Microsoft's Head in the Clouds

Description: What fascinating insights do we obtain from examining 3.4 million four-digit PINs? What plans are already underway as a backup for today's vulnerable GPS technology? How many Passkeys will websites store per account? And what's all this about Microsoft promising to get serious about their cloud-based services security?

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-974.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-974-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson's here. He's ready. He's champing at the bit, excited to get the show on the road. He's going to talk about what we learned by examining 3.4 million four-digit PINs. You guys have some bad habits. He'll also talk about an interesting old-school approach to solving the GPS fuzzing problem. And then Microsoft and how they're getting serious about cloud-based services security. Or are they? It's a Big Yellow Taxi moment coming up next. He'll explain on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 974, recorded Tuesday, May 14th, 2024: Microsoft's Head in the Clouds.

It's time for Security Now!, the show where we cover the latest news from the security-verse, as it were.

Steve Gibson: Second verse, same as the first.

Leo: There you go. That's Steve Gibson, the guru around here when it comes to security, privacy, and technology in general. Hi, Steve.

Steve: Yo, Leo. It's great to be with you again for this May 14th episode of Security Now! as we continue to approach 999.

Leo: Wow.

Steve: I actually had intended to make time to take some of these whiskers off, but I ran out of time, so our audience will have to - those unfortunate enough to be watching the video can just bear with us.

Leo: You're with it. That's hip now to have a little bit of a beard and scruff.

Steve: Oh, hey. See, if you just wait long enough...

Leo: It all comes around. Are you a three-blade, four-blade, or a five-blade guy?

Steve: Oh, god. There was the greatest piece that SNL did back in the day. Remember that?

Leo: That's why I asked.

Steve: And the tagline was "Because they'll believe anything."

Leo: It was their response, I think, to two blades. But anyway, it was so old that they thought five blades was funny. And then in fact that's exactly what we've got now.

Steve: I think the more the merrier as far as the blades go.

Leo: All the blades. All the blades.

Steve: That's right. So, okay. This is another of those episodes where there are such rich topics to discuss that we've going to do a few of them rather than a gazillion little tiddy bitty topics.

Leo: Careful there.

Steve: So we're going to look at what fascinating insights we have obtained from the examination of 3.4 million four-digit PINs.

Leo: Oh, I love this. I know what you're talking about. I saw this.

Steve: I posted this picture, a heat map that we'll be describing in detail here in a minute, on Twitter yesterday. Nothing I have ever posted before has generated so many little heart things, "likes" I guess we call them when we're a hipster like I am with my unshaven face. So that's going to be a lot of fun.

We're also going to look at an interesting surprise which is the plan that is already underway as a backup for today's vulnerable GPS technology which we talked about, we opened the show with last week, was like looking at what Russia is doing over in the Baltics and the vulnerability that we may not be taking seriously enough. It turns out we've got an answer for that.

Also there was a lot of feedback from our listeners who are avid Passkeys users about their experiences. I want to share some of those and essentially correct the record about one aspect that was wrong from last week.

And then we're going to take, as the title of today's podcast suggests, which is "Microsoft's Head in the Clouds," a look at a topic that everybody else in the industry has already covered, but we haven't yet here. And we're going to do it in our own way, as we always do, which is Microsoft's promise to get serious about their cloud-based services security. What happened? What has been found? And we have an interesting take, as we always do. So I think our listeners are going to have a great podcast. You know, surprise.

Leo: Surprise, surprise. You're going to have a great podcast. Well, we're very excited about that. Meanwhile, let me talk about one of our great sponsors as we get ready. Picture of the Week coming up, as well. All right. I'm ready for picture time.

Steve: So, yeah, this is just a quick simple cartoon. We've got two people sitting behind their laptops. One is sort of curious about what the other one is doing. And she looks over at his screen and says, "What are you doing on the dark web?" And his reply is, "I forgot my password, so I'm looking it up."

Leo: Of course. The NSA knows. And it's not been the hackers. Right.

Steve: That's where you'll find your password, on the dark web.

Leo: You bet, you bet, you bet. That's awesome.

Steve: Okay. So this, as I said, this is just a great chart. This is from the Information Is Beautiful project, which, you know, demonstrates that, if you graph things in creative ways, you can learn a lot. And this is a perfect example of that. 3.4 million four-digit PINs which were obtained from multiple data breaches were aggregated. Now, you know, this is a wonderfully enlightening graphic chart that I want to share. Unfortunately, the terms "graphic chart" and "listeners" are at odds.

Leo: You're going to have to describe it, Steve.

Steve: Yeah. I'm going to note that this delightful chart is at the top of this week's show notes. I tweeted it.

Leo: Yeah. You need to see it.

Steve: And I gave it a permanent GRC shortcut of pin, P-I-N. So anybody can see it at any time from grc.sc/pin, P-I-N. Okay. But, you know, I can do this verbally also. Okay. So this chart, as I said, takes 3.4 million four-digit PINs which were recovered from, and disclosed by, multiple data breaches.

Now, of course a four-digit PIN can have any value between 0000 and 9999. So there are 10,000 possible PINS. And this wonderful chart contains 10,000 little itty-bitty squares

arranged in a flat two-dimensional map. So it's got rows, you know, it's got 100 rows and 100 columns. And of course 100 times 100 is 10,000. So one way to think of this is that the first two digits of the PIN - which, you know, 00 through 99, specify one axis, and the last pair of digits specify the other. So every single possible four-digit PIN has its own square on this chart. And within this 3.4 million PIN dataset, the relative number of times every single possible PIN appears in the dataset determines the brightness of its square on the chart.

Okay. So what do we learn from this? Okay. Possibly the most prominent feature is a bright diagonal line running from the lower left corner of the chart, where both of the first two and the last two digits are 00, to the chart's upper right corner, where the first two and the last two digits are both 99. The diagonal line, then, is formed by all of the intermediate squares where their first two and last two digits are identical. And naturally like 00 in the far lower left, that's bright because a lot of people just chose 0000 as their PIN. And similarly, the very far upper right corner, also very bright because 9999 is many people's PIN. So there is some variation in the brightness along the diagonal which is interesting. You know, and of course human nature being what it is, the PIN 6969 appears to be overrepresented relative to its neighbors. No surprise.

Two other solitary bright spots would also not surprise anyone. They are the locations of the 1234 and 4321 PINs. Not very creative, and thus bright on the map. Another really interesting prominent line is the 20th line up from the bottom. Since lines are numbered from 0, the 20th line is the line for all PINs beginning with 19. And what's so interesting is that the line gets gradually brighter as it moves to the right, then dims a bit toward the end and wraps around a bit to the 20 line on the left. So what's going on here? Well, if you guessed people's birth year, you would be correct. PINs often begin, it turns out, with 19 and they appear to be brightest somewhere around 1980 seems to be the place where it's most, you know, most people have their PINs clustered there.

Leo: A lot of 40 year olds.

Steve: Exactly.

Leo: I would have thought it would be the baby boomers that would be the brightest, but maybe not.

Steve: Yeah, it's kind of fading out for us, Leo. And on the other hand, then so are we.

Leo: Yes.

Steve: Yeah. Another notable feature is a generally brighter region down at the lower left of the chart. This would be where both the first two and the last two digits form low numbers. Okay. Why? Because people used their month and day of birth within the month, running from 1 to 12 of course for the month, and then day of month 1 through 31. And what's interesting, there's a brighter horizontal stopping at 12 than the vertical stopping at 12, both which however are clear. This indicates that most people chose the ordering with the month first and the day of month second as their PIN.

Now, stepping way back from it and looking at the overall illumination, there's a top-to-bottom brightness variation, with it being brighter at the top and dimmer toward the bottom, suggesting that most PINs have low starting numbers. But there's less left-to-

right variation. So people are generally choosing four-digit PINs with, as I said, smaller first two digits, but for some reason more randomly distributed last two digits.

And the final really interesting observation is that whereas most of the chart shows varying shades of illumination, there are around 40 distinct cells that are black or nearly black. Like, I mean, dramatic contrast against their neighbors. In other words, out of all 10,000 possible four-digit PINs, there are around 40 of those that are significantly underrepresented.

Leo: Isn't that weird.

Steve: Isn't that? It's so odd.

Leo: Yeah.

Steve: For some reason...

Leo: It looks kind of randomly distributed, but maybe not.

Steve: Yeah, well, most of them have high...

Leo: They're mostly above 60, yeah.

Steve: Yeah. Almost all of them are in the upper third of the chart, so their first two digits are larger. For some reason, for example, very few people have chosen 6806. So if you're looking for a lesser chosen four-digit PIN, there you go.

Leo: They're all in there.

Steve: That's right.

Leo: Or 68 whatever this one is. You know, it's interesting. There are three dots on the 68 line.

Steve: Yeah. And in fact that first one on the 68 line was the 6806 that I just chose to highlight. But you're right. And looks like there's also three on the 60 or on the 70 line.

Leo: Yeah. Huh. I mean, how odd is that?

Steve: It's really non-random in that area. Okay. So, and as for the extremely low entropy skewing observed in the dataset, again, low entropy skewing, get this. Just the top 20, the top 20 most used PINs out of, remember, 10,000 that are possible, right, just the top 20 account for 27% of all PINs observed in use.

Leo: Oh. Oh, that's terrible.

Steve: Those top 20 are 1234, 0000, 7777, 2000, 2222, 9999, 5555, 1122, 8888, 2001 1111, 1212, 1004, 4444, 6969, 3333, 6666 1313, 4321, 1010.

Leo: If any of those sound like your PINs, you're in trouble.

Steve: Yeah, just very, very...

Leo: It means you can guess, you know, 10 or 20, and have a one in four chance of being right.

Steve: Right. If, for example, something prevented you from brute-forcing all 10,000, you would absolutely want to go for those 20 as your first 20 guesses.

Leo: It also means you should use more than four digits in your PIN; right?

Steve: Yeah. So I think we're still at four-digit PINs purely for historical reasons. It's just, you know, it's because that's, you know, once upon a time we didn't have computers, and people had to actually remember them. And I'm sure all the people used, you know, their month and day of birth, or the last four digits of their Social Security number, or digits from their license plate or, you know, something. The point being four digits was all they could actually remember. We didn't have technology to say, oh, yeah, I know, here's a string of 20 digits, you know, repeat after me.

Leo: Pick something, you know, what I always do is I pick the last four digits of a phone number, not my current phone number, but maybe my childhood phone number or phone number I particular remember because those are mostly pretty random. They certainly don't have anything to do with my birth date. I don't know. Or just pick something random. You can remember four digits. Or better yet, use an alphanumeric password, not a PIN.

Steve: Yeah. Well, and of course back once upon a time - or, no. I was going to say "once upon a time" we were keying them into our touchtone phones in order to authenticate ourselves.

Leo: Oh, yeah, right.

Steve: But even then, unless you used Q - I think was Q missing?

Leo: Q was missing; that's right.

Steve: There are a couple things that were not there on our models.

Leo: Well, you know where these are mostly still used is on ATM machines. I don't know of any ATM machine that uses more than four digits.

Steve: Yeah.

Leo: Right?

Steve: Again, because there's some backend, some old creaky backend machine that can only take four digits. Anyway, this was a huge win for our audience, who got a big kick out of it. So again, if you want to see what we were talking about, grc.sc/pin, and that will bounce you over to my site. I grabbed the - I actually could have just pointed to it. The original source was over on Reddit, and that got tweeted to me. But I was afraid that that might not last, you know, it could disappear. So I grabbed it and stuck it on GRC's server just because it's just such a cool infographic.

Okay. We started off last week with the piece in Wired about the growing threat to GPS. While the mischief Russia has been getting up to in the Baltic region is quite localized, we also noted that space is, sadly, not necessarily a benign environment anymore. A piece of our listener feedback which was generated by this discussion last week led me to look at what's being done about this.

Shaun Merrigan wrote. He said: "Steve, regarding SN-973 and GPS vulnerability, the U.S. is testing an updated version of the LORAN system which was shut down in the 1980s, called eLORAN. I've been monitoring the eLORAN test signals on 100kHz since August of 2023. My ancient LORAN receivers woke up and started giving me timing signals output again at that time, and have been receiving continuously ever since."

Okay. So this note from Shaun got me to poke around a bit, and I quickly learned that, indeed, there is an acute recognition of the inherent vulnerability of any satellite-based navigation system. LORAN is an abbreviation for Long Range Navigation, and the "e" in eLORAN stands for "enhanced." The original LORAN dates back from World War II. It's a ground-based navigation system that operates entirely differently from GPS. And of course entirely differently is what you want in something that's going to withstand an attack on GPS. You want something very orthogonal to the thing that you're trying to create a second solution for.

I found an interesting summary on the site GPS World. The article's title was "eLORAN: Part of the solution to GNSS vulnerability." Under the heading "Opposite and Complementary," the article leads with: "Though marvelous, GNSS are also highly vulnerable. eLORAN, which has no common failure modes with GNSS, could provide continuity of essential timing and navigation services in a crisis."

So here's what they explain. They said: "GPS fits Arthur C. Clarke's famous third law: 'Any sufficiently advanced technology is indistinguishable from magic.' Yet, it also has several well-known vulnerabilities including unintentional and intentional RF interference, the latter known as jamming; spoofing; solar flares; the accidental destruction of satellites by space debris; and their intentional destruction in an act of war; system anomalies and failures; and problems with satellite launches and the ground segment.

"Over the past two decades, many reports have been written on these vulnerabilities, and calls have been made to fund and develop complementary positioning, navigation, and timing, which are collectively referred to as PNT - Positioning, Navigation, and Timing - PNT systems. In recent years, as vast sectors of our economy and many of our daily

activities have become dependent on GNSS, these calls have intensified. A key component of any continent-wide complementary PNT would be a low frequency, very high power, ground-based system because it does not have any common failure modes with GNSS" - collectively meaning satellite based - "which are high frequency, very low power, and space-based. Such a system already exists, in principle. It is LORAN, which was the international PNT gold standard almost 50 years prior to GPS becoming operational in 1995. At that point, LORAN-C was scheduled for termination at the end of 2000.

"However, beginning in 1997, Congress provided more than $160 million to convert the U.S. portion of the North American LORAN-C service over to enhanced LORAN. In 2010, when the U.S. LORAN-C service ended, it was almost completely built out in the continental United States and Alaska. During the following five years, Canada, Japan, and European countries followed the United States' lead in terminating their LORAN-C programs. Today, however, eLORAN is one of several PNT systems proposed as a backup for GPS."

Okay. So first of all, it's great news that the U.S. has been seriously looking into a backup technology. Since I think our listeners will find this interesting, I'll share a little bit of background: "In the 1980s," this author writes, "I used LORAN-C to navigate on sailing trips off the U.S. East Coast. It had an accuracy of a few hundred feet and required interpreting blue, magenta, black, and green lines that were overprinted on nautical charts." And we'll get to why that is here in a minute. "The system was a modernized version of what was originally launched in 1958, a radio navigation system first deployed for U.S. ship convoys crossing the Atlantic during World War II. Its repeatability was greater than its accuracy. Lobster trappers could rely on it to return to the same spots where they'd been successful before, though they may have had some offset from the actual latitude and longitude.

"By contrast, eLORAN has an accuracy of better than 20 meters, and in many cases better than 10. It was developed by the U.S. and British governments, in collaboration with various industry and academic groups, to provide coverage over extremely wide areas using a part of the RF spectrum protected worldwide. Unlike GNSS" - which is to say GPS - "eLORAN can penetrate to some degree indoors, under very thick canopy, underwater, and underground. And it is exceptionally hard to disrupt, jam, or spoof. Unlike LORAN-C, eLORAN is synchronized to UTC and includes one or more data channels for low-rate data messaging, added integrity, differential corrections, navigation messages, and other communications.

"Additionally, modern LORAN receivers allow users to mix and match signals from all eLORAN transmitters and GNSS satellites in view. For the eLORAN system to cover the contiguous United States, between four and six transmission sites could provide overlapping timing coverage, and 18 transmission sites could provide overlapping positioning and navigation."

Okay. The article quoted Charles A. Schue, the CEO of UrsaNav. He said: "Think of a resiliency triad, consisting of GNSS, global; eLORAN, continental; and an inertial measurement unit with a precise clock. It is extremely difficult to jam or spoof all three sources of location and time at the same time, in the same direction, and to the same amount." In other words, great for protecting ourselves.

So it's cool that Shaun's ancient LORAN receivers woke up and began picking up LORAN signals. I don't know where he's located, but the intention is to cover the continental U.S. with multiple overlapping transmitters. The author of that article: "It had an accuracy of a few hundred feet and required interpreting blue, magenta, black, and green lines that were overprinted on nautical charts." Right? Why these fancy charts?

Imagine for nautical navigation, so you're out on the ocean somewhere, that two synchronized radio transmitters have been placed on the coast, several hundred miles apart. These two stations both emit a pulse of radio frequency energy at precisely the same time, and the pulses radiate outward spherically from each station at the speed of light, so 186,000 miles per second. So the ship at sea will receive these two pulses, but it does not know when they were sent. So it doesn't know its distance from these transmitters. The only thing it knows is the relative timing separation between them when they arrived.

Now, you can get out a pencil and paper and play with this a bit, but the LORAN system is called a "hyperbolic positioning system" because any given pulse separation describes a hyperbola. In other words, when a ship received a pair of pulses, their relative spacing would tell the ship's navigator which of many possible hyperbola, plotted on their navigational charts, the ship was currently sitting on. It would not yet have any way of knowing where it was sitting along that hyperbola, but it would have that one piece of information. The ship would get a fix on its position along that hyperbola by tuning to a different pair of transmitters. It would get another pulse spacing, which would identify another hyperbola on the navigation chart, and its location would be at the intersection of the first and second hyperbola.

So that's the way we located ourselves back during World War II. The good news is that today we have far more advanced technology with integrated circuits and fancy computers that can do all of this for us. But what hasn't changed is the decision to use low frequency, high power terrestrial transmitters to provide precise timing and location data as a backup for GPS. It's dispiriting to imagine that we might need it, but what's been going on over in the Baltics with Russia and GPS probably helped to get those projects funded here in the United States. So just a little very cool bit of technology.

Leo: That is really interesting, yeah, very cool.

Steve: Yeah. Hyperbolic positioning system. And on that note let's take a non-hyperbolic break, and then we're going to talk about Passkeys.

Leo: Yes, indeed. Coming up, we cover security from A to Z with this guy right here, Steve Gibson. And on we go with the show.

Steve: So, okay. A number of bits of feedback from our listeners. Jeff Urlwin, he said: "Just listened to SN. Passkeys are even worse based upon website implementation. Some sites" - get this - "use a cookie to 'know'" - he has in air quotes - "they issued you a Passkey. So even with 1Password, which supports and synchronizes among Passkeys," or among browsers, he says, "I can't use the Passkey from a different browser than originally set." He says: "CVS pharmacy is one with this bad implementation. Thanks for all your great shows."

RG tweeted: "Regarding Passkeys, for what it's worth, every website I have set up with a Passkey has let me set up multiple Passkeys, so I have not been limited to a single ecosystem."

Lachlan Hunt tweeted: "Regarding what you said in Episode 973 about Passkeys, you'll be happy to hear that every single account for which I've been able to register a Passkey and store in 1Password has been able to support registering multiple Passkeys. For some of my most important accounts, I've registered additional Passkeys stored on my YubiKeys. In my experience, storing Passkeys in 1Password has been fantastic. The only

major issue I've encountered has been with certain sites, for example PayPal and LinkedIn, that do browser sniffing to unnecessarily prevent Passkeys from being used within Firefox. This can usually be worked around by simply spoofing the User-Agent string." But again, you know...

Leo: Why would you do that? That's weird.

Steve: We've talked about poor implementations. So poor implementations certainly exist. And Miguel Frade said: "Hi, Steve. In SN-973 you read Dave Brenton's questions about using a backup YubiKey. To complement your answer, I'd like to share my personal experience. I've owned two YubiKeys for several years, one with me all the time, and a backup stored in a safe place. Some services, like Gmail, GitHub, and Bitwarden, allow us to register more than one YubiKey. In case of Bitwarden's Family plan it allows registering up to five YubiKeys. I guess it should be the same for Bitwarden's individual premium plan.

"Unfortunately," he writes, "PayPal only allows registering one YubiKey. Regarding the question 'Can the same key be applied to two different people,'" he says, "the answer is yes, if we're talking about the physical key YubiKey. Each service will use one of the 25 available slots inside the YubiKey, regardless of the person owning the account." He finishes, "I hope this information can be useful to other SN listeners. All the best, Miguel."

Okay. So last week's discussion of this generated, as I've said, significant feedback from our listeners. And the thing that stood out more than anything was that everyone showed a somewhat different set of facts. Some said that WebAuthn/FIDO2 providers would allow any number of Passkeys to be registered with a service. Some said that only one could be. And others like Miguel noted that this varied by provider, with PayPal, for example, only allowing for a single registration. If this were true, it would mean that separate "his" and "her" YubiKeys could not be used with some services. But all other listeners noted that they had never encountered a site that did not allow for any number of Passkey registrations. And doing so is part of the Passkeys specification. So all sites should.

I went over to PayPal to take a look, and their Passkeys management page makes it very clear that they support multiple Passkeys without any trouble. However, PayPal appears to only support Passkeys generated by iOS and Android devices. Its FAQ is quite clear about that, and there's no mention of YubiKeys. So perhaps that's what stopped Miguel. He didn't actually register a Passkey at all over on PayPal. He was using PayPal's longstanding and much older...

Leo: Football.

Steve: ...multifactor authentication.

Leo: Football.

Steve: Right.

Leo: My first football, yeah.

Steve: The multifactor authentication over on PayPal. But this further demonstrates the mess that we're currently working through. The fact that something stopped Miguel even though he has a perfectly secure authentication device, arguably even though smartphones are now very secure, but you can argue that - you can make a strong argument that the YubiKey being so focused and single purpose and simpler, you know, and doesn't have multiple radios hooked to it, is more secure than the two smartphone brands that PayPal does support. But this all shows that we're still in the early days of this technology. You get a YubiKey which supports Passkeys. PayPal supports Passkeys. But PayPal won't support a Passkey generated by a YubiKey.

One thing that all the feedback made very clear was that many of our listeners have jumped into the Passkeys world with both feet. They like them, and I think that's great. Really. I think that those of us in the industry who are grousing at the moment - and Paul Thurrott, for example, went on a nice rant again about this last week.

Leo: Easily triggered.

Steve: Yes. He's doing so, well, no, all of the users are doing - those of us in the industry are ranting because we're disappointed with the rollout and are impatient for Passkeys to live up to their potential. We know that change takes time, and that this is still the very early days for this new technology. Browser and browser extension support for original username and password authentication has created a system that's mostly good enough for now, with second-factor authentication adding additional protection where needed. Your football, Leo.

None of us can predict the future, and today's Passkeys support remains really disappointing. But in the grand, you know, if nothing else, in the fact that so few sites, you know, have jumped on the bandwagon. On the other hand, why would they if it's not urgent for them? But in the grand scheme, relative to how slowly new technology is adopted, Passkeys only became available yesterday. Once the various kinks are ironed out and any device we wish to use can supply a previously generated Passkey to a website, the traditional problems with passwords will begin to fade.

I think that the most compelling use case of all is the typical user, you know, and there are a lot more of those typical users than there are listeners to Security Now!, the typical user who has no interest whatsoever in any of this. They could care less. They're using an iOS or Android smartphone, a Mac or Windows device having strong biometric hardware authentication. They visit a site which newly supports Passkeys, and the site says: "Hey. How would you like to never need to use a username or password to log in with this device ever again?" You know, who's not going to click yes? Any regular user will think, "That's great. Passwords are annoying as hell. If I don't need to use one here anymore, count me in."

Presumably, and this is what remains unknown, whether and to what extent additional sites will offer this support over time. If it does succeed in setting a new standard, then Passkeys will just gradually and organically seep into the world and become the way Internet users authenticate. I think, you know, we're excited by the potential, those of us who are into the technology, and we want it to happen immediately. But it's just going to take some time. And clearly a lot of the listeners of this podcast have been curious about this. And mostly their experiences have been all good, which I think is great.

Okay. Microsoft's Head in the Clouds. SC Magazine's headline read: "Sweeping cybersecurity improvements pledged by Microsoft," and follows with "Numerous

cybersecurity incidents" - I'm sorry. Numerous cyber secure - but there were incidents. But they wrote: "Numerous cybersecurity enhancements..."

Leo: Enhancements in response to incidents, that's what it was.

Steve: Uh-huh, "...will be adopted by Microsoft to address the woeful security failures driven by poor cybersecurity practices and lax corporate culture identified in a report issued by the Cyber Safety Review Board last month." And SecurityWeek carried the headline: "Microsoft Overhauls Cybersecurity Strategy After Scathing CSRB" - that's the same, the Cyber Safety Review Board - "Report." And then they follow with "Microsoft Security Chief Charlie Bell pledges significant reforms and a strategic shift to prioritize security above all other product features," basically saying we're going to stop with the features here. Although one could argue that they're not stopping with their AI push. But otherwise, we're really prioritizing security.

Now, anyone who's been following this podcast for the past year will have heard me go off on Microsoft over their truly astonishing apparent lack of concern or accountability over egregious security practices. Doing so always leaves me feeling a bit odd, since I'm sitting in front of Windows machines, all of my coding for the PC has been for Microsoft operating systems, from DOS through desktop and server, and I love the Windows working and development environments.

But as we've clearly documented on this podcast over and over, security researchers repeatedly hand Microsoft every detail, complete with working proofs of concept demonstrations, for various vulnerabilities which Microsoft will seemingly ignore for months and even years until that vulnerability is actually used to cause a highly public catastrophe. And only then will Microsoft apparently think, "Huh. Why does that exploit path have a familiar ring to it?" Right. You know? And we understand why; right? Microsoft is a monopoly. You cannot build a large modern enterprise without Microsoft glue. Too many things require Microsoft. So the simple fact is, Microsoft does not have to care. And we've seen example after example of Microsoft not doing anything it does not want to do.

All of this makes Microsoft's recent pronouncements about their new focus upon security all the more interesting. Two weeks ago, the "Cybersecurity Dive" site posted an article with a headline that caught my eye. They wrote: "At Microsoft, years of security debt come crashing down," and the subhead was "Critics say negligence, misguided investments, and hubris have left the enterprise giant on its back foot."

They wrote: "Years of accumulated security debt at Microsoft are seemingly crashing down upon the company in a manner that many critics warned about, but few ever believed would actually come to light. Microsoft is an entrenched enterprise provider, owning nearly one-quarter of the global cloud infrastructure services market and, as of the first quarter last year, nearly 20% of the worldwide Software as a Service application market. Though not immune from scandal, in the wake of two nation-state security breaches of its core enterprise platforms, Microsoft is facing one of its most serious reputational crises. Adam Meyers, Senior Vice President at CrowdStrike, said: 'It's certainly not the first time a nation-state adversary has breached Microsoft's cloud environments. After so many instances, empty promises of improved security are no longer enough.'"

Okay. Now, to review a bit, in January, Microsoft said a Russia-backed threat group called Midnight Blizzard gained access to emails, credentials, and other sensitive information from top Microsoft executives as well as certain corporate customers and a number of federal agencies. We're going to see that the numbers were actually

somewhat worse than that. Then in early April, the federal Cyber Safety Review Board released a long-anticipated report which showed the company failed to prevent a massive 2023 hack of its Microsoft Exchange Online environment. The hack by a People's Republic of China-linked espionage actor led to the theft of 60,000 State Department emails and gained access to other high-profile officials, actually many.

Just weeks ago, CISA issued an emergency directive to order federal civilian agencies to mitigate vulnerabilities in their networks, analyze the content of stolen emails, reset credentials, and take additional steps to secure Microsoft's Azure accounts. While the order only applies to Federal Civilian Executive Branch agencies, CISA warned other organizations could be impacted. For many critics of Microsoft, the events of the past nine months are the logical conclusion of a company that has ridden the wave of market dominance for decades and ignored years of warnings that its product security and practices failed to meet the most basic standards.

"AJ Grotto, the director of the Program of Geopolitics, Technology, and Governance at the Stanford Cyber Policy Center and a former White House Director for Cyber Policy said: 'In a healthy marketplace, these would be fireable offenses. Regrettably, the marketplace is far from healthy. Microsoft has the government locked in as a customer, so the government's options for forcing change at Microsoft are limited, at least in the short term.

"The concern was, and is, that Microsoft's security gaps would potentially lead to catastrophic outcomes. According to Karan Sondhi, CTO at Trellix: 'Microsoft needs to dedicate its internal resources towards zero-trust initiatives and make new investments in its infrastructure. Currently,' he says, 'Microsoft directs the vast majority of their security investments toward revenue-generating roles instead of internal security roles.'" And we'll come back and talk about that here toward the end.

"Microsoft has a considerable stake in the cloud security space. Not only is Microsoft one of the world's largest cloud providers, but according to Microsoft's CEO Satya Nadella during the company's fiscal second quarter conference call in January, 'It is also a major security provider to the enterprise. Microsoft has more than one million security customers, with 700,000 using four or more of its security products. Microsoft generates more than $20 billion in revenue per year from its security business." In other words, by selling security that one could argue ought to be baked in.

Okay, now, I should note here for the record that I don't have any feelings at all of schadenfreude. Really. I'm not the least bit happy that it took some seriously frightening and damaging security lapses within Microsoft to get them to finally start thinking about taking security seriously. It would have been better for everyone if those breaches never occurred. But, unfortunately, all evidence suggests that nothing would have changed at Microsoft ever but for those breaches. So the way things have been going, it was probably inevitable.

The trouble they've fallen into feels like the result of a cycle, a cultural cycle within Microsoft. We've witnessed such cycles within Microsoft in the past. I think that happens when a company grows so much that it keeps creating very wealthy upper management, who then, no longer needing to work, eventually leave the company. But they're not the only things that leave. What leaves with them is their deep understanding of the culture their leadership created while they were there. Those who replace them think they know how to keep everything running. But not having created it, they lack the same deep experience-based understanding of what's important. And then, over time, the ship drifts off course. Since I cannot even conceive of captaining a ship the size and complexity of Microsoft, it doesn't surprise me that it might lose its way from time to time; you know? I'm amazed it's still afloat.

Early last month, the Department of Homeland Security's CSRB, as I've mentioned, the Cyber Safety Review Board, released their findings following a deep and detailed investigation into Microsoft's recent security breach troubles. I'm going to share the summary of that report from the Cybersecurity Dive people.

Leo: On we go with the sad tale of Microsoft finally waking up to the security needs. This, by the way, was not their first time.

Steve: No.

Leo: At this rodeo. They keep waking up to security requirements. It's like, one more memo. Okay. That'll fix it.

Steve: Yeah. And as we'll see here when I sort of summarize and wrap this all up, that it's not clear what this means, but we can hope.

Leo: Pray.

Steve: Yeah. The thing that this article from - I've lost my cursor - from the Cybersecurity Dive brings is some additional color and quotes and background from other people. They wrote: "The CSRB report laid out a blistering assessment of a corporate culture that has failed for years to take cybersecurity seriously. The report was designed to assess the company's response to the summer 2023 breach from the People's Republic of China-linked threat actor that breached the company's Microsoft Online Exchange environment. However, it also laid out a security culture that failed to adhere to the most basic standards, given the enormous market power that Microsoft yields across modern business applications in government and the private sector.

"One of the more damaging findings was that Microsoft learned of the attacks only because the State Department had set up an internal alert system after purchasing from Microsoft at additional cost a G5 license. Customers who failed to purchase the enhanced security license were not able to see the extensive logging capabilities that would have alerted them to a breach." And we'll get back to the implications of that also.

"Many in the security community see the CSRB report and the recent CISA emergency directive as direct indictments, not only of Microsoft's security culture, but a government that has allowed Microsoft to maintain lucrative government contracts with no fear of competition across many of its services. Mark Montgomery, senior director at the Center on Cyber and Technology Innovation at the Foundation for Defense of Democracies said: 'The federal government gets off the hook a little easy in this report. Despite significant encouragement from outside experts, the Biden administration and its predecessors have failed to treat cloud computing as a national critical infrastructure that is itself critical to maintaining the security of our other national critical infrastructures.'" So it's a, you know, a critical infrastructure infrastructure.

"Senator Ron Wyden, who called for a federal investigation following the State Department email hack, said the federal government shared responsibility for the negligent behavior disclosed in the report. Wyden said Microsoft has been rewarded with billions of dollars in federal contracts, while not being held to account for even the most basic security standards." Wyden told the author of this article: "The government's dependence on Microsoft poses a serious national security threat, which requires strong

action." Now, think about that for a minute. "The government's dependence on Microsoft poses a serious national security threat." I know that the practice of politics generates a great deal of rhetoric, but that's not something you want a well-placed and respected U.S. Senator saying about your company.

And speaking of rhetoric, Microsoft knows how to play the game with the best of them. "Microsoft officials said they understand the larger concerns raised by the summer 2023 attacks, as well as the continued threat from Midnight Blizzard and other nation-state actors. The company is working to make extensive changes in its engineering processes, improve its relationships with the security community" - wow, listen to the security community? What a concept - "and its responsiveness to customer needs.

"Bret Arsenault, corporate VP and chief cybersecurity advisor at Microsoft, said in a statement: 'We're energized and focused on executing Microsoft's Secure Future Initiative commitments. And this is just the beginning. We commit to sharing transparent learnings..."

Leo: They love that word, I don't know why.

Steve: Oh, god, I hate that word.

Leo: I know.

Steve: "...transparent learnings." There's got to, you know, I thought, Leo, isn't there a better word? But you've turned an activity into a noun.

Leo: Right.

Steve: And I guess there's no helping you.

Leo: There's no helping you.

Steve: After you've done that. Yes. I just wish that some of Microsoft's customers would have some walkings, and would walk away.

Leo: From earnings.

Steve: Yeah, my god. Okay. So interestingly, one of the problems with being transparent about what's being fixed is that the process of enumerating all the improvements also serves to enumerate just how bad things had been allowed to become.

Leo: Oh, yeah, exactly.

Steve: Uh-huh. Listen to these numbers. Bret said that since the launch of the company's Secure Future Initiative, the company has sped up related engineering work

in several areas. Okay, he calls it a "speed up." Well, he lists four. He says: "Microsoft has accelerated the lifecycle management of tenants, with a focus on either unused or older systems. The company eliminated more than 1.7 million Entra ID systems related to used, aging, or legacy technology." In other words, there were 1.7 million Entra ID systems that could be eliminated, but had not been. They were just, you know, hanging around, waiting to be abused. "It has also made multifactor authentication enforcement automatic across more than one million Entra ID tenants." Which, again, says that they weren't before.

Also: "More than 730,000 apps have been removed across production and corporate tenants that were either out of lifecycle or were no longer meeting current standards." Nearly three quarters of a million apps were just, again, you know, left alone. Left there. Even though they were no longer serving any purpose. As we know, fundamental to security is taking an employee's badge and then removing all their passwords from the system before they have a chance to use them. Three quarters of a million apps were left there.

Also Microsoft said: "New employees and vendors are now being given short-term credentials to make impersonation and credential theft more difficult. More than 270,000 have been implemented thus far."

And finally: "The company's internal multifactor authentication implementation using Microsoft authenticator has been enhanced by eliminating a call feature and relying on an in-app login feature. This change covers more than 300,000 employees and vendors." Again, 300,000 employees and vendors were using an insecure feature of the multifactor authentication that likely made it easier to use, but was less secure. So, okay. Gee. I guess we should fix that.

Okay. So I've observed for some time here on the podcast that one of the reasons Microsoft has been acting the way it has, has been able to act the way it has for so long without correction, is that until now its negligence had no consequence, exactly as Senator Ron Wyden observed. For this article, Dante Stella, an attorney at Dykema and a specialist in incident response, said that enterprise customers do not usually walk away in the face of nation-state threats against Microsoft, in part due to its enormous presence as a cloud provider.

Dante was quoted: "Many switched to Exchange Online or Microsoft 365 to get away from on-prem servers and managed service providers. If the only other choice is going 'back,' or a potentially disruptive switch to another platform like Google Workspace, they will most often just ride it out and trust Microsoft to fix the issues." Right. The customers may be unhappy; but, due to Microsoft's dominance in the market, that unhappiness is never reflected in Microsoft's bottom line. So why change anything?

As we know, I always want to go to the source. So after reading this piece I was curious to see the report from the Cyber Safety Review Board. Now, the full report I'm not going to share. It's 34 pages of quite eye-opening content. But the short Executive Summary at the start paints the picture. Here's what the review board found. This is the actual report from this CSRB.

They wrote: "In May and June 2023, a threat actor compromised the Microsoft Exchange Online mailboxes of 22 organizations and over 500 individuals around the world. The actor known as Storm-0558, hereinafter simply as Storm, and assessed to be affiliated with the People's Republic of China in pursuit of espionage objectives, accessed the accounts using authentication tokens that were signed by a key Microsoft had created in 2016." In other words, that key had never expired or been rotated in seven years.

Okay. They say: "This intrusion compromised senior United States government representatives working on national security matters, including the email accounts of Commerce Secretary Gina Raimondo, United States Ambassador to the People's Republic of China R. Nicholas Burns, and Congressman Don Bacon.

"Signing keys, used for secure authentication into remote systems, are the cryptographic equivalent of crown jewels for any cloud service provider. As occurred in the course of this incident, an adversary in possession of a valid signing key can grant itself permission to access any information or systems within that key's domain. A single key's reach can be enormous, and in this case the stolen key had extraordinary power. In fact, when combined with another flaw in Microsoft's authentication system, the key permitted Storm to gain full access to essentially any Exchange Online account, anywhere in the world. As of the date of this report, Microsoft does not know how or when Storm obtained the signing key.

"This was not the first intrusion perpetrated by Storm, nor is it the first time Storm displayed interest in compromising cloud providers or stealing authentication keys. Industry links Storm to the 2009 Operation Aurora campaign that targeted over two dozen companies, including Google; and the 2011 RSA SecurID incident, in which the actor stole secret keys used to generate authentication codes for SecurID tokens, which were used by tens of millions of users at that time. Indeed, security researchers have tracked Storm's activities for over 20 years.

"On August 11, 2023, Secretary of Homeland Security Alejandro Mayorkas announced that the Cyber Safety Review Board (CSRB, or the Board) would 'assess the recent Microsoft Exchange Online intrusion and conduct a broader review of issues relating to cloud-based identity and authentication infrastructure affecting applicable cloud service providers and their customers.'

"The Board conducted extensive fact-finding into the Microsoft intrusion, interviewing 20 organizations to gather relevant information. Microsoft fully cooperated with the Board and provided extensive in-person and virtual briefings, as well as written submissions. The Board also interviewed an array of leading cloud service providers to gain insight into prevailing industry practices for security controls and governance around authentication and identity in the cloud." In other words, they really did look at the entire industry in order to support their conclusion that Microsoft stood out as negligent. This wasn't common practice, what Microsoft was doing, the way Microsoft was operating.

They wrote: "The Board finds that this intrusion was preventable and should have never occurred. The Board also concludes that Microsoft's security culture was inadequate and requires an overhaul, particularly in light of the company's centrality in the technology ecosystem and the level of trust customers place in the company to protect their data and operations. The Board reaches this conclusion based on seven points. One, the cascade of Microsoft's avoidable errors that allowed this intrusion to succeed; second, Microsoft's failure to detect the compromise of its cryptographic crown jewels on its own, relying instead on a customer to reach out to identify anomalies the customer had observed."

And I'll just take a moment here to say, elsewhere they explain: "The State Department was the first victim to discover the intrusion when, on June 15th, 2023, State's security operations center (SOC) detected anomalies in access to its mail systems. The next day, State observed multiple security alerts from a custom rule it had created, known internally as 'Big Yellow Taxi,' that analyzes data from a log known as MailItemsAccessed, which tracks access to Microsoft Exchange Online mailboxes. State was able to access the MailItemsAccessed log to set up these particular Big Yellow Taxi alerts because it had purchased Microsoft's government agency-focused G5 license that

includes enhanced logging capabilities through a product called Microsoft Purview Audit Premium. The MailItemsAccessed log was not accessible without that premium service."

Leo: Oh. This is why Microsoft gets in trouble, because they demand you pay for security.

Steve: Exactly.

Leo: But wait a minute. Big Yellow Taxi is the name of the tool?

Steve: Big Yellow Taxi is the name they gave to the intrusion detection rules for determining whether the Microsoft Exchange online mailboxes were being maliciously accessed.

Leo: It's the name of a Joni Mitchell song. But I don't really know why they used that. That's weird. Okay.

Steve: Big Yellow Taxi. The Big Yellow Taxi alert went off, Leo, and they thought, uh-oh. Something's wrong.

Leo: Okay. It's hard to take that seriously, I'll be honest with you.

Steve: That's why we normally don't get those internal names exposed to the public.

Leo: Yeah, yeah.

Steve: Yeah. When we find out that the State Department has named their intrusion rule "Big Yellow Taxi," it's like, uh.

Leo: Somebody's a Joni Mitchell fan. That's all I can say.

Steve: Is it a miracle you guys discovered this? Anyway. Also, they said, "The Board's assessment of security practices at other cloud service providers, which maintained security controls that Microsoft does not. Fourth, Microsoft's failure to detect a compromise of an employee's laptop from a recently acquired company prior to allowing it to connect to Microsoft's corporate network in 2021." So a compromised laptop was hooked up to Microsoft's network after Microsoft acquired a company, and that was a problem.

Also, number five: "Microsoft's decision not to correct, in a timely manner, its inaccurate public statements about this incident, including a corporate statement that Microsoft believed it had determined the likely root cause of the intrusion when in fact it still has not," meaning even to this day. And "Even though Microsoft acknowledged to the Board in November 2023 that its September 6th, 2023 blog post about the root cause was inaccurate, it did not update that post until March 12th of 2024, as the Board was

concluding its review, and only after the Board's repeated questioning about Microsoft's plans to issue a correction." In other words, what? Oh. Oh, you mean, what we immediately said back in September? Yeah, we've been meaning to change that. But, gee, you know, we just haven't gotten around to it.

Number six: "The Board's observation of a separate incident, disclosed by Microsoft in January of this year, 2024, the investigation of which was not in the purview of the Board's review, which revealed a compromise that allowed a different nation-state actor to access highly-sensitive Microsoft corporate email accounts, source code repositories, and internal systems."

And, finally, number seven: "How Microsoft's ubiquitous and critical products, which underpin essential services that support national security, the foundations of our economy, and public health and safety, require the company to demonstrate the highest standards of security, accountability, and transparency." As opposed, obviously, to the lowest.

"Throughout this review," they wrote, "the Board identified a series of Microsoft operational and strategic decisions that collectively point to a corporate culture that deprioritized both enterprise security investments and rigorous risk management." Deprioritized. "To drive the rapid cultural change that is needed within Microsoft, the Board believes" - and I love the fact that here's the government telling - this board on the government suggesting how Microsoft should run things. They said: "The Board believes that Microsoft's customers would benefit from its CEO and Board of Directors directly focusing on the company's security culture and developing and sharing publicly a plan with specific timelines to make fundamental, security-focused reforms across the company and its full suite of products. The Board recommends that Microsoft's CEO hold senior officers accountable for delivery against this plan.

"In the meantime, Microsoft leadership should consider directing internal Microsoft teams to deprioritize feature developments across the company's cloud infrastructure and product suite until substantial security improvements have been made in order to preclude competition for resources." In other words, if you don't have enough people available to fix your security, why don't you just hold off on all those new improvements that you were planning and get your security house in order first.

Leo: Yeah.

Steve: How would that be?

Leo: Yeah.

Steve: Because, you know, national security and billions of dollars in contracts that we keep providing to you and rolling over year after year after year.

Leo: And Big Yellow Taxi.

Steve: How about fixing?

Leo: Yeah.

Steve: How about fixing some things?

Leo: Yeah.

Steve: That's right. Okay, Leo. Let's take our last break.

Leo: Okay.

Steve: Then I want to look back at mainframe computing and why where we are is like where we were then.

Leo: Oh, that's interesting. A little bit of history to tie into the present.

Steve: Yup.

Leo: I like it. Now let us conclude our journey down the highway of insecure operating systems with Steve Gibson.

Steve: So one of the earliest breakthroughs in computing was the introduction of a concept that came to be called "timesharing."

Leo: Yes.

Steve: Back then, mainframe computers were incredibly expensive to purchase and operate. A single machine installation was planned years in advance. Electrical power and cooling was plumbed. Large rooms were set aside. And these machines had their own staff and managers. The bean counters, who occupied the upper floors, quickly realized that their costs were the same, whether or not the monstrously expensive machine in the basement was busily working for them or sitting idle. So the question soon became how do we keep this massive investment of ours busy? And the answer was timesharing.

Timesharing meant that a great many people could share the machine's time. This worked because most people spent most of their time staring at the screen of their timesharing terminal reading what had just been displayed, deciding what to do next, and then slowly punching out the next command they wished to issue. If it had been just one person, the mainframe would have been bored to death. But the bean counters perked right up when they learned that their machine in the basement could keep thousands of their employees, literally everyone in the building, busily poking away at their keyboards and never waiting long for their next screen of data to be presented.

Most of the company's thousands of employees never visited the basement. They weren't allowed to. Security was high because too much was at stake. All of the company's jewels had been concentrated into a single small region, and those who had privileged access wore white coats and prominently displayed ID tags. To most of the rest of the company, these tenders of the machine did not appear to speak English, and what

exactly they did down there in the basement was shrouded in rumor and mystery, with some not appearing to emerge for days on end.

I've painted this picture of the past because it's interesting that it's a close approximation of what has gradually and organically re-evolved today, mostly of its own accord. Part of it is upside down because, instead of computing being done in the basement, today it's being done in the clouds. But we have a very similar concentration of value into a small, high-security, tightly controlled area to which few people have access. And the concept of resource sharing exists pervasively. Thanks to the miracle of the global Internet, the networking wires that interconnect the servers are literally being shared by everyone in the world. And the use of virtual machine technology, which shares physical processor resources among a great many more virtual processors, is the essence of timesharing. No single virtual machine needs to, or can, keep a high-powered cluster of processor cores completely busy; so a much larger number of virtual machines can simultaneously share that single powerful resource with many others.

This move to the cloud does not feel like yet another phase. This feels like an inevitable evolution. Earlier I noted that Dante Stella had been quoted saying: "Many switched to Exchange Online or Microsoft 365 to get away from on-premises servers." I think this represents an inevitable evolution because, just as happened in the past era of mainframe computing, the computational resource we were able to create far outstripped the needs of the typical user. Today's processors are so powerful that most PC users today are only using a small fraction of their system's capabilities. When this is scaled up to an enterprise of 10,000 employees, the wasted resources are astonishing. Since most people today are, just as they were 50 years ago, staring at a screen, taking the time to figure out what it says, then poking away at their keyboard to indicate what they want to do next, we've returned to the mainframe era, and what we're sharing are cloud-based resources.

And I'll just note that the recent evolution of interactive cloud-based AI models represents another example where sharing a single massive resource among many users is vastly more economical than giving each user their own instance. And even though local mini-models can be used, thanks to our astonishing computing power, the best models will be continuously training, which requires massive connectivity and a far greater level of processing.

Okay. So how did Microsoft get into trouble? There's that old observation, which I've heard isn't actually true, but it makes for a great example nevertheless, that if you toss a frog into a pot of boiling water it will immediately jump out. But if the frog is placed into cold water and the temperature is slowly increased, it won't notice the change.

What this report makes clear is that the world has awoken to just how utterly dependent we have become upon computing in the cloud. It happened so gradually, so incrementally and slowly, with one day following after the next, with one company after another deciding that the economics of moving their communications infrastructure into the cloud made the most sense, that, just as with the apocryphal frog, we've arrived at a position where the security of our cloud computing can no longer be considered an afterthought, and it can no longer be taken for granted.

I initially skipped past the opening statement from the chair and deputy chair of the CSRB's report because now we have some context that they had when they wrote it. They said: "It is not an exaggeration to say that cloud computing has become an indispensable resource to this nation and, indeed, to much of the world. Numerous companies, government agencies, and even some entire countries rely on this infrastructure to run their critical operations, such as providing essential services to customers and citizens. Driven by productivity, efficiency, and cost benefits, adoption of

these services has skyrocketed over the past decade; and, in some cases, they have become as indispensable as electricity.

"As a result, cloud service providers (CSPs) have become custodians of nearly unimaginable amounts of data. Everything from Americans' personal information to communications of U.S. diplomats and other senior government officials, as well as commercial trade secrets and intellectual property, now resides in the geographically distributed data centers that comprise what the world now calls the 'cloud.'

"The cloud creates enormous efficiencies and benefits; but, precisely because of its ubiquity, it is now a high-value target for a broad range of adversaries, including nation-state threat actors. An attacker that can compromise a CSP can quickly position itself to compromise the data or networks of that CSP's customers. In effect, the CSPs have become one of our most important critical infrastructure industries. As a result, these companies must invest in and prioritize security consistent with this 'new normal,' for the protection of their customers and our most critical economic and security interests."

So, getting back to your comment, Leo, what will all this mean to Microsoft, and what will it mean to us? I have no idea, and neither does anyone else. For one thing, big changes take time. What Microsoft's rhetoric promises is a major reorganization of their corporate priorities. They're saying this because it has become clear to everyone that a major reorganization of their corporate priorities is exactly what will be needed.

I want to conclude our look at this by sharing the report of Microsoft's actions once the State Department's "Big Yellow Taxi" honked its horn, indeed noting that there was a problem. I want to share it because it reads like a detective novel, which I know our listeners will enjoy; and because, while it's part of the same scathing report, it paints Microsoft in a good light and shows what this behemoth is capable of doing when it wants to, or maybe needs to.

The report wrote: "Though the alerts showed activity that could have been considered normal - and indeed State had seen false-positive Big Yellow Taxi detections in the past - State investigated these incidents and ultimately determined that the alert indicated malicious activity. State triaged the alert as a moderate-level event; and on Friday, June 16th, 2023, so coming up on a year ago a month from now, its security team contacted Microsoft. Microsoft opened and conducted an investigation of its own, and over the next 10 days ultimately confirmed that Storm-0558 had gained entry to certain user emails through State's Outlook Web Access. Concurrently, Microsoft expanded its investigation to identify the 21 additional impacted organizations and 503 related users impacted by the attack and worked to identify and notify impacted U.S. government agencies.

"Microsoft initially assumed that Storm had gained access to State Department accounts through traditional threat vectors, such as compromised devices or stolen credentials. However, on June 26th, 10 days after the initial alert, Microsoft discovered that the threat actor had used OWA (Outlook Web Access) to access emails directly using tokens that authenticated Storm as valid users. Such tokens should only come from Microsoft's identity system, yet these had not. Moreover, tokens used by the threat actor had been digitally signed with a Microsoft Services Account (MSA) cryptographic key that Microsoft had issued in 2016. This particular MSA key should only have been able to sign tokens that worked in consumer OWA, not Enterprise Exchange Online. And this 2016 MSA key was originally intended to be retired in March of 2021, but its removal was delayed due to unforeseen challenges associated with hardening the consumer key systems." Whatever that means.

"This was the moment that Microsoft realized it had major overlapping problems. First, someone was using a Microsoft signing key to issue their own tokens; second, the 2016 MSA key in question was no longer supposed to be signing new tokens; and, third,

someone was using these consumer key-signed tokens to gain access to enterprise email accounts. According to Microsoft, this discovery triggered an all-hands-on-deck investigation by Microsoft that ran overnight" - oh my god, Leo, somebody lost some sleep over this.

Leo: Well, maybe not. Maybe they just ran it overnight, went to bed.

Steve: Oh, that's possible, you're right. "It ran overnight from June 26th into June 27th, 2023, focusing on the 2016 MSA key that had issued the token, as well as the access token itself. By the end of that day, Microsoft had high confidence that the threat actor was able to forge tokens using a stolen consumer signing key. Microsoft then escalated this intrusion internally, assigning it the highest urgency level and coordinating its investigation across multiple company teams. As a result, Microsoft developed 46 hypotheses to investigate, including some scenarios as wide-ranging as the adversary possessing a theoretical quantum computing capability to break public-key cryptography, or an insider who stole the key during its creation. Microsoft then assigned teams for each of the 46 hypotheses to try to prove how the theft occurred..."

Leo: How interesting.

Steve: Yeah.

Leo: What an approach, yeah.

Steve: "...prove it could no longer occur in the same way now; and to prove Microsoft would detect it if it happened again. Nine months after the discovery of the intrusion, Microsoft says that its investigation into these hypotheses remains ongoing." Another way of phrasing this would be "Microsoft still has no idea exactly how this happened." They know what, but not in detail exactly how.

The report continues: "Microsoft began notifying potentially impacted organizations and individuals on or about June 19th and July 4th, respectively. As detailed below, this effort had varying degrees of success. Ultimately, Microsoft determined that Storm-0558 used an acquired MSA consumer token signing key to forge tokens to access Microsoft Exchange Online accounts for 22 enterprise organizations, as well as 503 related personal accounts worldwide. Of the 503 personal accounts reported by Microsoft, at least 391 were in the U.S. and included those of former government officials, while others were linked to Western Europe, Asia-Pacific, Latin America, and Middle Eastern countries and associated victim organizations.

"Microsoft found no sign of an intrusion into its identity system and, as of the conclusion of this review, has not been able to determine how Storm-0558 had obtained the 2016 MSA key. It did find a flaw in the token validation logic used by Exchange Online that could allow a consumer key to access enterprise Exchange accounts if those Exchange accounts were not coded to reject a consumer key. By June 27th, 2023, Microsoft believed it had identified the technique used to access victim accounts and rapidly cleared related caching data in various downstream Microsoft systems to invalidate all credentials derived from the stolen key.

"Microsoft believed that this mitigation was effective, as it almost immediately observed Storm beginning to use phishing to try to regain access to the email boxes it had

previously compromised. However, by the conclusion of this review, Microsoft was still unable to demonstrate to the Board that it knew how Storm-0558 had obtained the 2016 MSA key."

So we've already seen that Microsoft has reversed its profit-motivated policy of charging its customers extra for security logging. We covered that earlier. And overall, a policy of charging anything extra in return for extra security seems similarly shortsighted. Security should be baked into all underlying aspects of any cloud deliverable. It should not be possible to "buy more security." It should be impossible to purchase less.

Only time will reveal what lessons Microsoft learns from all this. The lesson we must all learn is that when we transfer our corporate assets to the cloud, we're also transferring the responsibility for the security of those assets to the cloud services provider. So it's important to recognize that doing so does come with some risk, and that the fine print of the provider's contract holds them harmless, regardless of fault.

Leo: What a - what a world. Do you feel like Microsoft's, however, recognized the issue and made the changes they need to make, and we won't have to do this all over again?

Steve: Everybody loves a project, Leo.

Leo: It's like a committee. It's very similar to a committee. You know, you don't want to make a decision, appoint a committee. You don't want to really solve something, create a project. Maybe 47 of them, working overnight.

Steve: Yeah. Everybody loves a project. I mean, they need something to do. You know, you could argue Windows is done. You know, cloud, you know, Exchange is done. Everybody, I mean, what is the refrain we hear? Leave it the eff alone.

Leo: Yeah, yeah, yeah.

Steve: Just why - so Microsoft, yes, put a freeze, a formal freeze on features because they keep breaking it. Right? I mean, how many times have we said they're never going to get rid of the bugs because as many as they fix, they introduce new ones.

Leo: Yes.

Steve: With new features. Because they're constantly, you know, adding features. Stop with the features already. How about considering security a feature? What a concept.

Leo: Yeah. Good point. Good point. This is why you listen to Security Now!, right, every Tuesday, 1:30 Pacific, 4:30 Eastern, 2030 UTC, to get the deets, the update, the straight talk. That's the really most important thing, without fear or favor. So many other places on the 'Net you can get information, but there's always this undercurrent of, like, well, you know, who's paying for this? With Steve, you know. Steve says what he thinks and is extraordinarily trustworthy, and I love that about this show. I'm glad you're here. Glad you like it, too. Steve is at GRC.com. That's

where you'll get SpinRite. This is his only, by the way, the only thing he has, his only bread-and-butter winner, the world's finest maintenance and recovery and, what, speed-up utility for mass storage?

Steve: Performance recovery.

Leo: Performance recovery.

Steve: Data recovery and performance recovery.

Leo: Yeah.