Security Now! #974 - 05-14-24

Microsoft’s head in the Clouds

This week on Security Now!

What fascinating insights do we obtain from examining 3.4 million 4-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?

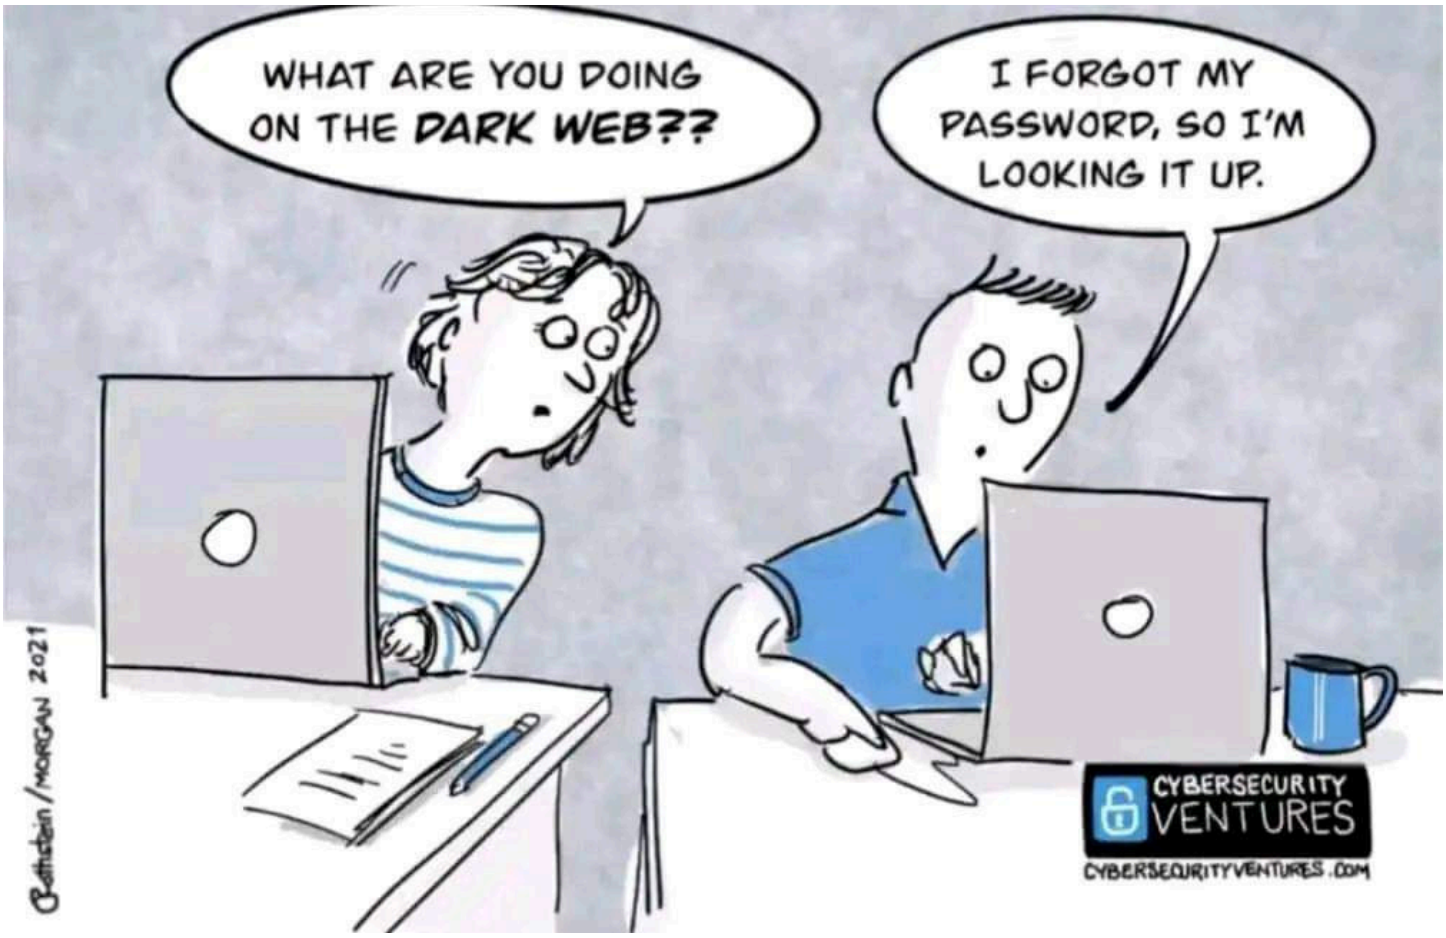

Security News

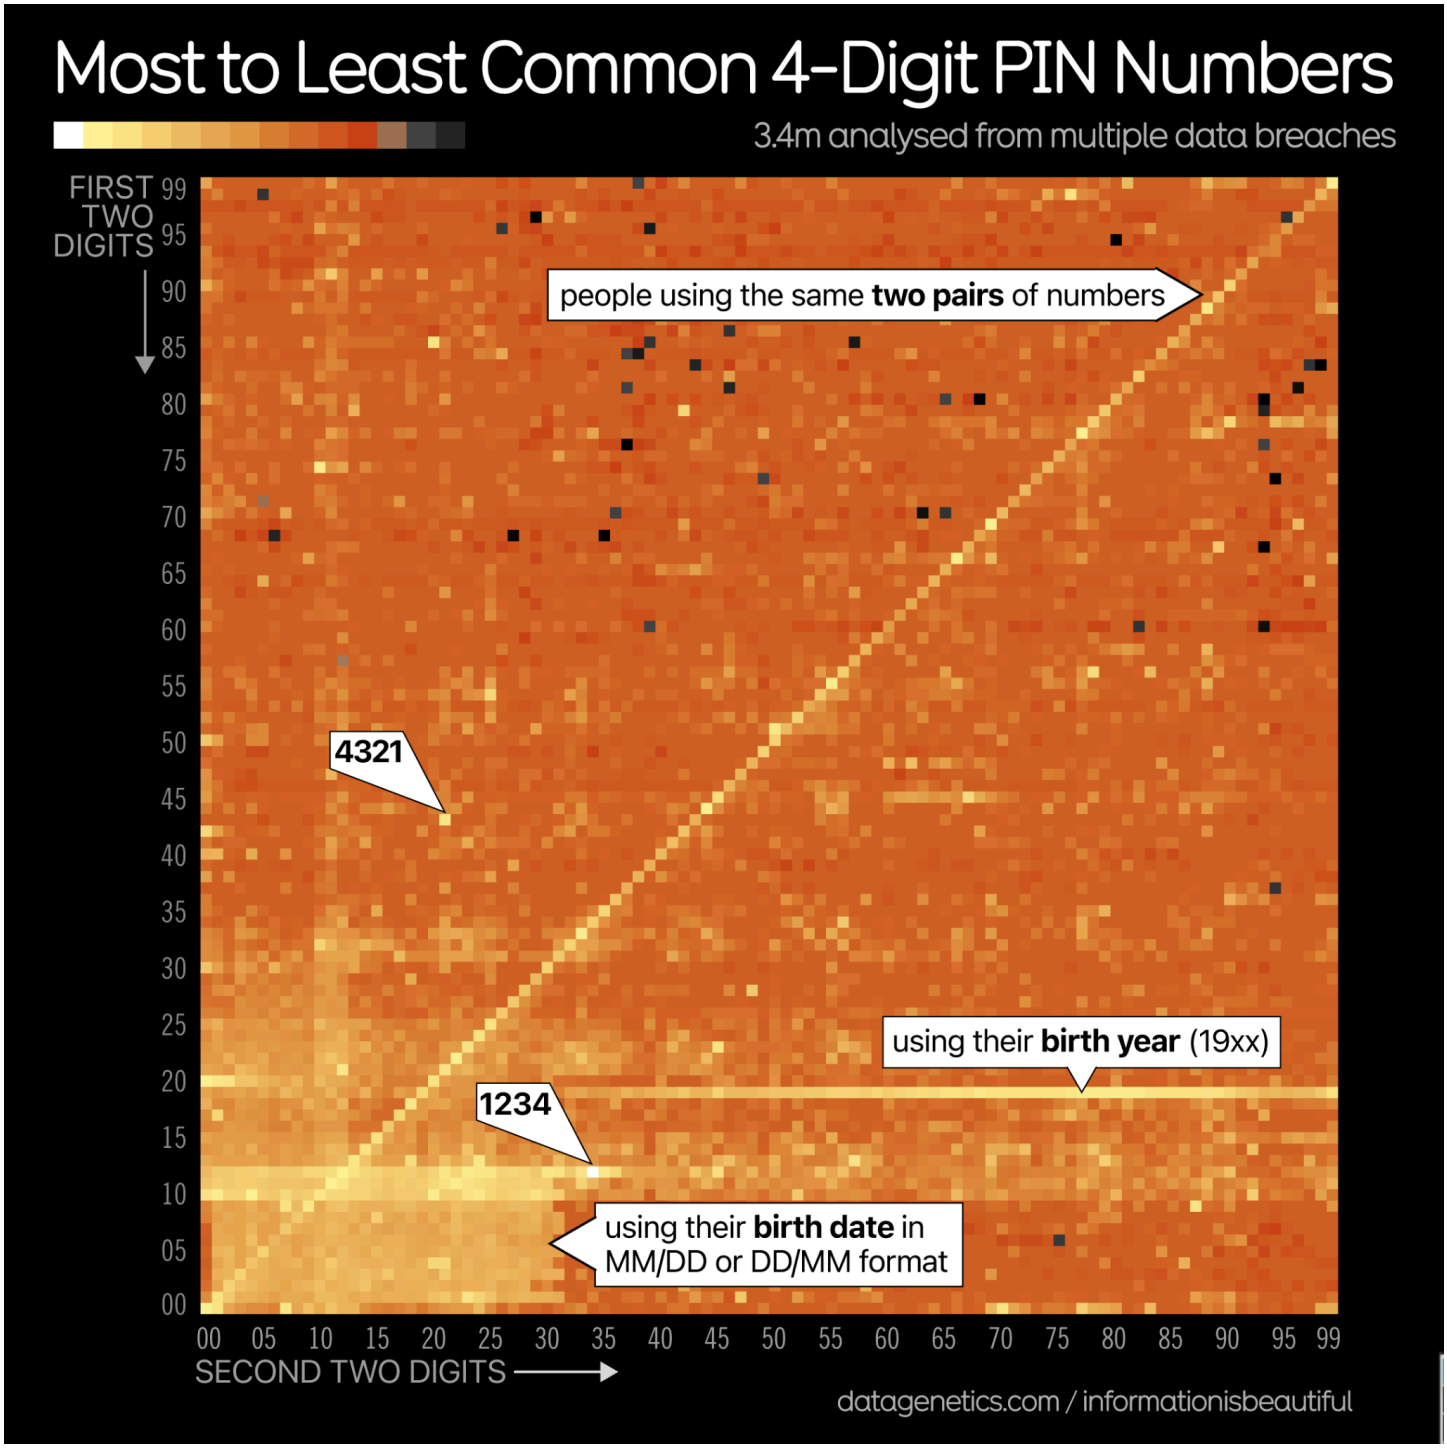

I wanted to share a wonderfully enlightening graphic chart with our listeners. Unfortunately, the terms "graphic chart" and "listeners" are at odds. So I'm going to note that this delightful chart is at the top of this week's show notes and I'll endeavor to explain to our listeners, who are commuting to work or moving the lawn as they listen to this, what it so wonderfully shows.

The chart takes 3.4 million 4-digit PINs recovered from, and disclosed in, multiple data breaches. A 4-digit PIN can have any value between 0000 and 9999. So there are ten thousand possible PINS… and this wonderful chart contains 10,000 little squares arranged in a flat two-dimensional

map having 100 rows and 100 columns. One way to think of this is that the first two digits of the PIN (00 through 99) specify one axis and the last pair of digits specify the other axis. So, every single possible 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, sets the brightness of its square on the chart. So what do we learn from this?

Probably the most prominent feature is a bright diagonal line running from the lower left corner of the chart, where both the first two and the last two digits are 00, to the chart's upper right corner where the first two and last two digits are both 99. The diagonal line, then, is formed by all of the intermediate squares when their first two and last two digits are identical. This bright diagonal is telling us that a great many people chose the same pair of digits twice. But there is also some variation in the brightness along the diagonal. Human nature being what it is, the PIN 6969 appears to be overrepresented relative to its neighbors.

Two other solitary bright spots would not not surprise anyone, they are the locations of the "1234" and "4321" PINs. Another prominent line is the 20th line 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. What's going here? If you guessed the people's birth year you would be correct. PINs often begin with 19 and they appear to be brightest around 1980.

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. Why? Because people used their month and day of birth with month running from 1 to 12 and day of month 1 through 31. There's a brighter horizontal stopping at 12 than the vertical stopping at 12. This indicates that most people chose the ordering with the month first and the day of month second.

Stepping back from it and looking at the overall illumination, there's a top to bottom brightness variation, but less of a left to right. So people are generally choosing 4-digit PINs with smaller first two digits but more randomly distributed last two.

And the final really interesting observation is that whereas most of the chart shows varying shades of illumination, there are around 40 cells that are black or nearly black. In other words, out of all 10,000 possible 4-digit PINs, there are around 40 of those that are significantly underrepresented. For some reason, for example, very few people have chosen 6806. So, if you're looking for a lesser chosen 4-digit PIN, there you go.

And as for the extremely low-entropy skewing observed in this dataset, get this: Just the 20 top most used PINs (out of those possible 10,000) account for 27% of all PINs observed in use. Those top 20 are: 1234 0000 7777 2000 2222 9999 5555 1122 8888 2001 1111 1212 1004 4444 6969 3333 6666 1313 4321 1010. I also had the bottom 20, but they were absolutely uninteresting since, not surprisingly, they all just looked like random numbers.

https://grc.sc/pin

Enhanced LORAN

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 localized, we also noted that space is, sadly, not necessarily a benign environment anymore. A piece of interesting listener feedback generated by this discussion last week led me to look at what's being done about this:

Shaun Merrigan / @shaun645D

Steve, regarding SN973 and GPS vulnerability: The US is testing an updated version of the LORAN system (which was shutdown in the 1980's) called eLORAN. I have 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. Shaun Merrigan

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 of Long Range Navigation and '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. I found an interesting summary on the site GPSWORLD. 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."* Here's what they explain:

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 (collectively referred to as "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, 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 for 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 $160M to convert the U.S. portion of the North American Loran-C service to enhanced Loran (eLoran). 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.

So first of all, it's great news that the US has been seriously looking into a backup technology. Since I think our listeners will find this interesting, I'll share a bit more 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. 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 had 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 meters. 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.*

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."*

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 US 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."* Why these fancy charts? Imagine for nautical navigation 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 at precisely the same time and the pulses radiate outward, spherically from each station at the speed of light, 186,000 miles per second. So the ship at sea will receive these two pulses, but it doesn't 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.

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 into 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 these projects funded here in the United States.

Closing The Loop

Jeff Urlwin / @jurlwin

Just listened to SN. Passkeys are even worse based upon website implementation. Some sites use a cookie to "know" they issued you a passkey. So even with 1Password which supports and synchronizes passkeys, I can't use the passkey from a different browser than originally set. CVS pharmacy is one with this bad implementation. Thanks for all your great shows.

... and ...

RG / @digitoxin

Regarding passkeys; for what it's worth, every website I have set up a passkey on has let me set up multiple passkeys, so I haven't been limited to a single ecosystem.

... and ...

Lachlan Hunt / @Lachy

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.

... and ...

Miguel Frade / @oMiguelFrade

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 have 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, Paypal only allows registering one YubiKey.

Regarding the question "Can the same key be applied to two different people?", the answer is yes, if we are 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.

I hope this information can be useful to other SN listeners. All the best, Miguel Frade

Last week's discussion of this generated significant feedback from our listeners. And the thing that stood out more than anything, was that everyone had a 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 for 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 spec.

I went over to Paypal to take a look, and their passkeys management page appears to 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. But this further demonstrates the mess we're currently working through. The fact that something stopped Miguel even though he has a perfectly secure authentication device – arguably more secure than the two smartphones Paypal does support – mostly just 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 it, and I think that's great. Really. I think that those of us in the industry who are grousing at the moment (Paul Thurrott went on a nice rant again about this last week) are doing so 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 "good enough" for now, with second-factor authentication adding additional protection where needed.

None of us can predict the future, and today's passkeys support remains disappointing. 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 is 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 with 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 login with this device ever again?" Who's not going to click "Do me!"?
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.

Microsoft's head in the Clouds

SCMagazine's headline reads: *"Sweeping cybersecurity improvements pledged by Microsoft"* And follows with: *"Numerous cybersecurity enhancements 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"*

SecurityWeek carried the headline: *"Microsoft Overhauls Cybersecurity Strategy After Scathing CSRB Report"* and follows with: Microsoft security chief Charlie Bell pledges significant reforms and a strategic shift to prioritize security above **all other product features**.

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 environment. 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 serious vulnerabilities which Microsoft will seemingly ignore for months and even years until that vulnerability is used to cause a highly public catastrophe. And only then will Microsoft apparently think "Huh... why does that exploit path have a familiar right to it." And I get it. 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 doesn't have to care. And we've seen example after example of Microsoft not doing anything it doesn't want to.

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: *"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 to scandal, in the wake of two major nation-state breaches of its core enterprise platforms, Microsoft is facing one of its most serious reputational crises.*
>
> *Adam Meyers, SVP at CrowdStrike, said: "It's certainly not the first time a nation-state adversary has breached Microsoft's cloud environments and after so many instances, empty promises of improved security are no longer enough."*
>
> ***To review:*** *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. 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.

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 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, Microsoft directs the vast majority of their security investments toward revenue generating roles instead of internal security roles."

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 1 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.

I should note for the record, that I don't have any feelings of schadenfreude here. 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 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. I'm amazed it's still afloat.

Early last month, the Department of Homeland Security's CSRB, the Cyber Safety Review Board, released their findings following a deep and detailed investigation into Microsoft's recent security breach troubles. Summarizing that report, the Cybersecurity Dive article 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 a G5 license from the company. Customers who failed to purchase the enhanced license, were not able to see the extensive logging capabilities that would have alerted them to a breach.

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."

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." 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 is 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 that 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 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 and future milestones as part of our efforts to strengthen all systems against attacks."

One of the problems with being transparent about what's being fixed is that the process of enumerating the improvements also enumerates just how bad things had been allowed to become. Bret said that since the launch of the company's Secure Future Initiative, the company has sped up related engineering work in several areas. He lists four:

- 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. It has also made multifactor authentication enforcement automatic across more than 1 million Entra ID tenants.

- More than 730,000 apps have been removed across production and corporate tenants that were either out of lifecycle or not meeting current SFI standards.

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

- The company's internal MFA 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.

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. 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 said: *"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. The full report is 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. 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 referred to a "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.

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.

The Board finds that this intrusion was preventable and should never have 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 (7 points):

1. *The cascade of Microsoft's avoidable errors that allowed this intrusion to succeed;*

2. *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;*

The State Department was the first victim to discover the intrusion when, on June 15, 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 MailItems-Accessed log was not accessible without that "premium" service.

3. *The Board's assessment of security practices at other cloud service providers, which maintained security controls that Microsoft did not;*

4. *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;*

5. *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; even though Microsoft acknowledged to the Board in November 2023 that its September 6, 2023 blog post about the root cause was inaccurate, it did not update that post until March 12, 2024, as the Board was concluding its review and only after the Board's repeated questioning about Microsoft's plans to issue a correction;*

6. *The Board's observation of a separate incident, disclosed by Microsoft in January 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*

7. *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.*

Throughout this review, 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.

To drive the rapid cultural change that is needed within Microsoft, 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 all instances, security risks should be fully and appropriately assessed and

addressed before new features are deployed.

Based on the lessons learned from its review and its fact-finding into prevailing security practices across the cloud services industry, the Board, in addition to the recommendations it makes to the President of the United States and Secretary of Homeland Security, also developed a series of broader recommendations for the community focused on improving the security of cloud identity and authentication across the government agencies responsible for driving better cybersecurity, cloud service providers, and their customers.

- *Cloud Service Provider Cybersecurity Practices: Cloud service providers should implement modern control mechanisms and baseline practices, informed by a rigorous threat model, across their digital identity and credential systems to substantially reduce the risk of system-level compromise.*

- *Audit Logging Norms: Cloud service providers should adopt a minimum standard for default audit logging in cloud services to enable the detection, prevention, and investigation of intrusions as a baseline and routine service offering without additional charge.*

- *Digital Identity Standards and Guidance: Cloud service providers should implement emerging digital identity standards to secure cloud services against prevailing threat vectors. Relevant standards bodies should refine, update, and incorporate these standards to address digital identity risks commonly exploited in the modern threat landscape.*

- *Cloud Service Provider Transparency: Cloud service providers should adopt incident and vulnerability disclosure practices to maximize transparency across and between their customers, stakeholders, and the United States government, even in the absence of a regulatory obligation to report.*

- *Victim Notification Processes: Cloud service providers should develop more effective victim notification and support mechanisms to drive information-sharing efforts and amplify pertinent information for investigating, remediating, and recovering from cybersecurity incidents.*

- *Security Standards and Compliance Frameworks: The United States government should update the Federal Risk Authorization Management Program and supporting frameworks and establish a process for conducting discretionary special reviews of the program's authorized Cloud Service Offerings following especially high-impact situations. The National Institute of Standards and Technology should also incorporate feedback about observed threats and incidents related to cloud provider security.*

One of the earliest breakthroughs in computing was the introduction of a concept that came to be called "timesharing". Back then, mainframe computers were incredibly expensive to purchase and operate. A single machine installation was planned years ahead. 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 that 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 the 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 cores can share that single powerful resource with many others.

This move to the cloud does not feel like 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-prem servers and managed service providers."*

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 ten thousand 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 have 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 then 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 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 the full 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, 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 what will all of 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 are 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 "Bill Yellow Taxi" log-reading intrusion detector decided that there was, indeed, a problem occurring. 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 when

it 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 16, 2023, 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 (OWA). 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 26, 2023 [10 days after the initial alert], Microsoft discovered that the threat actor had used OWA 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 2021, but its removal was delayed due to unforeseen challenges associated with hardening the consumer key systems.*
>
> *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 from June 26 into June 27, 2023, focusing on the 2016 MSA key that had issued the token as well as the access token itself. By the end of the 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 hypothesis to try to: prove how the theft occurred; prove it could no longer occur in the same way now; and to prove Microsoft would detect it if it happened today.*
>
> *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 19 and July 4, 2023, 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 European, Asia-Pacific (APAC), Latin American, 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 27, 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 begin to use phishing to try to gain 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.*

We've already seen that Microsoft has reversed its profit-motivated policy of charging its customers extra for security logging. 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 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.