Stamos on "Microsoft Security"

Description: What changes will the EU's soon-to-be-in-force Digital Markets Act be bringing to Apple's traditional iOS policies? What OS is ransomware unable to infect? What has HP done now with their printer ink policy? How many stolen user database records will fit in 12TB? Can't you just delete that incriminating chat stream? Did Mercedes-Benz leave their doors unlocked? What's the latest on ransom payments rates? And after entertaining some questions from our terrific listeners and a long-awaited announcement from me, we're going to take a look at Alex Stamos's reaction to Microsoft's most recent security incident response.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-959.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-959-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here with lots to talk about. Apple's response to the EU Digital Markets Act. Yes, the browser ballot is back. How many stolen user database records will fit in 12TB? I'll give you a hint, it's more than the total number of people in the entire world. And finally, Alex Stamos explains what we knew, what we suspected all along. Microsoft has not been fully forthright about this most recent data breach. Learn how it may affect you. All coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 959, recorded Tuesday, January 30th, 2024: Stamos on "Microsoft Security."

It's time for Security Now!. Hello. I know, you waited all week for this guy to show up. Here he is. All you have to do is say his name three times - Steve Gibson, Steve Gibson, Steve Gibson - and he appears. Mr. G.

Steve Gibson: I actually do have clicking one's heels together three times later in the show.

Leo: Oh, my god. Great minds.

Steve: So what do you know. That's right. I'm sure that explains it, Leo. That's definitely how it happened. Okay. So we're here, last show of, what month is this, still January, with a bunch of stuff to talk about. We're going to answer some questions, as we always do. What changes will the EU's soon-to-be-in-force Digital Markets Act be bringing to Apple's traditional iOS policies? I know that's been a big topic for many of the podcasts on TWiT recently. There are a couple specific aspects that affect us that we'll be talking

about. What OS is ransomware unable to infect? I could give you, not 10, not 20, not even 50 questions, or attempts to answer that. And I think you wouldn't pick up on this interesting result.

Leo: Except that I want to use that OS.

Steve: Uh, well, you will until you hear what it is.

Leo: Oh.

Steve: What has HP done now with their printer ink policy? More nonsense. How many stolen user database records will fit into 12TB? We have the answer. Can't you just delete that incriminating chat stream? Maybe not. Did Mercedes-Benz leave their doors unlocked? What's the latest on ransom payments rates? And after entertaining some questions from our terrific listeners and, yes, a long-awaited announcement from me, we're going to take a look...

Leo: What? That was a little subtle tease.

Steve: I bet you can guess what it is. That one you only need one answer for the question.

Leo: I won't say anything. I won't say anything.

Steve: We're going to take a look - all of our listeners know, too. We're going to take a look at Alex Stamos's reaction to Microsoft's most recent security incident response.

Leo: Oh, I can't wait.

Steve: Update, oh, it's fun.

Leo: I can't wait.

Steve: So, yeah, it's definitely earned itself our title topic for the week. And of course we've got a fun Picture of the Week. So I think another great podcast for our listeners.

Leo: I'm a big fan of Alex's. He's been on TWiG many times.

Steve: Yeah.

Leo: And TWiT. And he's just so great. He has a real job now. They sold the Krebs Stamos security firm.

Steve: Yup.

Leo: So he's got to work for a living.

Steve: Got absorbed into SentinelOne.

Leo: Yeah.

Steve: Yeah, he's now Chief Trust Officer at SentinelOne.

Leo: Great, that's great. Oh, I can't wait to hear what he has to say. Steve, you got a picture for me? I haven't looked.

Steve: I do.

Leo: I haven't looked.

Steve: This week we answer the question, how do you get hipsters to obey those "Keep Out" warnings?

Leo: Well, now, wait a minute, let me see? Danger. Danger, danger.

Steve: And you can see we have a hipster approaching in the distance, and there's no way she's going to cross that yellow tape line because the sign makes it clear: "Danger. No Wi-Fi beyond this point."

Leo: And there's also, it looks like, another one. I wish we could read this other one. It starts "Danger, Lost Art," but I don't know...

Steve: Yeah, I wondered what that other one down in the corner said.

Leo: Something, Lost Art something In Progress is what I think it says. That's hysterical. So clearly this is a coffee shop with a sense of humor, I think.

Steve: Somebody's got a sense of humor. It's got like sort of an SR license number at something or other.

Leo: Yeah, that's all made up, yeah.

Steve: So, yeah. And, yeah, I don't know what that is. But anyway, I thought that was great. And once again, thanks to our listeners for providing me with a constant flow. I've got some in the queue that are really wonderful, too.

Leo: This is really great. I love it.

Steve: Okay. So in perfect timing, following on from last week's discussion of Mozilla's complaints against Apple, Google, and Microsoft, we have the news, initially reported by The Verge, that Apple will be changing their browser engine policy for the EU, the European Union. Under their headline, Apple is finally allowing full versions of Chrome and Firefox to run on the iPhone. And I'll have a lot more to say about this after I share what The Verge said.

What they said was: "With iOS 17.4" - and what are we on, .3 now I think, right, or .3.something. Anyway, we all updated last week because there was a zero-day, the first one of the year that Apple had to remediate.

Leo: It's hard to keep track. We're on, yeah, 17.3, yeah.

Steve: Yeah. We had 20 last year, and so we're at the first one here at the end of January.

Leo: Amazing, yeah, yeah.

Steve: Anyway, "Apple is making a number," they wrote, "of huge changes to the way its mobile operating system works in order to comply with new regulations in the EU. One of them is an important product shift. For the first time, Apple is going to allow alternative browser engines to run on iOS but only for users in the EU."

Verge says: "Since the beginning of the App Store, Apple has allowed lots of browsers, but only one browser engine: WebKit. WebKit's the technology that underpins Safari, but it's far from the only engine on the market. Google's Chrome is based on an engine called Blink, which is also part of the overall Chromium project that's used by most other browsers on the market. Edge, Brave, Arc, Opera, and many others all use Chromium and Blink. Mozilla's Firefox runs on its own engine, Gecko. On iOS, though, all those browsers have been forced to run on WebKit instead." And of course this is exactly what we were talking about last week with Mozilla's complaint, that's like all we can do is put a skin on this thing.

Anyway, "Which means," writes The Verge, "many features and extensions simply won't work anymore." Actually, never would work. "That changes with iOS 17.4." But does it? We'll see. "Anyone building a browser, or building an in-app browser for their app, can use a non-WebKit engine if they wish." And those are the important three words. "Each developer will have to be authorized by Apple to switch engines 'after meeting specific criteria and committing to a number of ongoing privacy and security mitigations,' Apple said in a release announcing the change, at which point they'll get access to features like Passkeys and multiprocessing. Apple's also adding a new choice screen to Safari so that, when you first open the browser, you'll be able to choose a different default if you wish.

"Apple is only doing this," they write, "because it is required to by the EU's new Digital Markets Act (DMA), which stipulates, among other things, that users should be allowed to

uninstall preinstalled apps including web browsers that 'steer them to the products and services of the gatekeeper.'" You know, meaning the platform provider. "In this case, iOS is the gatekeeper, and WebKit and Safari are Apple's products and services. The same section of the DMA also means Microsoft has to let people disable Bing web search and uninstall Edge, and it will cause other changes, too." Yeah, we'll see how that goes.

"Even in its release announcing the new features, Apple makes it clear that it is not pleased: 'This change is a result of the DMA's requirements, and means that EU users will be confronted'" - oh no - "'with a list of default browsers before they have the opportunity to understand the options available to them.'" Oh, boy. "'The screen,' continues Apple, 'also interrupts EU users' experience the first time they open Safari intending to navigate to a webpage.'" Oh, the horrors. "Apple's argument for the App Store has always amounted to only Apple can provide a good, safe, happy user experience on the iPhone. Regulators don't see it that way, however, and Apple's furious about it."

The Verge says: "Again, these changes only apply for iPhone users in the EU. Apple says it allows European users to venture forth, to travel without breaking their browser engine. But it will make sure only accounts belonging to people who actually live in the EU will get these new engines." Or at least that obstructive experience of having to choose one. "Elsewhere in the world, you'll still be getting WebKit Chrome" - oh, yeah, right, WebKit's look - "and WebKit itself. Apple argues, without merit or evidence," says The Verge, "that these other engines pose a security and performance risk, and that only WebKit is truly optimized and safe for iPhone users.

"In the EU, we're likely to see these revamped browsers in the App Store as soon as iOS 17.4 drops" - this is still The Verge talking because I don't think so - "as soon as iOS 17.4 drops in March." As in, like, you know, five weeks from now. "Google," they write, "for one, has been working on a non-WebKit version of Chrome for at least a year." Which is really interesting. Makes you wonder what they knew that we didn't. "European users are about to get a serious browser war on their iPhones." Or maybe not.

Okay. So before The Verge's piece ended when they suggested that users would see new browsers in March, I was already, as I was reading this the first time, shaking my head, since porting an entire browser engine to a new OS is no small task. Putting a browser skin over the WebKit engine, which is what everyone else has done until now, is entirely different from running Chromium's Blink or Firefox's Gecko engines under iOS.

And if this will only be allowed for users having accounts based in the EU, I'm wondering why anyone really cares. The branding skin is all anyone sees. And I was watching some of MacBreak Weekly, your previous podcast today, Leo, and I saw your very Apple-knowledgeable co-hosts, you know, nodding their heads, like Andy was like - or and Alex, yeah. You know I don't care what's under there.

Leo: Right.

Steve: So, and again, the branding skin is all anyone sees. And I mentioned last week I use Firefox on my iPhone and multiple iPads. To me, it looks and acts like Firefox, and I appreciate its various features. For example, it's possible for my login to my Mozilla account to function and then sync tabs across my various devices, including my iOS devices. Just now, when I picked it up to double check on that, it asked me whether I wanted to pick up editing this podcast's Google document on my iPad. So again, you know, and we know that, for example, add-ons that are Firefox add-ons have no prayer of operating there. But it's still, you know, it's got the Firefox little logo, and it looks like Firefox.

So I presume that compatibility with the EU's DMA, their Digital Markets Act, will mean that the sort of link-stealing behavior Mozilla was complaining about, which we talked about last week, you know, which continually pulls users back to Safari, that'll probably actually have to disappear. If Mozilla is resource constrained, as they presumably are, I would much prefer to have them keep their focus on the desktop, where it matters, and ignore this as a distraction which, after all, only applies to EU territory. So, you know, like why would they bother to go to all the trouble of porting the true Firefox code base over to iOS for a piece of the world when it already looks like everybody has Firefox who wants to.

Okay, now, before I share the big worry that this story prompted in me, I want to share a bit more about this. The day after The Verge's coverage, MacRumors followed up with additional coverage under their headline "Apple Further Explains iOS 17.4's New Default Browser Prompt in the EU." So MacRumors said: "After updating to iOS 17.4, which is currently in beta, iPhone users in the EU will be prompted to choose a default web browser when they first open Safari. In an email today, Apple shared additional details about how this process will work. Apple said iPhone users in the EU will be presented with a list of the 12 most popular web browsers from their country's local App Store at the time, and noted that the options will be shown in random order for every user." You know, that's much as like ballots in the U.S. have their orders randomized in order to remove the bias of just picking the first one if you're in a hurry.

"Apple shared," as an example, "an alphabetical list of the browsers that will currently be shown in every EU country." Okay. So that's a long list, 12 of the most popular browsers. So MacRumors said: "We've elected to highlight browsers that will be shown in France, Germany, Italy, and Spain as examples."

So, okay. You're in France. You open Safari for the first time. And so these are not randomized. These are alphabetical; right? We have the Aloha browser. What? I guess it's better than Sayonara. We have Brave.

Leo: It's hello and goodbye. That's pretty good.

Steve: We've got Brave, Chrome, DuckDuckGo, Ecosia? Okay.

Leo: Oh, yeah, yeah, yeah. That's a browser that is a NetZero browser, so they have a carbon neutral browser. Ecosia, yeah. They also have a search that's carbon neutral, yeah.

Steve: Meaning what? You have to only inhale when you're using it?

Leo: Yeah, something like that. They probably, I mean, you know, maybe they have green network centers, or more likely they buy carbon credits. But anyway, that's their pitch.

Steve: Well, okay. Then we have Edge, we've got Firefox, we have the Onion Browser. I thought that was interesting. I didn't know there was a - yeah.

Leo: Yeah.

Steve: Then Opera, Private Browser Deluxe - ooh.

Leo: Ooh.

Steve: Yeah. Qwant and Safari.

Leo: Yeah.

Steve: Now, okay, that's in France. Germans, oh, they get their own list. Aloha, some are very familiar by now, Aloha, Brave, Chrome, DuckDuckGo, Ecosia - keep inhaling - Edge, Firefox, Ivanti Web@Work, then the Onion Browser, Opera, Safari, and You.com AI Search Assistant made the top 12 over in Germany.

Leo: Yeah.

Steve: Italy looks pretty much the same as Germany. Let's see, Spain, yeah, pretty much the same as the other guys. So anyway, sort of an overwhelming, like, what browser would you like to pretend to be using?

Leo: They've done this before. This is a so-called "browser ballot." And they made Microsoft do this for a while.

Steve: Wow.

Leo: Yeah.

Steve: And were there many when Microsoft did this?

Leo: Yeah, yeah.

Steve: Did you have to, like, you had to scroll the page in order to...

Leo: Yeah. Yeah.

Steve: Wow.

Leo: Is Sleipnir on any of these? Because that was the one, the Scandinavian browser that really got a lot of benefit from...

Steve: Well, there are 23 other countries in the EU that this change applies to. So each country will get its own list.

Leo: Each one will get its own. Wow.

Steve: And, yeah, you may be able to find your favorite, depending on where you are. So MacRumors says: "It's been possible to change an iPhone's default web browser through the Settings app since iOS 14. Apple has now gone a step further and added the default browser prompt in Safari to comply with new regulations under the EU's DMA. In the EU, iOS 17.4 also allows web browsers to use web engines other than Apple's WebKit," which of course is of most interest to us. And again, in March this is going to happen.

Okay. So, okay, this clarifies a few things. If all of these browsers are currently the top 12 most popular in each EU country's regional Apple App Store, then they're all, as we know, currently using simple skins over WebKit since that's all that's been possible until now. That means that users will likely not initially be changing the underlying engine. Just the skin. They'll be prompted to proactively pick a skin. And I imagine that a great many of them, I don't know, given a choice, will opt for Chrome, you know, since that's the browser that dominates the desktop. Then after that, if specific browser vendors see some reason to invest in porting the Chromium or Mozilla engines over to iOS, then that might happen.

Now, the one reason I can see Google investing in a full Chrome port is that they badly need their Privacy Sandbox API to be running everywhere. If tracking - now, that's assuming that Apple doesn't port the Privacy Sandbox API into WebKit, which I hope they do, much in the same way that I hope that Mozilla adds it to Firefox. It's open source. And, you know, Google would like to see this thing become a standard. But if Apple doesn't do that, Google needs their Privacy Sandbox API to be running on iOS devices so that when users choose the Chrome skin, it's more than just skin. And of course that can happen a ways down the road; right?

So in March people in the EU get this choice, and they say, well, what do you know, now I want to go with Chrome. And initially - now, I should also mention that Google's been working on this port for a year, as I mentioned before. Maybe they knew something. So before long it looks like it will be possible to actually have real Chrome and the extended advertising user interest API over on iOS platforms. But again, it's really difficult to imagine that Mozilla would either have the resources or would choose to spend them because this is only happening in the EU. We don't know what the future holds for this.

I mentioned a big worry that this announcement triggered in me. What we see here is Apple capitulating to the demands imposed by a regional legal framework. I suppose they have no choice if they wish to continue operating in the EU. The Verge made it clear that Apple is furious about this, but capitulating they are. And this reminded me of the pending European eIDAS 2.0 legislation, remember, the one which intends to compel the world's web browsers and operating systems to accept, without recourse, any and all root certificates that the EU may choose to require browsers to honor. The EU's DMA is about competition and antitrust. Its aim is to water down Apple's vise grip on its traditional heavy-handed business practices there.

So it's not directly comparable to the eIDAS 2.0 legislation, but I do get a sinking feeling about this. We may be - it may be a good thing that I'm no longer committed as my first and only priority to getting SpinRite finished because there may be a need to keep our browsers clean of EU-enforced certificates if our browsers when we're not in the EU try to impose those on us. We'll see.

Okay, now, Leo. We're going to answer the question, what operating system cannot be infected by ransomware? Under the heading "Dodged a Bullet," we have the news...

Leo: Okay.

Steve: We have the news that the third largest bank in the world, which is China's ICBC, was hit with a ransomware attack which got into and would have compromised their entire network, except for one tiny detail.

Leo: Running Windows 4. I don't know.

Steve: Believe it or not, in this Year of Our Lord 2024...

Leo: 3.1?

Steve: ...the critical currency trading network used by China's ICBC bank was being run by a Novell Netware server.

Leo: There you go. Nice choice. Excellent job.

Steve: As they say, if it's not broken. Anyway, a Novell Netware server was so entirely alien to the ransomware, which had no idea how to infect the server or get up to any other mischief, that nothing happened. Consequently, the bank just shrugged off the attack, cleaned some modern workstations that had succumbed, and got on with their day in this Year of the Dragon.

Leo: Wow, Netware, wow.

Steve: Yeah. Yeah. You know, they probably have text, you know, ASCII text terminals, and they're typing, you know, wow. I guess it wouldn't be ASCII, though, in China, would it. It would be...

Leo: Oh, yeah, that's right. I don't know. That's a good question.

Steve: Okay. So the news that many of our listeners forwarded to me recently was that HP has once again been bricking their printers when those printers are found to contain third-party ink cartridges. Here's what 9to5Mac wrote under the headline "Third-party ink cartridges brick HP printers after 'anti-virus' update." They said: "HP is pushing over-the-air firmware updates to its printers, bricking them if they are using third-party ink cartridges. But don't worry, it's not a money-grab, says the company. It's just trying to protect you from the well-known risk of viruses embedded in ink cartridges." What?

Leo: What? Well-known.

Steve: What? What? "HP has long been known," they write, "for sketchy practices in its attempt to turn ink purchases into a subscription service. If you cancel a subscription, for

example, the company will immediately stop the printer using the ink you've already paid for." In other words, disconnect from the network before you cancel your subscription. Wow. This is just so wrong.

"HP's CEO Enrique Lores somehow managed to keep a straight face," they said, "while explaining to CNBC that the company was only trying to protect users from viruses which might be embedded into aftermarket ink cartridges. He said: 'It can create issues where the printers stop working because the inks have not been designed to be used in our printers, to then create security issues. We have seen that you can embed viruses in the cartridges, and through the cartridge, go to the printer, from the printer to the network.'" And then it takes over the world.

Ars Technica asked several security experts, actual experts, whether this could happen. And they said this is so far out there, it would have to be a nation-state attack on a specific individual. You know, like who somehow got a cartridge sent to them from Russia.

Leo: Your free cartridges are here. Congratulations.

Steve: Why is the label printed in Russian? Well, we don't know. But trust us. Put it in.

Leo: It's going to be good, yeah.

Steve: Perfectly good ink. So one expert said: "Purely from a threat-modeling perspective, I'm skeptical, unless it's a nation-state doing a tailored attack."

Another expert said: "As someone who works for a different inkjet print company, I'd say it's pretty terrible engine design if you could maliciously craft a cartridge to contain a virus." And we're not talking about a liquid virus; right? It's not something that you don't want to inhale. This is like, you know, a computer virus on, like, what? The chip that monitors the ink level?

Anyway, he said: "The amount of information which needs to be stored on the cartridge is fairly small," like a serial number; right? "If the data is not in the format you expect, reject it as invalid." And as a matter of fact, HP's known to be quite good at rejecting such things.

The last expert asked said: "I've seen and done some truly wacky hardware stuff in my life, including hiding data in SPD EEPROMs on memory DIMMs and replacing them with microcontrollers for similar shenanigans. So believe me when I say that this claim is wildly implausible even in a lab setting, let alone in the wild, and let alone at any scale that impacts businesses or individuals rather than selected political actors." So these experts are recognizing that, well, you know, given enough motivation in designing a custom attack, in a lab, somehow getting a specific cartridge into someone's printer, eh, maybe. But not like HP is protecting their customer base and the world because, you know, those cartridges.

Anyway, HP is facing a class action lawsuit - no surprise there - for deploying the bricking code without informing printer buyers of its intention to do so. The lawsuit explains: "This is a class action brought against HP, Inc., for requiring customers who had purchased certain brands of printers to use only HP-branded replacement ink cartridges, rather than purchasing ink replacements from its competitors. HP accomplished this through firmware updates it distributed electronically to all registered owners of the printers,

which effectively disabled the printer if the user installed a replacement ink cartridge that was not HP-branded. In the same time period, HP raised prices on the HP-branded replacement ink cartridges. In effect, HP used the software update to create a monopoly in the aftermarket for replacement cartridges, permitting it to raise prices without fear of being undercut by competitors." So sometimes you get what you deserve.

Leo: That's pretty appalling. It's just appalling.

Steve: In this case, it is just over-to-top bad.

Leo: It's greed.

Steve: Yeah. And actually we will be discussing greed a little bit later today because it seems to be a...

Leo: Comes up a lot, yeah.

Steve: ...recurring topic, unfortunately.

Leo: Yes.

Steve: So we have a leak that's being called MOAB because it's the Mother Of All Breaches, M-O-A-B, totaling an astounding 12TB of data contained within 26 billion (with a B) database records. So if you were wondering whether you would ever have an actual need for that 15TB hard drive you purchased recently, well, yes. There would be three terabytes left over after you transferred the MOAB breach onto that drive.

The super-massive leak, as it's being called, contains data from numerous previous breaches, including data from LinkedIn, Twitter, Weibo, Tencent, and other platforms' user data, making it the largest collection of stolen user data ever discovered. The data includes records from thousands of meticulously compiled and reindexed leaks, breaches, and privately sold databases.

Bob Diachenko, we've mentioned him before. He's the guy who appears to specialize in discovering open and exposed databases online. He was behind this discovery. Although the owner of the database was initially unknown, an outfit known as Leak-Lookup, which is a data breach search engine, said it was the holder of the leaked dataset. The platform posted a message on X, saying the problem behind the leak was a "firewall misconfiguration" - and how - which was fixed.

While the leaked dataset contains mostly information from past data breaches, it almost certainly holds some new data that has never been published before. For example, the Cybernews data leak checker, which relies on data from all major data leaks, contains information from over 2,500 data breaches with 15 billion records. But MOAB contains 26 billion records, that's an additional 11 billion records...

Leo: That's more people than there are in the world.

Steve: Yeah. There's probably some duplicates.

Leo: Duplicates, yeah.

Steve: Yes. That is organized in 3,800 folders, presumably 3,800 individual data breaches occurring through time, with each folder corresponding to a separate data breach. While this doesn't mean that the difference between the two automatically translates to previously unpublished data, billions of new records point to a very high probability of that being the case, that there is some never-before-seen information. Researchers believe that the owner of the MOAB has a vested interest in storing large amounts of data and, therefore, could be a malicious actor, data broker, or some service that works with large amounts of data.

The researchers said: "The dataset is extremely dangerous as threat actors could leverage the aggregated data for a wide range of attacks, including identity theft, sophisticated phishing schemes, targeted cyberattacks, and unauthorized access to personal and sensitive accounts." While the team identified over 26 billion records, duplicates are also likely. However, the leaked data contains far more information than just credentials. Most of the exposed data is sensitive and, therefore, valuable for malicious actors.

A quick run through the data tree reveals an astoundingly large number of records compiled from previous breaches. The largest number of records, 1.4 billion, just one source, 1.4 billion, comes from Tencent QQ, a Chinese instant messaging app. However, there are 504 million from Weibo, 360 million from MySpace, 281 million from Twitter, 258 million from Deezer, 251 million from LinkedIn, 220 million from AdultFriendFinder, 153 million from Adobe, 143 million from Canva, 101 million from VK, 86 million from Daily Motion, 69 million from Dropbox, 41 million from Telegram, and on and on down the list.

In addition to data on individuals, the leak also includes records of various government organizations in the U.S., Brazil, Germany, Philippines, Turkey, and others. If anyone wonders where and how targeted credential stuffing attacks originate, one would need to look no further. The database contains names and addresses, very personal information, password hashes, and in-the-clear email addresses.

So, you know, the people who discovered this are to some degree, I think, understandably hyping it up a bit. But this is not to say that it's not a seriously worrisome collection of potentially potent data. We should also keep in mind, however, that it is a collection of data gathered from all previous data breaches. That means that it is aging and is no longer current, at least much of it is no longer current. You know, no one should ever have their security breached. But anyone who is still using "123456" as their single global password - which, fortunately, is becoming quite difficult to do any longer, thanks to password policies. If you're still using 123456, you should not be surprised if your account were breached. And really, no big database is required, you know, to even use - to even try using 123456 in order to get in.

So there is - we know that there's something to be said for pulling all this together and indexing it and making it rapidly searchable. That's I think the threat, where for example who knows whether the people behind the 23andMe hack took the account data that they could see, you know, pulled from a massively large and indexed database like this to obtain a bunch of hashes which had been reversed and then used those for password spraying. Or, speaking of spraying, the same thing could go for Microsoft. So anyway, just interesting that we're talking about a single one-stop shop for 12TB of data. You know, you ask what it contains, the answer is, well, what doesn't it contain? Because

apparently everything that has ever been breached, that's ever been put out on the dark web or elsewhere has been pulled together.

Federal investigators are warning companies which are under investigation that they may not and must not delete chats, and that they must arrange to preserve conversations that have taken place via business collaboration and ephemeral messaging platforms. In dual coordinated press releases last Friday, the U.S. Department of Justice and the U.S. Federal Trade Commission announced updated language in their preservation letters and specifications, documents they send to companies which fall under federal investigation. The new language updates evidence preservation procedures to cover modern tech stacks such as Slack, Microsoft Teams, and Signal. I guess without that they figured that, you know, unless they were very clear what they meant by "thou shall not delete," companies could say, oh, well, we didn't know you meant that.

Companies that receive subpoenas or other legal notifications must take steps to preserve chat logs and disappearing IM messages, and any who do not will be subject to obstruction of justice charges. The problem, of course, is that being charged with obstruction of justice might be better than revealing what they were deliberately, you know, talking about, and then chose to delete. The Deputy Assistant Attorney General of the Justice Department's Antitrust Division said: "These updates to our legal process will ensure that neither opposing counsel nor their clients can feign ignorance when their clients or companies choose to conduct business through ephemeral messaging."

And this updated guidance comes as the DOJ face difficulties pursuing its antitrust lawsuits against Google and Amazon. So, you know, there were some targets they had in mind. February last year, the DOJ accused Google of lying when it claimed it auto-suspended its chat auto-deletion feature. In addition, the DOJ claimed that for a period of four years, Google trained employees to delete internal chats and move conversations to off-the-record platforms because it anticipated facing antitrust litigation in the near future. Later, in November last year, the FTC accused Amazon of deleting more than two years' worth of internal Signal employee chats after the agency started a multi-state antitrust lawsuit.

I have a representative snippet of the DOJ's evidence-hiding complaint in their antitrust case against Google. It's just a few lines. It reads, and so this is from the federal complaint, says: "The newly produced chats reveal a company-wide culture" - speaking of Google - "of concealment coming from the very top, including CEO Sundar Pichai, who is a custodian in this case. In one chat, Mr. Pichai began discussing a substantive topic, and then immediately wrote 'also can we change the setting of this group to history off.' Then nine seconds later, Mr. Pichai apparently attempted unsuccessfully to delete this incriminating message."

And there's a reference to a piece of evidence where there's like a serial number, but it does have the words GOOG-PLAY. So one wonders what they're talking about there. The complaint continues: "When asked under oath about the attempted deletion of the message, Mr. Pichai had no explanation, testifying 'I definitely don't know' and 'I don't recall.'

"Like Mr. Pichai, other key Google employees, including those in leadership roles, routinely opted to move from history-on rooms to history-off chats to hold sensitive conversations, even though they knew they were subject to legal holds." Meaning after they had been told they need to retain all records. This thing says: "Indeed, they did so even when discussing topics they knew were covered by the litigation holds in order to avoid leaving a record that could be produced in litigation. As the examples below make clear, Google destroyed innumerable chats with the intent to deprive Plaintiffs" - meaning the federal government - "and other litigants of the use of these documents in litigation."

Okay. So the federal government is making it very clear that digital recordings of private conversations may not be deleted from the moment of notification of pending litigation. If executives wish to hold private off-the-record conversations, they're going to need to do it the old-fashioned way, face to face in a private setting with no one recording. And, you know, Leo, it sort of begs the question, too, what if you use a system where one of its features is not to record long-term history?

Leo: I guess you're not allowed to, is the point.

Steve: Right.

Leo: But are you allowed to then have a private - you always said take your clothes off, go in the middle of a field, and have a [crosstalk]...

Steve: And throw a big thick comforter, a blanket over yourself, yeah.

Leo: I've been watching an old movie called "The Yards," with Joaquin Phoenix. And he was meeting with - he's kind of a mobster meeting with a city councilman, or no, I guess it was a borough president. And the borough president made him take off all his clothes, and then he took off all his clothes, before they had a conversation, to make sure that they weren't a wire. So there's a longstanding precedent for this, I guess. I don't - it's pretty funny.

Steve: It's the only way to be sure.

Leo: But I'm guessing that they can't say you may not have any conversations in person. The court can't say that.

Steve: Right, right.

Leo: But they can say you may not use any technological means that don't leave a paper trail. Right? I guess they can.

Steve: I think that must be what they can say, yeah.

Leo: Yeah, yeah. Isn't that interesting, yeah.

Steve: Wow. Yeah. Okay. So here's one that'll really, well, really ruined Mercedes's day. Mercedes-Benz accidentally exposed, putting it mildly would be a trove of internal data.

Leo: Not a MOAB, but a trove.

Steve: It's not a MOAB, no. It's a trove. By leaving a private key online that gave unrestricted access to the company's source code.

Leo: Yikes.

Steve: And that key was there, exposed, for more than 90 days, almost 120 days, before it was discovered and responsibly reported by a cofounder and the Chief Technology Officer for the London-based group RedHunt, as in "Hunt for Red October," RedHunt Labs.

Leo: Let me guess. They pushed it into a git and published it.

Steve: Uh-huh.

Leo: Did they really? Oh, that's hysterical.

Steve: Yep. What RedHunt discovered during a...

Leo: [Crosstalk] by accident, I've got to say.

Steve: Yes, yes. What they discovered during a routine Internet data scan earlier this month, earlier January, was a Mercedes employee's authentication token sitting in a public GitHub repository. This token served as an alternative to using a password for authenticating to GitHub. As such, it would and did grant anyone full access to Mercedes's GitHub Enterprise Server, which would in turn allow the download of the company's entire collection of private source code repositories.

RedHunt said that the GitHub token gave "unrestricted" and "unmonitored" access to the entire source code hosted at the internal GitHub Enterprise Server. The repositories - now, here's where it goes from bad to worse. The repositories include a large amount, shy of a MOAB, I agree, but still, of intellectual property. Get this: connection strings, cloud access keys, blueprints, design documents, single sign-on passwords, API keys, and other critical internal information. RedHunt provided evidence that the exposed repositories contained Microsoft Azure and Amazon Web Services keys, a SQL database, and Mercedes source code. It's not known if any customer data was contained within the repositories.

Leo: This is why we're glad we drive BMWs.

Steve: Oh.

Leo: Wow.

Steve: Wow.

Leo: But it's so easy to do. You go git add, and then you push it, and it's all there.

Steve: Yes. Last Monday TechCrunch, serving as a middleman for RedHunt, disclosed the security issue to Mercedes. On Wednesday, a Mercedes spokesperson confirmed that the company "revoked the respective API token and removed the public repository immediately." They said also: "We can confirm that internal source code was published on a public GitHub repository by human error. The security of our organization, products, and services is one of our top priorities. We will continue to analyze this case according to our normal processes. Depending on this, we'll implement remedial measures."

Now, since the exposed key was first published last September, it sat there through the balance of September, all of October, all of November, all of December, and most of January of this year. What's not known is whether anyone besides RedHunt may have discovered and taken advantage of the exposed key. And, you know, weren't forthcoming as RedHunt immediately was. Mercedes declined to say whether it is aware of any third-party access to the exposed data or whether the company has the technical ability, such as through access logs, to determine if there was any improper access to its data repositories. The spokesperson cited unspecified security reasons. Uh-huh.

We've previously, of course, and to your point, Leo, covered that GitHub has begun proactively scanning repositories for these sorts of inadvertent disclosures when they recognize them. But, you know, they don't know everybody's format of anything that they would consider sensitive. So they're doing the best job they can.

Leo: It's so easy to do.

Steve: Yes, it is.

Leo: I guess the question is why are companies using GitHub?

Steve: Yes.

Leo: Instead of their own Git servers.

Steve: Yes. You know, our software and intellectual property management systems have become so complex and interdependent that they have also become brittle to these sorts of human errors. And I don't see that changing; you know? We're moving in this direction. And they provide a great deal of power and flexibility and leverage. But as you said, boy, if you make a mistake, it also amplifies the mistake just as much as it amplifies the power that it provides when it's all working correctly. Wow.

The good news is fewer ransoms are being paid. The number of ransomware victims who opted to pay ransoms fell to an all-time low by the end of last year. The cybersecurity firm Coveware - we've talked about them before, they track these things - estimates that only 29% of victims paid ransoms during the fourth quarter of 2023. That's down from 85% who were choosing to pay back when we began talking about this and when they started tracking it, which was the first quarter of 2019, you know, so four years ago when this all began to really ramp up. So 85% initially. Now we're down to 29%. So that's great.

Coveware attributes the fall to improved data backup and recovery strategies in corporate environments, and companies getting smarter about not trusting empty promises made by ransomware groups. So it's like, yeah, I mean, here we are years downstream. Not only does the CIO absolutely know about this, but there's no way that the CEOs and COOs in these organizations are not all aware of the threat posed and somewhere along the way said, you know, to the CIO, you know, what resources do you need that you don't have? If we get hit by ransomware, we don't want to be taken down. So the world has changed. That's good.

Okay. I have some feedback from our listeners. Conradical, he tweeted: "Steve, please take a deeper dive into the technology behind verified camera images. My gut reaction," he writes, "is you've overlooked something because public key cryptography should allow the images to be verifiable and unmodifiable."

Okay. There are without a doubt many amazing things that public key crypto can do. And I am deeply enamored of them. That's what I built the whole SQRL system around. It was all public key crypto based. But in the case of the verified camera images, you must ask yourself what could a camera contain that cannot be copied by someone who gets their clutches on such a camera? I contend that anything a camera can know, someone can find a way to pry out of that camera to duplicate whatever it knows; and, in doing so, duplicate its ability to make a strong assertion of an image's origin. In other words, this entire system depends upon the camera, which is out in public, being able to keep a secret. And everything we know tells us that's almost certainly not possible if someone is sufficiently motivated.

The most common application we have today of public key crypto is the dynamic creation of secure connections to remote web servers, where those servers are asserting their identity. Only one thing allows that system to work, which is that those servers are not accessible to others. If they were, the secrets they're protecting could be stolen, and others could impersonate them. That's the difference in the security model of the camera versus a remote server. It's the remoteness of the server that allows it to protect its secrets, the fact that it can only be accessed through a carefully managed TCP connection. The infamous Heartbleed vulnerability demonstrated what would happen if server secrets could be accessed through a side channel. The server's secrets would be compromised.

So it's not that the public key crypto doesn't still require secrets. It does. It's just that only one side of the transaction needs to be able to keep something secret. Unfortunately, when a camera is signing the pictures it takes, it's the private key that the camera is using to perform the signing that needs to be kept secret. Building a state-of-the-art hardware security module into the camera, which is I'm sure what they've done, will likely make it as difficult as possible to extract the HSM's key. The unanswered question is, will it be difficult enough? And only time will tell.

Leo: So this is one of the first cameras to do this. This is the Leica M11-P, which is supporting the content authenticity initiative.

Steve: Yup.

Leo: Launched by Adobe, Twitter, and The New York Times. And it's exactly as you say. And by the way, it may be that it is difficult to do that, and that's why it's such - this is a $10,000 camera - why it's, you know, starting here is that they probably did build in a Secure Enclave. I mean, they must have; right?

Steve: Oh, yeah. It's got, I mean, I'm sure they went to every length they could to keep anybody from ever extracting its private key.

Leo: And the thing of course that happens is you can strip it out. You can get a JPG of the image that doesn't contain it. What's interesting about this, though, I think it's kind of cool, is it's signed when you take the picture. And it's a form of metadata, but it's not in the Exif. It's signed. And then when you modify anything, that's also recorded when you modify it in Lightroom.

Steve: Yep. Yep.

Leo: So there's a chain of custody, which I think is very interesting.

Steve: Yeah. The Adobe software maintains a complete audit trail of any changes that are made to the image. So you always have those, and you're able to rewind it all the way back.

Leo: So I think the main point I guess is that you'll certainly see images that don't have these credentials. But if you see an image that has the credentials, you're supposedly going to be able to say, oh, I see who took this, and I see how it was modified. Is this an actual recording of something that really happened?

Steve: Right.

Leo: And, yeah, I guess the only way that could be forged is they'd have to get the camera and somehow get the credentials out of the camera.

Steve: Well, yeah. And that's just it is that everybody, you know, these cameras will be floating around. Now, maybe, and it's probably the case, that every camera has a unique private key.

Leo: I think that's the case, yes.

Steve: So it would be, you know, this camera signed this, you know, alleging that it signed this image. So what they could do is make any extraction of the key a destructive process. You know, so the point being that the only way to really pull off a spoof would be to arrange to extract the key and leave the camera still intact, and its owner not knowing that anything had been done. But in general the problem is we're being, you know, the idea is that it's being put forth as a means of detecting any spoofing of the image. And like so the images signed with the CAI system are trusted at a higher level.

Leo: Yeah.

Steve: And, you know, that will be true until I'm sharing a piece of research a couple years from now from the guys from the University...

Leo: DVD John. Yeah, yeah.

Steve: The University of the Negev, you know, who were able to hear a conversation across the quad by bouncing a laser off of a plant leaf. Those guys are going to say, oh, yeah, you know, unfortunately there's a side channel attack that's available. We were able to aim our sniffer at it while we took a picture, and now we know what the key is.

Leo: Steve, you do, I mean, you trust the - like a credential from an iPhone; right? I mean, that's a similar thing with a Secure Enclave and so forth. That hasn't been cracked; right?

Steve: Yeah.

Leo: So I think it's similar to this. Anyway, you've given a good excuse to buy this $10,000 camera and report back to you. So I will. I'll just pick one up, and...

Steve: Leo, you need one. It is crucial...

Leo: It's now tax deductible, thanks to you.

Steve: It's crucial that the pictures you take of Lisa's birthday party be authenticated.

Leo: Authentic. Only...

Steve: Absolutely.

Leo: Yes. And unmodified. Okay. So, yeah.

Steve: Okay. Anyway...

Leo: Now I have to get - I've got your approval. I just have to get Lisa's now, and we're set.

Steve: That's right. Good luck with that. So...

Leo: It's for work, honey.

Steve: Jg1212G, he tweeted: "Hi, Steve. I was just listening to Security Now! and got hooked into the $15 per week flashlight story. I had to look into it. I found it on the Play Store and followed the link to their website, simplemobiletools.com." He said: "I thought, very strange, the site says open source and free. So I clicked on the GitHub link at the

bottom. Sure enough, it's open source. So I looked at the developer's page on GitHub." And that's github.com/tibbi. He says: "Wow. His graph shows he was extremely active up until the end of October 2023, then completely stopped." He says: "That's so strange. I would really like to know what happened to him. If you hear any news, please let us know. I love a good mystery. Thanks, Jason."

Jason, ask and you shall receive. Our listener megascrapper brings an end to the mystery. Megascrapper, tweeting from @megascrapper, says: "Hi, Steve. I'd like to follow up on last week's listener feedback about the absurd subscription prices for a flashlight app. I was made aware of the entire Simple Mobile Tools suite, which includes Simple Flashlight, after watching a video by Brodie Robertson." And we have a link in the show notes for anyone who's curious.

"Unfortunately," he writes, "what happened with Simple Flashlight was exactly what you presumed in your reply to that listener last week. With very little notice, the owner/primary maintainer of the app sold the entire suite to an Israeli publisher, ZipoApps" - good old Zipo - "which is notorious for the practice of acquiring existing apps and slapping on an outrageously expensive subscription plan.

"But not all hope is lost," he says. "The entire suite is open source, GPLv3 licensed. And one of the maintainers already forked it under a project called Fossify, including Simple Flashlight. It seems to be still in early development, and I can't find the app on Google Play Store, but keep an eye out when it gets released. Thank you very much for your work. Look forward to 999 and beyond."

So, thank you, megascrapper, for your follow-up on this and for the confirmation that this is the sort of thing that happens with highly popular apps in the Google Play Store. The description for the video that he linked says: "I was a fan of the Simple Mobile Tools suite for a really long time, and then out of nowhere the developer Tibor Kaputa just sold the entire project and ran away with the bag."

Leo: Would you say it's "kaputa"?

Steve: Kaputa, that's very good, Leo. Yes. Wow. So anyway, we would agree it's certainly Tibor's right to do whatever he wanted to with his own intellectual property. It's clear that since the entire project is open source, it was his project's developer keys that was of actual value because they allowed its purchaser to take over the official popular app and then upgrade it into the existing channel of owners.

And, boy, this caught everybody's attention. Jon Dagle, he tweeted: "Hey, Steve. In response to the Flashlight app story. First, to access the flashlight brightness, swipe down from the top right to get Control Center." Okay, so now he's talking about iOS rather than Android. He says: "Long-press the flashlight icon. Solved." Okay. I tried it on my iPhone, and I was amazed.

Leo: I tried to tell you this last week.

Steve: You did tell me. But I didn't know you meant iPhone. It has, Leo, it has four levels of brightness. I never knew.

Leo: Yeah.

Steve: Now, this is super useful to me since the flashlight defaults...

Leo: Really bright, yes.

Steve: ...to a setting that should be labeled "Visible from Orbit." All I want to do is I want to read the menu in a darkened restaurant. I'm not trying to signal aliens for pickup. So I immediately set it to its lowest level, which will be much more appropriate in the future, and it won't blind my fellow diners if I inadvertently pass its laser beacon across their vision.

Okay. Before I get to Jon's second point, I just want to mention something that's quite annoying.

Leo: Yes.

Steve: I have this dull sense that there is vastly more available from today's iPhone than I'm aware of. You know? But how would I ever discover this on my own? I guess I just have to sit around and press on everything to see if anything happens, which is annoying. The original concept of the graphical user interface was that it was discoverable. You know, that's what, you know, what was so cool about it was that you had nested drop-down menus running along the top of the screen. Unlike the text command interface that preceded them, you could sit down and run the mouse around the screen and find everything that you might need. Today, it's easy to do the basic things with a phone, but it's annoying to imagine just how much more remains hidden behind the need, you know, and here's where I said at the top of the show, you need to click your heels together three times in order to discover something. How would you know? Just as Dorothy had no idea how to get home to Kansas.

Jon's second comment: "I recently ran across another long-time trusted app that was sold. It's the super excellent Network Toolbox on iOS. I think it's been mentioned on Security Now! in the past. It has a host of powerful networking tools. But the longtime developer sold the app sometime late in 2023. When I first opened the app after recently resetting all settings, I got the 'Network Toolbox wants to track you across websites' alert." He said: "The sale/transfer was all silent as far as I'm aware. Considering the app has a lot of sensitive functions, the trustworthiness of the developer is rather important. So beware."

And all of these stories got me to thinking that perhaps Google and Apple really ought to consider adding proactive notification to apps when their ownership changes hands.

Leo: Oh, I like that, yes.

Steve: Yeah. I've never participated in such a transfer, so I don't have any clear sense for whether a developer might simply turn over their entire online identity to a third-party purchaser, or whether there's some more formal and controlled process for doing so. But if it's knowable to Google or Apple, it would seem useful to add a bit of friction and visibility to this otherwise very slippery and transparent process. This is, you know, there's a lot of trust that's built up over time. So if that publisher changes, it seems to me that those being asked to trust someone new should know.

Leo: Yeah. And this is a big issue, as you mentioned, in Google Chrome extensions, where this seems to happen a lot.

Steve: Right.

Leo: So, yeah. I think, I mean, I think with the Apple thing you kind of have - you have to have a developer count, so they should know.

Steve: Right.

Leo: That should be clear, I think.

Steve: If it moves to a different developer.

Leo: Yeah, yeah, yeah.

Steve: Right. Brian Doyle tweeted, well, he brings us another example of Mozilla which they can add to their growing list of "good luck with that" grievances against the tactics being employed by those who wish to use their own platforms to their competitive advantage. He tweeted: "Hi, Steve. I came across this message while looking for a way to save full web pages into OneNote, and had to laugh at Microsoft implying that Firefox is not a 'modern' browser. Thought you might enjoy. Here is the original site." So, and I went there, too, and grabbed a screenshot. It's onenote.com/clipper, C-L-I-P-P-E-R. And if you go there in Firefox, up comes a state-of-the-art-looking website, and it says right up at the top of the screen: "OneNote Web Clipper is no longer supported on Firefox browser and works best using a modern browser like Microsoft Edge."

Leo: Oh, boy. Oh, boy.

Steve: Yeah, not that stinky old Firefox.

Leo: Oh, yes, so out of date, Firefox. Oh, boy.

Steve: Yeah, uh-huh. That's right.

Leo: Come on. Geez.

Steve: So, wow.

Leo: Sad.

Steve: Just, you know, cheap shot. It really is.

Leo: It is a cheap shot.

Steve: Yeah. Someone whose Twitter handle I didn't really get until I said it phonetically is ShipRkt, clearly Shipwrecked. He said: "Hello, Steve. I hope you don't mind me sending you a message." No, that's Twitter. "Could you discuss on a future Security Now! episode why Credit Karma is storing over 1GB of data on my iPhone?"

Leo: Yikes.

Steve: "What on earth" - yeah. "What on earth uses that much data for a credit app? Thanks for your time." Okay. So as I said, I don't mind receiving messages, which is why I go in search of them every week for the podcast. But neither do I have any idea why the Credit Karma app might be storing over 1GB of data on anyone's iPhone. One thought I had was to wonder whether this 1GB might include the app itself in that total. One of the sad trends we see is applications becoming increasingly and, in fact, obscenely bloated. They evidence no respect whatsoever for the user of their apps. I'm sure that very few consumers are even aware of this, which is why there's little cost associated, you know, reputational cost associated with being so careless with the consumption of other people's storage.

Anyway. Curmudgeon rant off. I would love to put this question to our listeners, who we know are quite resourceful. I poked around a bit, but I didn't find anything obvious about Credit Karma's iOS app resource consumption. So if anything finds anything, I'd be happy to share it.

Leo: I would just, I mean, you can just delete it. When you delete something from the iPhone it says "Do you want to delete the associated data?" You say yes. Install it again, login again, and I don't think you've lost anything from Credit Karma. It may be that it's recording every transaction you make, and that's added up over time. But I think starting over probably wouldn't hurt. They should be storing that on their servers, not yours, personally.

Steve: That's a very good point. Mark Guy, and it looks like from his Twitter handle, it's @SDTwitGuy, he appears to be a fan of the network.

Leo: Or Twitter. A lot of people call Twitter "Twit," which bugs the hell out of me.

Steve: Oh, that's a very good point. But he is also listening to Security Now!. So he said: "I heard your comment about staying on Windows 7 on the 1/23/24 podcast last week." He said: "My main system is Windows 7 Ultimate. They'll have to pry it out of my cold, dead hands." He said: "It's stable. It runs perfectly. I subscribe to 0patch and still get updates for MS Security Essentials, plus I use Malwarebytes Premium. Never had any problems. Plus I know where everything is. I bought a used Windows 10 laptop, and I can barely find anything. I also am an avid fan of Windows Media Center. Nothing else comes close to its functionality. It's how I watch and record TV, so I will never update my system."

Leo: Yeah, because it will remove it if you do. So he's probably right to stick with 7, yeah.

Steve: Yeah. He said: "I'm also a huge sci-fi fan, and I LOVE [in all caps] that you and Leo talk about your fave sci-fi authors, books, and series. Thank you."

Okay. So I wanted to mention two things from Mark. I know that Mark and I are far from alone among the listeners of this podcast. Just like with that Novell Netware server, it's working, so don't mess with it. And yes, at some point I'll rebuild my machine around Windows 10. Since I'm an MSDN developer I could still actually register a new machine as Windows 7. But I'm not totally insane. I'm typing this into a Windows 7 workstation mostly because moving to Windows 10 would take a non-zero amount of time. And like Mark and many of our listeners, why bother when this 64-bit edition of Windows 7 is working fine for me.

The second thing I wanted to mention follows from Mark's comment where he said "I'm also a huge sci-fi fan" and love it that we talk about this stuff. I've been intending to mention that after investing in about six of those Aeon 14 novels, you know, the ones invariably featuring voluptuous, heavily armed female commandos on their covers, despite the fact that another hundred or so of those novels remain, I had finally reached my limit.

Following a number of recommendations, I gave the "Expeditionary Force" novels a try, but they just didn't grab me. They're written in a first-person narrative style that just didn't work for me, and I kept waiting for something to happen. Now, my trouble might be that they're a bit too realistic, and not that much actually happens in life. Once you've read much of Peter Hamilton's work, you're somewhat cut loose from the need for an excess of reality. Which, you know, never kept Peter from telling a story.

But in the meantime, Ryk Brown, spelled R-Y-K Brown, the prodigious author of the Frontiers Saga series, had dropped a few more books in his third of five planned 15-book story arcs. Since we're up to book 10 in arc #3, we've passed the halfway point. You know, 75 books is what he's got planned. I've turned a number of my very close friends and family members onto this series, and I have been unable to shake them loose. They want nothing to do with anything else. They just want more Ryk Brown. As we know, I've wandered around while I've been waiting for more. I happily consumed the entire Silver Ships series following another recommendation from a listener, and of course some of those Aeon 14 series.

Anyway, I'm bringing all this up because Ryk Brown's writing style, his deep characterization, his perfect management of a large and growing number of very different and distinct characters, and the fact that you never need to wait long for some action, continues after 40 books, which is how many I've consumed, to be absolutely enjoyable and gratifying. All the books are available under Amazon's Kindle Unlimited plan and as audio books from Amazon. Through the 19 years of this podcast we've shared our discoveries of many terrific books. For sheer solid entertainment value, I think this series deserves everyone's attention. So I just wanted to be sure that, again, it was on everyone's radar.

A couple last things. Dizzle Von Dazzle tweeted: "Quick question." Yes, Mr. Dazzle. Dizzle von Dazzle. He said: "As you are an avid user of Windows 7, how do I continue to use websites that use HSTS?" You know, that's HTTP Secure Transfer Security. STS. I've forgotten what it's - I know what it is. I forgot the abbreviation.

Leo: Yeah, something like that, yeah.

Steve: HSTS.

Leo: Yeah.

Steve: Secure, uh, anyway, someone will tell us. Anyway, he said: "It's a new install on an old-ish Lenovo IdeaPad All in One." So he recently installed Windows 7. He says: "Is there a way to update the SSL libraries, as none of the update managers for different music production applications I own seem to work either." He finishes: "Keep up the amazing work on SpinRite, and here's to episode 999."

Okay. I'm having no trouble with Windows 7 and HSTS sites, such as GRC, which was one of the earliest to adopt HSTS, even though I've forgotten what it stands for, and its permanent registration in Chrome, which GRC also was an early adopter of. Under my Win7 setup notes, I have a subdirectory named "Before registering or installing Win7 updates," and that subdirectory contains three specific Microsoft updates. There's an SHA-256 update, a Servicing Stack update, and an update which is KB3102810.

From my notes, it appears that you should find those three individual standalone updates and install them in that order. Then you can successfully bring Windows 7 current, and all should be well. So the fact is Windows 7, which was first published in 2008, yes, it's showing its age. It didn't support signing things, things being signed with SHA-256. It only knew SHA-1. And so that would explain why he's unable to update his other, what was it, different music production applications. Their updating systems are probably using SHA-256, which Windows 7 does not support out of the box. You need to install the SHA-256 update for Windows 7. Then it will probably work. And are we finally - almost.

Leo: A couple more.

Steve: Two left. A listener who asked to remain anonymous said: "Steve, my company is switching to Bitwarden from LastPass as a result of me raising the issue a year ago, which is a result of your discussion on the podcast. Please keep this anonymous if you mention it on the SN podcast. My question is, can I get a readout from you on the advisability of adding TOTP" - you know, Time-based One-Time Password - "codes and secrets into Bitwarden so that it can fill in the field on sites you're logging into?" He says: "Personally, it gives me a 'Gibsonian response' and feels like all your eggs are in one basket if you do that. What do you think? Regards, long-time listener, et cetera."

Okay. So we've mentioned before about this, but it's worth just covering again. And I know that, Leo, you concur since I've heard you say the same thing on other podcasts. But I've also had some time to think about this and perhaps to mellow about it a bit. I understand the convenience. But it's also true that it represents a classic tradeoff between convenience and maximum security. My honest feeling is that the actual risk of having all the eggs in one basket is likely less significant than the benefit that comes from ease of use. That is, the risk is less of a concern than the benefit it provides in terms of ease of use. So if, for example, it was ever a matter of not registering and using a Time-based One-Time Password due to the inconvenience of needing to use a second authentication device, which would be more secure - and, for example, last week Paul Thurrott was explaining that his wife has absolutely zero interest in anything that gets in her way then yes, it would be better to have Bitwarden able to automatically fill-in the one-time password field than not to use time-based multi-factor authentication at all.

I have no problem keeping my one-time password tokens in my iPhone and in manually transcribing them. I don't have to do it that often, and I really appreciate the real, the true sense of security I get from that. But that's just me. So better to use any one-time password than none, even if it's being automatically filled in by the browser. And if given a choice, it's better not to have the browser filling it in, even though the actual danger, I think, is realistically small. You know, Bitwarden wouldn't have done it if it was really a big problem. They did it because it's like, okay, you know, you still get all these benefits from a one-time password. Why burden the user if we can do it for them?

Leo: If there were some way they could like separate, like keep the TOTP secrets in a vault somewhere on a different country from where the, you know, in case there's a breach, that's the fear; right? In the LastPass breach, if somebody had gotten both LastPass password vault and the TOTP secrets, which many people did use in LastPass, then you would not be quite as secure. So if there could be, you know, you can with Bitwarden, in a personal account, anyway, store your vault yourself. Maybe if you just stored - I wonder if you could separate the TOTP database out. I'll have to look into that. But you're right. You know what?

Steve: Exactly.

Leo: I'm sure it's fine. Use it, you know what, use Scrypt, or I guess you have to use Argon2 as your PBKDF2. Use a really long, you know, miserable master password, or better yet, passkeys. I've been using passkeys now with Bitwarden for passwordless login. And while it doesn't work everywhere, it is very convenient where I can use it. You know, we're putting passkeys in Bitwarden, might as well put your TOTP secrets in there, too; right?

Steve: Yup.

Leo: Yeah. I agree with you. Convenience.

Steve: Okay. Yeah, again, it's better than not using them at all. Way stronger than just a username and password. So, you know, if that's what it takes, do it.

Leo: Depends on your threat model, you know. If he's working for the NSA, well, then, you know, you should do something else.

Steve: Yeah, yeah. George Palfi, that's @PalfiGeorge, he said: "Steve. I'm a devoted listener and longtime SpinRite owner, though I wish it worked on Macs. I gave up Windows completely years ago."

Okay, George. The good news is I made some changes a few months ago to allow SpinRite 6.1 to run on Macs where it can. "Where it can" means Intel Macs, where it's possible to boot from a USB or a CD, typically through Boot Camp. The previous trouble with SpinRite on Macs had been with the keyboard, since SpinRite was accessing the keyboard hardware rather than using the BIOS. I changed that so that SpinRite could work with some less PC-compatible Dell machines, and we got Mac compatibility in the bargain.

Leo: Nice.

Steve: A number of testers have confirmed their ability to now run SpinRite on Macs.

Leo: You were using the A line for the BIOS on the keyboard?

Steve: It's, yeah, they are, well, no. Even PCs that use USB keyboards do take the time to stuff the keyboard data into low RAM.

Leo: Wow. Wow. Wow.

Steve: Which is what the hardware does. Mac doesn't do it. So the Mac's PC emulation is slightly less compatible than all the others. And I am finally able to announce...

Leo: I need drum rolls, please.

Steve: Yes, that after more than three years of work, I am completely satisfied that SpinRite 6.1 is as good as it can be, and that it is finally ready for release. There is nothing left I'm aware of that could be done to further improve SpinRite's functions. I could keep fussing with it forever, adding this or that convenience feature around the edges, but it's already received a large collection of new convenience features, and it is by far the best SpinRite that's ever been created.

Leo: Woohoo!

Steve: It's been proven to work in every environment that it's been placed in by more than 818 testers who've registered with our GitLab instance and who have obtained it through my release announcements in GRC's web forums. It's finally done.

Leo: Hallelujah. Wow.

Steve: Officially, its code still calls itself Release Candidate 6. And it makes sense to let it rest for a bit before it's moved to Final Release 1 since I would prefer not to have to be tweaking the code after it's been released. And there's really no hurry. While the paint is still wet and drying, I'll be working on SpinRite's documentation, which will all be browsable and explorable online. Since many people prefer to click on a video than to read text, I'll also create some video walkthroughs as I did for ReadSpeed, so that someone can get a feel for what SpinRite looks like while it's running.

Leo: Nice.

Steve: Next, once the documentation is finished, I'll be bringing GRC's long-awaited email facility online to get our promised incoming email bag setup to receive incoming mail from this podcast's listeners. So many people have written that they had to painfully

login to Twitter just to get a note to me. Anyway, I get it. That'll finally be changing. And I'll create a weekly mailing list for this podcast so that those who would like to receive a weekly summary and link to the show notes will be able to get that, as well. I'm sure I will continue posting on Twitter. I'm not yet sure whether I'll continue monitoring incoming tweets and DMs there. I'll just play that by ear. I'd very much like to consolidate the channels that I need to follow, and email is the most universal medium which we all share. I've had so much positive feedback from people saying, yes, yes, yes, please, just let me use email.

And once all of that is in place, I'll finally begin the process of notifying all 20 years' worth of SpinRite's past purchasers. Since I imagine many of those 20-year-old email addresses are no longer valid, my plan is to send out the announcements off the free availability of 6.1, starting from the most recent and gradually heading toward the least recent so that, you know, so that I'm not seeing...

Leo: What's the oldest? What's the oldest account?

Steve: ...100% bounces of everything, yeah.

Leo: Who's the longest owner of SpinRite, after you?

Steve: Well, I mean, this is only - my online database only goes back to 2003. And it's funny, too, because Lorrie and I had Sue, my employee for the past 40 years, over to dinner. And we were talking about this because once 6.1 is available, we're not going to continue allowing people to upgrade from earlier versions of SpinRite because, you know, it's been 21 years.

Leo: Yeah, I think that's fair.

Steve: Since, you know, I mean, it's like it's just ridiculous. It's like, come on, guys. So she was excited because that meant she no longer needed to run FoxPro in a DOS box. Our original database, where from every single person who ever bought any copy of SpinRite is in a FoxPro database which we call Dino because, yes, it is a dinosaur.

Leo: It's a dinosaur. Well, you know, dinner parties at the Gibsons', they're really - conversation is fascinating. Actually there's quite a few people listening right now who would love to have been at that party. Well, that's great. Congratulations, Steve. That's really good news. That's great. So look for your email, if you still have that account from all those years ago.

Steve: Well, and the reason I'm bringing it up this way, I did some research in the last week about emailing. And the world has become such a sewer with spam that GRC cannot just suddenly start sending out email in great volume because we don't have a reputation. So the beauty is, by using the Security Now! email and having people sign up, and I'll send them a confirmation, and then that way the world will start seeing GRC doing mail which is valid and it's being accepted by people, instead of like, what the hell is this? And that'll allow us to establish a reputation. And it turns out reputation matters in the same way that it does with the use of a digital certificate to sign software. So, and

I'm not in a big hurry. I don't have to notify everybody in one day. I'll just let it kind of dribble out slowly over time so that, again, we're not setting off any alarms in the anti-spam centers of all ISPs in the world.

Leo: Actually, I should have asked you about this because starting February 1st, Gmail is going to require that all messages are authenticated with DMARC and SPF and DKIM.

Steve: We already have been for quite a while. We're SPF, DKIM, and DMARC.

Leo: Yeah.

Steve: In fact, there is an amazing site, you have to go look at it, it's - I think it's called LearnDMARC. Let me see if I can bring it up.

Leo: Not the easiest thing in the world, I can tell you right now.

Steve: It is so overdesigned, I just - oh, there it is. Yeah, LearnDMARC.com. This guy, whoever he is, is a man after my own heart. He so overdesigned this thing. It is just - it is a gorgeous site. So he gives you a one-time use email address to which you send any piece of email. That allows him to pick it up and then check all of your security settings - SPF, DKIM and DMARC - which basically uses SPF and DKIM, and verifies all the proper security settings of whoever it is who's sending out email on your behalf. And I passed all of the tests with flying colors. But it was just a beautiful experience. He's got stuff floating around the screen to fill in the fields. It was just - it was delightful. So Learn...

Leo: [Crosstalk] send them some email.

Steve: Yeah.

Leo: I'm pretty sure our - I use of course Fastmail, our sponsor. And I'm pretty sure that Fastmail is doing that all from [crosstalk].

Steve: I didn't fill out anything. I didn't fill out anything. And he said what you just got, he says, "Waiting for incoming mail." And after waiting a while, you don't have to write us an entire love letter.

Leo: Wait a minute. Did I not send it?

Steve: This guy is great.

Leo: You call it a beautiful site. It is, let's say, text heavy.

Steve: Just wait. Just wait.

Leo: Oh, yeah, look. It's getting pretty quiet here. Waiting. Oh, so after you send it, it gets pretty; huh?

Steve: There it is.

Leo: Oh, it came, yeah. Wake up, Neo. Hi, there. LearnDMARC. My name is Deo. Oh, this is so cute.

Steve: Oh, just wait. It gets better. Look at that.

Leo: Oh, you're right. This is nice.

Steve: It's so well done.

Leo: Oh, this is nice. Oh, my gosh. You're right. I take it back, Steve. This is a beautiful site. Holy cow. Let me shrink it down so you can see the whole thing. Press any key. Okay. Is used to look up the domain's SPF policy, running SPF, yes, Laporte.email has an SPF policy. It should have all of these - DKIM, DMARC, and SPF - because I use Fastmail. It's one of the reasons I have my domains hosted by Fastmail. Oh, this is really nice. He did a great job.

Steve: Yeah.

Leo: And this is nontrivial. I've got to say the stuff behind the scenes that he's doing to do the validation, that's great. Very nice.

Steve: Anyway, it keeps going as you press keys, and runs through, performs all the tests to demonstrate that the person doing your mailing has got their act together.

Leo: Somebody says, outofsync in our Discord says, that looks like a web dev's rsum. You're probably right, yeah.

Steve: That's good, yeah.

Leo: Yeah, you're probably right. A lot of CSS. Oh, look at that sliding on over. Look at that.

Steve: Ah. It's just beautiful.

Leo: Yes, my DKIM I think is okay. I see you've included a DKIM signature. I've retained - the signature passed validation. All right. There's a pass. Yeah, this is almost like a videogame. This is hysterical.

Steve: It's just gorgeous.

Leo: So I know I'm going to pass all three because I'm using Fastmail. But still, that's great.

Steve: Yup.

Leo: This is sponsored by URIports. Uh-oh.

Steve: Because you scrolled...

Leo: Did I scroll?

Steve: Yeah, you scrolled.

Leo: Don't scroll, kids. Do not scroll. Do not do what Leo did. Alignment is - yeah, looks like I'm okay; right? Yeah, there we go. Now I'm getting passes. It's because I scrolled up and couldn't find the pass. Okay, you need to work on that CSS a little bit. Yeah, we got passes everywhere. Very nice.

Steve: Yeah.

Leo: Very nice. So the bottom line is, because of spam, you're going to have to really start making sure your email's provider is doing all this. Of course if you're going Gmail to Gmail, you don't have to worry about that. But if you're going any other email provider including especially your ISP, you might want to check and make sure it's okay. I like this. What is it again, LearnDMARC.com.

Steve: DMARC, for those who are listening, yeah. Really, really great. I just stumbled on it because I was doing a little research into what am I going to have to do in order to not get, you know, not get blocked on DNS blocklists and all that crap.

Leo: This is actually, I guess, URIports is the company that does this DMARC monitoring. So that's why they're so good at it.

Steve: Okay. So recall that "Midnight Blizzard" is the dramatic renaming Microsoft gave to the Russian state-sponsored group, originally known as Nobelium, which most recently managed to crawl inside Microsoft's network to obtain access to data belonging to their uppermost top-level executives. As we covered last week, late Friday night before last, Microsoft slipped out the news that a lesser-protected system had succumbed to the

Russians after being sufficiently sprayed with passwords. What Microsoft shared at the time left no one feeling satisfied.

So last Thursday on the 25th, Microsoft attempted to offer additional useful information. Most observers, however, have still been left wanting. The reading between the lines that we did last week appears to have been correct.

At the top of last Thursday's lengthy update, Microsoft wrote: "As stated in the MSRC blog, given the reality of threat actors that are well resourced and funded by nation states, we are shifting the balance we need to strike between security and business risk. The traditional sort of calculus is simply no longer sufficient. For Microsoft, this incident has highlighted the urgent need to move even faster. If the same team were to deploy the legacy tenant today, mandatory Microsoft policy and workflows would ensure MFA and our active protections are enabled to comply with current policies and guidance, resulting in better protection against these sorts of attacks.

"Microsoft was able to identify these attacks in log data by reviewing Exchange Web Services (EWS) activity and using our audit logging features, combined with our extensive knowledge of Midnight Blizzard. In this blog, we provide more details on Midnight Blizzard, our preliminary and ongoing analysis of the techniques they used, and how you may use this information pragmatically to protect, detect, and respond to similar threats in your own environment.

"Using the information gained from Microsoft's investigation into Midnight Blizzard, Microsoft Threat Intelligence has identified the same actor has been targeting other organizations and, as part of our usual notification process, we have begun notifying these targeted organizations."

Okay, now, that's enough of that. As I noted, many if not most observers of Microsoft's handling of this incident have come away less than impressed. So I wanted to share the highlights of an important industry-shaping interview Alex Stamos conducted with CNBC last Friday. To remind everyone, Alex is a computer scientist. He obtained his EECS degree from Berkeley. Today he's an adjunct professor and lecturer at Stanford University's Center for International Security and Cooperation.

He first popped onto our map when he left Facebook after serving as their Chief Security Officer and then, in 2021, teamed up with ex-CISA director Chris Krebs. Recall that Chris was fired from his position as director of CISA by President Trump after CISA put out a statement declaring that the 2020 U.S. Presidential election had been the most secure election in American history. So Chris and Alex were both free, and they formed the Krebs Stamos Group. That group later became part of SentinelOne where Alex now has the title of Chief Trust Officer. He often serves as an expert witness in court and provides expert testimony to Congress.

Okay. So Alex's credentials are well established within the industry and government. The following, which I wanted to share, is what he posted last Friday following his interview on CNBC and following Microsoft's updated breach disclosure. The title Alex gave his LinkedIn posting was "Microsoft's Dangerous Addiction to Security Revenue." Under that headline, he wrote this.

He said: "On Monday, CNBC gave me a chance to discuss Microsoft's Friday night news dump of a new breach by Russian intelligence services, in which I called for more details from Microsoft so that other organizations could defend themselves. Yesterday, we gained a bit more transparency in the form of a blog post from 'Microsoft Security'" - again in air quotes - "the commercial security division of Microsoft."

"Some reactions," he wrote. "First, Microsoft buries the lead with this paragraph." Then he quotes them. "Using the information gained from Microsoft's investigation into Midnight Blizzard, Microsoft Threat Intelligence has identified that the same actor has been targeting other organizations and, as part of our usual notification processes, we've begun notifying these targeted organizations."

Alex says: "Translation: Since the techniques outlined in the blog only work on Microsoft-hosted cloud identity and email services" - whoops - "this means that other companies were compromised using the same flaws in Entra, better known as Azure Active Directory, and Microsoft 365. Microsoft's language here plays this up as a big favor they are doing the ecosystem by sharing their 'extensive knowledge of Midnight Blizzard' when in fact what they are announcing is that this breach has affected multiple tenants in their cloud products."

Leo: Oh, my god.

Steve: Uh-huh. And, he says, in a subsequent update to his original posting, Alex notes that Joseph Menn of the Washington Post has several sources indicating that at least 10 other companies were breached and will be disclosing those breaches soon.

Leo: This isn't the same as the Exchange Server vulnerabilities we've been talking about.

Steve: Not the one that China used...

Leo: This is Entra. This is different.

Steve: Yeah. Yeah.

Leo: Oh, boy.

Steve: Ten more breaches, at least 10 that were part of this.

Leo: They sure downplayed this. I knew they were, too. You could tell they were burying it, yup.

Steve: Yup. Yup. Second point he makes: "Microsoft continues to downplay the attack by abusing the term 'legacy.'" He says: "One of the big open questions from last week was how an attack against a 'legacy non-production test tenant' could lead to access to the emails of key Microsoft executives."

Leo: Yeah. Yeah. How did that happen?

Steve: He says: "We get a bit more detail in this paragraph." And we quote them now: "Midnight Blizzard leveraged their initial access to identify and compromise a legacy test

OAuth application that had elevated access to the Microsoft corporate environment. The actor created additional malicious OAuth applications. They created a new user account to grant consent in the Microsoft corporate environment to the actor-controlled malicious OAuth applications. The threat actor then used the legacy test OAuth application to grant them Office 365 Exchange Online full_access_as_app role, which allows access to mailboxes."

Alex says: "I've seen this fundamental problem in multiple investigations, including the one that Microsoft worked so hard to label as 'SolarWinds Incident.'" He says: "Azure AD is overly complex, and lacks a user experience that allows for administrators to easily understand the web of security relationships and dependencies that attackers are becoming accustomed to exploiting. In many organizations, Azure AD is deployed in hybrid mode, which combines the vulnerability of cloud," he says, "(external password sprays) and on-premise (NTLM and Mimikatz)" - meaning combining vulnerabilities both outside and in. He says: "Identity technologies in a combination that smart attackers utilize to bounce between domains, escalate privilege, and establish persistence."

He says: "Calling this a 'legacy tenant' is a dodge. This system was clearly configured to allow for production access as of a couple of weeks ago, and Microsoft has an obligation to secure their legacy products and tenants just as well as ones provisioned today. It's not clear what they mean by 'legacy'; but whatever Microsoft's definition, it is likely to be representative of how thousands of their customers are utilizing their" - meaning Microsoft's - "products today. Microsoft does," he says, "however, offer all of us some solutions."

Which brings us to point number three, which he labels: "Microsoft is using their own security flaws as an opportunity to upsell." He writes: "These sentences in the blog post deserve a nomination to the Cybersecurity Chutzpah Hall of Fame."

Leo: Oh, I love you, Alex.

Steve: "As Microsoft recommends that potential victims of this attack against their cloud-hosted infrastructure first detect, investigate, and remediate identity-based attacks using solutions like Microsoft Entra ID Protection."

Leo: Oh, yeah, you need that, yeah.

Steve: That's right. "Number two, investigate compromised accounts using Microsoft Purview Audit Premium."

Leo: Oh, yeah. You don't have that yet? Oh, we should get that, too.

Steve: Got to get that. And three, "Enforce on-premises Microsoft Entra Password Protection" - don't want to get sprayed - "for Microsoft Active Directory Domain Services." He says: "In other words, Microsoft is using this announcement as an opportunity to upsell customers on their" - meaning Microsoft's - "security products, which are apparently necessary to run their identity and collaboration products safely." He says: "This is morally indefensible, just as it would be for car companies to charge for seat belts or airplane manufacturers" - you know where he's going - "to charge for properly tightened door bolts."

He says: "It has become clear over the past few years that Microsoft's addiction to security product revenue has seriously warped their product design decisions, where they hold back completely necessary functionality for the most expensive license packs or as add-on purchases."

Okay, and I'm going to interrupt Alex for a moment to note that while all of this is highfalutin' enterprise stuff, I've long made the same point about Microsoft leveraging the insecurity of their "out of support" operating systems. They blithely offer additional years of extended security support for their otherwise "out of support" operating systems to their enterprise customers, while at the same time starving the end users of those same operating systems of that vital security in a bald effort to force users to move to newer operating systems which they neither need nor want. If the security updates are available anyway, deliberately withholding as ransom the patches to your defective operating system because you can is morally indefensible and reprehensible. Anyway, my two cents.

Referring to Microsoft's two recent posts, Alex says: "While these two arrogant and circumspect posts do, at least, admit 'the urgent need to move even faster' in securing their products," he says, "I would argue that Microsoft has a much deeper cultural problem to solve as the world's most important IT company. They need to discard this poisonous idea of security as a separate profit center and rededicate themselves to shipping products that are secure by default, while providing all security features to all customers." He says: "I understand the need to charge for log storage or human services, but we should no longer accept the idea that Microsoft's basic enterprise offerings, including those paid for by the U.S. taxpayer, should lack the basic features necessary to protect against likely attacks."

He says: "My current employer competes against some of these products from Microsoft. But if Microsoft did a better job by default, that would reduce the need for SentinelOne and other security vendors to provide basic safety protections. For all the language about the sophistication of the hackers behind this attack, there's nothing here that is outside the norm for ransomware groups attacking Microsoft's technologies, and Microsoft customers of all sizes should be concerned that these techniques will be deployed against them if they do not pay extra for the secure version of Microsoft's cloud products."

Leo: Wow.

Steve: "Twenty-one years after the Trustworthy Computing memo, it's once again time for some soul searching in Redmond."

Leo: You go, Alex.

Steve: So bravo, Alex.

Leo: Yeah, yeah.

Steve: You know, I love the system of free enterprise we enjoy in these United States. The profit motive provides strong impetus to innovate and provide value. But the lure of increased profit carries a danger when an executive faces a decision about whether to include a desirable and important feature in the base product, or to charge extra for it.

A crucial feature that's necessary for this system of free enterprise to deliver its maximum value to the public at large, rather than to simply further line the pockets of those executives and their shareholders, is competition. While it's an enviable position to be in, Microsoft is only able to get away with these usury practices because they have no real competition in the markets they dominate. This has been a problem for them in the past, and it may be again in the future.

Leo: Yikes.

Steve: I'm glad that people like Alex are saying this on CNBC and posting this in his LinkedIn feed because, you know, we need to shine a light on this.

Leo: We'll have to talk about it tomorrow on Windows Weekly, too, because this is really a much worse attack than we had hypothesized, at least the...

Steve: It is much worse than we thought.

Leo: Yeah, yeah.

Steve: It is broader, and apparently a huge number of Microsoft's customers...

Leo: Well, 10 anyway, yeah.

Steve: Yeah.

Leo: That's just the tip of the iceberg, I'm sure.

Steve: Bye.