to get in. Thank you for your hard work."

So he's referring to the use of ROP, Return Oriented Programming, which we mentioned and talked a bit about last week. It's a "living off the land" practice of using bits of code that's already present in the target device to obtain the effects that are needed for the attack. I'm certain Apple has Return Oriented Programming attack prevention in place, as must any highly secure attack-prone operating environment these days. But while ROP makes attacks far more difficult by scrambling and randomizing the memory locations occupied by code, that code is still present in memory. It's just been moved at load time into initially unknown locations.

One of the things the Kaspersky guys noted was that a huge amount of the malware bulk was spent examining the system's memory. Now we know why. So that would likely have been code designed to locate the bits of executable code that they needed in order to execute their exploit. So while ROP can make attacks much more difficult, it's also not a perfect solution. We still don't have one except let's not have any bugs. And we certainly haven't gotten there yet.

Two notes. Robin Ramaekers, he said: "On SN-955 you had Ethan Stone giving you a quick note that he had problems closing the Edge browser. While it is true that when you close the window, the Edge processes keep running, there is an easy way to close the browser. If you click the ellipsis in the upper right you find the option" - it's actually way down at the bottom, the very bottom - "to close Microsoft Edge all the way down. This is in contrast to clicking the X close box at the upper right, which may only close the user interface."

And someone calling himself Warwagon, who posts often, he's a well-known contributor to GRC's newsgroups, he wrote: "Here's a fix for Edge running in the background. Open Edge. Click the three dots in the top right and click Settings in the dropdown menu. In

the top left do a search for 'startup,'" S-T-A-R-T-U-P, one word. And for me I had to wait a while. It takes Edge like a surprising long time to produce any results, but it does. He says: "On the right you'll see two options, 'Startup Boost' and 'Continue running background extensions and apps when Microsoft Edge is closed.' Turn both of those off. That should remove Edge from the Task Manager when it's closed unless you also have some website notifications enabled."

Okay. So those are some great suggestions. Here's what I found: When I closed my instance of Edge, like just using the standard X close box in the upper right, it did not continue, my Edge did not continue running anything in the background. Following Warwagon's advice, I found that I had the Startup Boost option turned off already. And that's what made the difference.

Leo: Ah.

Steve: With Startup Boost turned on, Edge does not close unless you open the ellipsis menu and choose "Close" down at the bottom.

Leo: So that's what Startup Boost means, just stay, just don't go away. Never stop.

Steve: Yes. Yes. Just clicking the UI's X box only closes the UI, and it definitely leaves a bunch of processes. I initially had 16 Edge processes running. It whittled itself down to 10 eventually.

Leo: That's crazy.

Steve: But still, it's like, it's just sitting there squatting on RAM and obviously taking up some time from your machine. So anyone who wants to truly close Edge will need to either turn off Startup Boost or use the ellipsis menu and select Close down at the bottom of that menu, if you've got Startup Boost turned on. If you have ample memory and would rather have Edge pop onto the screen instantly because basically it's always actually running, then you can turn on Startup Boost, and that's what you'll get. So thank you guys for the feedback. I'm glad that we got some closure there.

Thomas Tomchak, he tweeted, he said: "I'm guessing this is out of scope for how SpinRite should be used, but I tried to boot my Windows VM into SpinRite because I wanted to run it on the internal drive of the VM. I first tried this with 6.0, and it booted up. So I then downloaded the pre-release Windows EXE and created a new ISO. I uploaded that to the vSphere host, attached it as a CD, and told the VM to boot using the BIOS. This time it went right into the attached screen."

And I have to tell you, when this thing came up in Twitter, what is it, my heart went into my throat or something, or my hope sank. Anyway, I thought, oh, no. Anyway, he said: "I'm sending it in case it's of any help to you, but understand I'm using the software in a way it wasn't intended to be used. Hopefully it helps in some way."

So anyway, yes. As I said, when I saw Thomas's screen capture showing that SpinRite had intercepted the processor's attempt to execute an illegal instruction, and I saw at the very top of the screen he was running the latest release 5.06 which is believed to have no such remaining loose ends, my first thought was "Oh no, now what?" But then I was

greatly relieved to read that this was the result of him attempting to run SpinRite 6.1 within a virtual machine.

Okay. I decided to share this question because there has been a great deal of interest in running SpinRite in virtual machines for various reasons. So I need to discourage and disabuse everyone of that idea. SpinRite 6.0 was and is a very tame and well-behaved "generic" DOS application. By comparison, SpinRite 6.1 really is not. SpinRite now assumes that it has access to true physical hardware, and it does things like briefly switch the processor from real mode into protected mode in DOS - in DOS - then directly alters the processor's memory management segmentation registers to remove real mode's traditional 64Kb segment limitations.

Leo: So what you're saying is don't run this through VirusTotal.

Steve: Well, and it's surprising. This version is tripping zero...

Leo: Oh, that's interesting.

Steve: ...of Virus Total's - but what I will say is don't try to run this on other than...

Leo: VM, yeah.

Steve: ...real, yes, on other than real hardware. You know, after it tweaks the hardware segmentation registers, it switches back into real mode using an oversight on Intel's part. They had an original bug in the 286. And because of that, bless their heart, they never changed that behavior, which allows hackers like me - and this was also the way very large games like Doom were able to run under DOS...

Leo: Oh, yeah. Oh, yeah.

Steve: ...is this allows you basically a flat 32-bit, 4GB address space, even though you're in real mode.

Leo: Right.

Steve: So it creates a hacked but quite reliable pseudo mode known as "flat real mode." And that allows SpinRite 6.1 to talk to the AHCI driver memory-mapped I/O up at the high end of the 32-bit address space which it would otherwise have no access to, and to be able to use 16MB or even larger buffers. So it's very safe to say that SpinRite 6.1 and any kind of emulated virtual environment are going to not be seeing eye to eye.

Guillermo Garca, he said: "Hi, Steve. Just listened to SN-955 and your description of the certificate discovery tool. As I consider using it, I'm wondering how to reinstall a certificate that I might erase or delete and later realize that I need?" Okay. So there's actually a very cool solution to this. The Windows Certificates Snap-In that you, Leo, demonstrated last week, has a number, a large number of preexisting folders; and it's possible to simply drag and drop certificates between the folders.

Leo: Oh.

Steve: There already is an Untrusted Certificates folder which on my Win10 machine contains a Certificate Trust List folder and I think it had one thing in it. But if you drag a certificate from the Trusted Root Certificate Authorities folder onto the Untrusted Certificates folder, the system will spontaneously create a nice new Certificates folder underneath the Untrusted Certificates folder which can contain and document any certificates that you have chosen not to trust.

In fact, you can experiment with using this on any of the expired CA certificates that are currently in the Trusted Root Certificate Authorities folder. It turns out there's a bunch of them, and they would not be trusted anyway because they're expired. If you sort by expiration date, you'll see that brings all the unexpired ones to the top. You'll find that there are a bunch in there that would never be valid anyway. So you could just, if you want, drag them over into the Untrusted Certificates folder, which makes them untrusted and takes them out of circulation. So anyway, dragging these certificates back and forth is simple, and I would expect that to be error free. And should you ever discover that you needed one that you had dragged into the Untrusted Certificates folder...

Leo: You still have it.

Steve: It's still there.

Leo: That's awesome.

Steve: Just drag it back.

Leo: Wow. Very simple answer.

Steve: So Guillermo, thank you for asking the question. A.J. Druda, he said: "Steve, do you list on GRC.com how to lock credit at the credit bureaus? All I keep finding are paid sites that will do it for me."

Leo: Oh, don't pay anyone to do that. Oh, boy.

Steve: No. Crazy. Believe it or not, this is all so messed up that the terms "lock" and "freeze" have important and different meanings. This listener used the term "lock," how to lock the credit, but a "freeze" is what everyone wants.

Leo: Right.

Steve: And be careful not to go for a "lock," since some of the services actually charge a fee, they're allowed to, for locking.

Leo: They used to charge for freezing. The Fed's fixed it.

Steve: Right, they are no longer allowed to charge for freezing.

Leo: They used to make freezes free, but then it was like $35 to unfreeze in some states. It was crazy.

Steve: Oh, I know.

Leo: So that is a federal law requires them to freeze and unfreeze unlimitedly for free.

Steve: Yup. So I don't have a page at GRC, but Investopedia has a terrifically clear page which explains all the details and provides very good links to each of the credit reporting bureaus. I have the full, long "How to freeze and unfreeze your credit" Investopedia link in the show notes, but it's also this week's GRC shortcut of the week, so anybody can easily find it. Just go grc.sc/956. And that'll bounce you to the Investopedia page where you'll find links to, I mean, I've checked them all, directly into the freeze and unfreeze pages of each of the credit bureaus.

Leo: Yeah. And so does the Federal Trade Commission. They have a very nice page, if you don't trust Investopedia, which you should. But there's a government page also that describes all this. And I use this all the time because once you freeze it, you can't apply for credit. You know, so you've got to unfreeze it. And when I just recently bought a new car, I unfroze it for three days. I said, "Which reporting service do you use?" They said TransUnion. They make it fairly easy. They don't want you to unfreeze or freeze because it costs them money.

Steve: And did you see an automatic expiration?

Leo: Yeah.

Steve: An automatic - yes. That is the cool - I think that's the cool thing. As I mentioned a couple weeks ago, I applied for an Amazon credit card since I purchase so much through Amazon it just seemed like it made sense. But that's the first time. I already had all of my bureaus frozen. So when I went to do it, they automatically have an automatic refreeze after an expiration time that you're able to set. So I think that's very cool.

Leo: Remember, these credit reporting agencies make their money by selling your information to credit cards and others so they can make you offers. So they don't want you to do this. But the Fed said no, you have to allow this. And so they do, somewhat grudgingly in some cases. The only one that didn't have an automatic unfreeze was Experian. For some reason they said, well, if you turn it off, it's going to be off, man. So I have to now go back to Experian and turn it back on. But everybody else had an automatic unfreeze. Of course they want that; right? They don't want you to have it frozen. Oh, no, but that's right. They don't want you to unfreeze.

Steve: They want you to forget and leave it unfrozen.

Leo: That's right, yeah.

Steve: Andre Couture said: "Hi, Steve. Regarding the Picture of the Week for Episode 955." Remember that was the EU to U.S. power converter made out of paper clips and baling wire and, you know, grandma's stockings. Anyway, he said: "Well, I remember having to do something very similar many years ago while traveling to Europe for a presentation I had to give. I had forgotten the European power adapter. So I used what I had on hand and in my luggage to establish a connection. You do what you have to do; right? LOL." And I, well, I suppose if there's no other choice, then, yeah, one does what one must. I can imagine that it would be something well remembered.

And Peter G. Chase tweeted: "Re Ad Hoc Adapter." He said: "I can't possibly be the first one to point out that, while different countries may vary, the EU voltage is usually 240, while of course the U.S. and Canada are both 120. So whatever appliance was on the other end of that cord very likely got fried almost immediately."

Leo: Actually, not so. Not so. Because - and the reason I know this is your laptop, many appliances now in the U.S. are rated for 110 to 240.

Steve: Are able to handle either voltage.

Leo: They can handle it, yeah, yeah.

Steve: Right.

Leo: So that's why in many cases you can just use an adapter. I would still not recommend that method, but you can use an adapter without a transformer. You don't need to go through that.

Steve: Again, I would say, "Everyone, don't try this at home."

Leo: No.

Steve: Or in this case don't fry this at home.

Leo: Please. I beg of you.

Steve: Okay. So finally, "The Inside Tracks." I'm feeling very good about where SpinRite is today. No new significant problems have arisen for several weeks despite significant and continual testing. And those people whose drives SpinRite was previously having trouble with have all reported back in that SpinRite's latest prerelease managed to plow through those drives' known sticky spots while effectively recovering and repairing

everything that it encountered. Several have publicly stated that they've been amazed and impressed. So it very much feels as though SpinRite is back, and that I will be letting it go shortly.

To share some sense for where my recent focus has been, I spent the past few days exploring whether I could improve SpinRite's remaining time-to-work prediction. Back when SpinRite was born, in the late '80s, drives were sectored like pie slices with radials stretching out from the center, which described the region of each "sector" around the circumference. In fact, we've grown so used to using the term "sector" that it has completely lost its original meaning. The term was born when these were literally angular sectors of a disc.

The problem with this simple sectoring was that the tracks at the outside of the disc were physically longer than the tracks nearer to the center; yet back then, all tracks contained the same amount of data. If all tracks contain the same amount of data, and the outer tracks have a longer circumference, and the inner tracks have a shorter circumference, that meant that the individual bits were being written with reduced density around the outer tracks and increased density around the inner tracks.

Disk drive read-and-write electronics were originally separated from the drive in an outboard controller, and drives had no intelligence at all. They were basically just some read/write electronics and a stepping motor. But IDE drives, where IDE stands for Integrated Drive Electronics, changed that by placing a drive's read-and-write electronics onto each drive. Once that was done, the drive was able to become something of a black box. It could simply declare how many sectors-worth of storage it contained, and everything about how it worked in detail could be kept internal.

Drive designers very quickly saw that this meant they could dispense with the whole original notion of sectoring as it once was done. If the outer tracks had a larger circumference, they could take advantage of that to store more data around those longer tracks. And this also allowed them to push tracks further inward toward the center of the drive by reducing the storage bit rate so as not to be cramming too many bits into too small a circumference. This in turn allowed them to squeeze every last bit of storage into each drive and to make more complete use out of each physical disk's surface.

There is, however, one cost to that which is often overlooked, which is that the data transfer rate drops as we move inward toward the inner tracks. If we think of the beginning of the disk's storage as the outer tracks, then the end of the drive is the inner tracks where things are slower. The reason this matters to us today is that - or especially to me, but also to SpinRite's users - is that SpinRite's original remaining time estimation system assumed a uniform data rate across the entire drive. In other words, it performs a linear estimation. It continually monitors the total elapsed time required to get however far along it has and projects its completion time assuming that the rest of the drive will be the same as the average of everything it has seen so far.

Now, that was accurate, and it worked well for SpinRite versions 1 and 2 and 3, but it has become less and less true as the end of drives have become slower and slower as advancing technology has allowed them to push more data closer in to the disk's center where tracks are the shortest. The result is that for today's spinning drives, SpinRite's estimation will always underestimate the total time it will require. So as I said, I've spent the last few days looking closely into this to see what I might be able to improve. I've learned some interesting things that I thought I'd share while they're still fresh in my mind. What I found, after examining a handful of different multi-terabyte spinning drives, is that the ends of those drives have half the performance of their beginning tracks.

Leo: Wow. That is a big drop-off.

Steve: Half, right.

Leo: We always, I mean, this goes back to the old days when I'd always put my swap drive at the beginning of the hard drive, right, because it was faster.

Steve: Exactly.

Leo: Than the internal drives.

Steve: We did know that that was the case. Now, okay, at first blush that sounds awful; right? But the decrease in performance is not linear. What we're really looking at is area rather than circumference. And as we know, area changes with the square of a circle's radius. What this means is that while a drive's data transfer performance does steadily decrease as we move inward toward the drive's end, the decrease is very gradual until we get much closer to the end of the drive, where it finally begins to drop significantly.

Okay. So let's put some numbers to this. In general, SpinRite's current linear estimator takes about a minute to stabilize, which is to say it needs 60 seconds of operation to have established a sufficient baseline of work in time and distance to settle into a prediction that no longer varies. And what I found through lots of experimentation on many different contemporary spinning drives is that it's necessary to add an additional 30% to SpinRite's initial front-of-drive-only linear estimation. So, for example, to make the math easy, say that SpinRite predicts a 10-hour run for a drive. The actual running time, due to the very end of the drive being much slower, will be 13 hours, so 10 plus 30% of 10. Did I say 30? Thirteen, 13 hours. So you're adding three hours to SpinRite's initial prediction of 10.

Okay, now, here's another interesting factoid that falls out from the math and which I've verified multiple times experimentally. The first 60% of the drive requires exactly half of the total running time. So the last 40% of the drive requires the second half of the total running time. Or expressed another way, whatever length of time is required for SpinRite to get 60% of the way through the entire drive is the amount of time that will be required for it to finish.

Now, I'm not certain yet exactly what I'm going to do with this information, but I needed to gather it to know what I was dealing with. This obviously does not apply to solid state storage since it's not spinning, or even to shingled magnetic storage, you know, SMR format drives, since both of those technologies track when memory has been written to, I should say when and if memory has been written to them, and so they don't read anything from their media when SpinRite checks to see what's there if nothing has ever been written, as is often the case at the end of those drives. So, and we've seen this. We've seen that later portions of those drives, both SSDs and SMR spinners, appear to be performing much faster at the ends than at the front where they have stored data. They don't slow down as we go along; they speed up.

And since drives are supposed to be "black boxes" which we just trust with our data now, there's no requirement for any drive to declare what technology it's using. And many, if not most, do not give SpinRite any indication of what lies beneath their interface. The 30% rule could just be a common rule of thumb for SpinRite's users. At least initially, they know better than SpinRite knows whether they're testing a spinner or a solid state drive. So the rule of thumb for spinning media would be to start SpinRite, give its

predictor a minute to settle down, see how long it expects to be running, then add 30% to that. And that's a pretty good indication of where you will be with a spinning drive.

The other thing I'm considering, and I think I'm probably going to do it, is changing SpinRite's label on the screen. It's currently right-justified, and it's got a bunch of spaces in front of the word "time." I'm thinking I'm going to change it to est. time, as in estimated time. Then, as soon as SpinRite gets to the 60% point, it will have acquired sufficient awareness of the drive, maybe having seen a gradual decrease in performance over that span of time, to get to 60%, to be able to reliably determine that the drive is spinning, and that as much time remains as has been spent so far. So at that point, it would adjust its timer and change the "est." to, I don't know, "real time" or "true time" or "good time," you know?

Leo: Good time's good, but confusing.

Steve: So that at that, yeah, at that point, as soon as SpinRite said "true time," then you would know that it was at the halfway point, and SpinRite would be able to project what its probable completion time would be much more accurately.

Leo: Yeah, makes sense.

Steve: So in any event, after I'm finished with this podcast ere today, which I will be in about one sentence, I plan to make those final changes, after which I believe I'll finally be content to declare SpinRite 6.1 finished and ready for the world.

Leo: Now, will you wait for another episode of Security Now! to announce that? Or will you just do it?

Steve: The nature of this is sort of a soft event. For example, I'll take - right now there's all this prerelease jargon all over the UI. So I'll take that out so it no longer says "prerelease." I'll declare an actual release candidate. I need to just let a day or two go by, or five, to like make sure, like see if anything happens. You know, who knows maybe SpinRite...

Leo: No showstoppers.

Steve: Yeah. Like SpinRite works better if the word "pre" is in front of release.

Leo: You never know. That would be a weird regression, but you never know.

Steve: It's computers, yeah.

Leo: I know. Turns out that was an important part of the lookup table. Oh, no. Well...

Steve: And we actually have an example from our own experience. Nobody knows to this day, I don't know, why SpinRite 5.0 is better at recovering data from diskettes than SpinRite 6. It is, like, spooky. I've stared at the 6.0 code. I didn't change anything. Actually, I kind of remember that I did, but I don't remember what it was. And but then I've gone back and looked, and I've got this 5.0 source, and I've got the 6.0 source. They look the same. But spooky, you know...

Leo: Something happened.

Steve: There's like for some reason - now, the good news is floppies are gone. But we actually do routinely recommend to 6.0 users who are having problems, like really need to recover data from a diskette, use 5.0. And all 6.0 and 6.1 owners have access to 5.0 if they really do need to recover something from a diskette.

Leo: Isn't that funny.

Steve: 5.0 is better, and no one knows why.

Leo: Computers. I tell you.

Steve: There's a little spookiness going on.

Leo: It's a great mystery. Steve Gibson, all the spookiness happens at GRC.com. While you're there pick up your copy of SpinRite 6. You'll get 6.1 pretty soon, I would guess. Free, free upgrade for all 6.0 buyers today at GRC.com, world's best hard drive maintenance and recovery utility, floppy disks and SSDs. So the works.