

Security Now! #977 - 06-04-24

A Large Language Model in Every Pot

This week on Security Now!

When is a simpler application better than something complex? How did the first week of GRC's new email system go? Have you been Pwned? And if so, how worried should you be? What's the latest new supply-chain attack vector? What certificate authority just lost all their TLS server business? And remember that early messaging service ICQ? - whatever became of it? Finally, after I share a tip about a perfect science fiction movie, two pieces of listener feedback and one user's happiness over SpinRite, we're going to look at what a prominent security researcher learned after using Microsoft's Recall for ten days, and why I think Microsoft is willing to bet the farm and risk the dire warnings of the entire security community over this unasked for capability.

“But Officer...”



Last Week

"Tornado Notes"

I wanted to thank all those listeners who correctly recalled that the random notes DOS app we were trying to remember last week was "Tornado Notes." And it was not Phil Katz of PKZip fame, it was a guy named Jim Lewis of MicroLogic Corporation. When I first encountered Tornado Notes from a company named "MicroLogic Corporation" of Hackensack , New Jersey, I wondered "why is that name so familiar?" And it turned out that was because the same guy had created one of the most useful sets of 8.5" x 11", double-sided plastic single-sheet processor instruction reference cards the world had ever encountered:



Upon the event of my death, my plan is for cremation, after first having whatever organs may still be functioning and useful to anyone, removed. But IF my plan were burial, I would want these processor instruction reference cards buried alongside me. I can't begin to express how important they were back when I was writing assembly code, first for Apple's and later Atari's 6502-based machines, and then when I started out programming the IBM PC to create "Flicker Free", my first commercial app for the PC. Someone over on Reddit commented that it was a good thing these were 100% plastic or he would have worn his out; they were indispensable. If anyone remembers these and is feeling nostalgic, searching for "Micro Logic Instruction Reference Cards" will turn up entirely legible PDFs of both sides of these indispensable treasures. And I have links to all four of them in the show notes for the 6502, Z80, 8088/86 & the 68000:

[6502 PDF](#)

[Z80 PDF](#)

[8088/86 PDF](#)

[68000 PDF](#)

Tornado Notes was for DOS and was utterly unique. When Windows happened, Jim tried to recreate the success of Tornado Notes with a product he named "Info Select". But Info Select was the victim of its own featuritis. The sublime beauty of Tornado Notes was that it was so simple. It did exactly and only one thing perfectly and instantly. It began as a massively overwhelming disorganized pile of rectangular notes. Then, as you typed successive characters of a string, all those notes that did not contain the substring that had been entered thus far disappeared. So you got this very satisfying, almost animated, real-time winnowing of your entire pile until you could see the note you knew was there somewhere. And notice that you also saw ALL the notes that also contained the same substring, which was often surprisingly useful.

Unfortunately, Jim did not understand that Tornado Notes succeeded due to the constraints imposed upon it by its DOS environment. So when he created its successor, "Info Select for Windows", he gave it hierarchies and categories and menus and formatted printing and everything else you can imagine that Windows made possible. I think there was even a kitchen sink tucked in there somewhere. We wanted the same thing for Windows that we had for DOS, but what we got was a monstrosity that required all manner of configuration and thought. Yes, it could do so much more than Tornado Notes. But the very thing that was so beautiful about Tornado Notes was everything it did not do, which, as it turned out in retrospect, made the one little thing it did do, so compelling and useful!

I'm mentioning this because there's a larger lesson here. One of the things the original designers of UNIX got exactly right was the idea of creating many simple commands that took some input, did something to it, then produced some output. To that, you add the simple ability to interconnect these individual small building blocks into a chain by piping the output of one into the input of another, and you're able to interactively create and assemble a much more complex ad hoc function.

And, Leo, while I'm not a LISP programmer, I have the sense that the same sort of approach can be used there, where you incrementally build up a much more complex solution that's assembled from many smaller functions. The point I hope to make here is that "more" is not always "better." And, sadly, this is not a lesson that the people who design the remote controls for A/V equipment appear to have ever learned.

Email @ GRC

I also wanted to follow up on last week's announcement of GRC's new email system, which has been a resounding success. If you missed last week's episode and don't know about it yet, you can go to our old grc.com/feedback page, which explains a bit about the nature of "web form spam" and contains a pointer to grc.com/mail.

The only post-announcement glitch we encountered was from users, mostly of gmail, but also a few other ISPs, who are using their own domains backed by those services. Since the email they send comes from that underlying service, like gmail, rather than from their domain alias, and since the incoming filter that's in front of "securitynow@grc.com" looks to see whether the sender is known to us, listeners need to register their underlying gmail account at GRC, not their aliased account which is the one shown in the From: header of their email. After I realized what was going on, I added a little note to that effect to the email registration page and registrations have been trouble free ever since.

Also, using an anonymizing email service address won't work. I received an email from a listener who was using the "SimpleLogin" email anonymizer service by Proton. When that listener sent email to GRC, the sender's email was a long one-time 54-character random account name in front of the @simplelogin.com domain name. So GRC's filter had never seen that before and never will again.

But to make a long story short, our listeners LOVE this simple solution: Register one time, optionally subscribe to whatever announcement lists, if any, you may wish, and from then on you can simply send email to "securitynow@grc.com". I've been overwhelmed with notes of thanks and congratulations from listeners I've never heard from before who were never going to sign up for Twitter just to maybe send me a note. Twitter is about so much more than that. It's about building a community and a following. I had just been using it as a point-to-point instant messaging service... which is exactly what email is.

Needless to say, I will **never** share anyone's email address publicly, and anyone requesting anonymity for their name, receives it. One of the nice things about GRC's blessedly retired web form was that it solicited our listener's location. It was nice being able to include that when sharing feedback, since it made it feel a bit more personal. So if you happen to think of it, let me know where you're writing from.

My current work now is on automating and catching realtime email bounces, so I can immediately inform someone when GRC can detect that it was unable to successfully deliver their authentication loop email. Once that's in place, I'll stick my toe in the water to begin actually sending email in today's spam-conscious climate, and we'll ramp up from there.

So, thank you, everyone, for your support. Your support and interest is the reason I became convinced that we need to keep this going past 999. And here we are, already, at 977 with our 20th birthday coming up in August!

Have I Been Pwned?

While I was writing the note above I received an email alert from Troy Hunt's "Have I Been Pwned?" email breach monitoring service. The email's Subject was: "16 emails on grc.com have been pwned in the **Telegram Combolists** data breach" The breach occurred one week ago on May 28th, 2024. In the breached data, 361,468,099 email accounts were found. And HIBP sent this email because 16 of those belonged to GRC.COM. The description of the breach is:

In May 2024, 2 Billion rows of data with 361 Million unique email addresses were collated from malicious Telegram channels. The data contained 122 Gigabytes across 1700 files with email addresses, usernames, passwords and in many cases, the website they were entered into. The data appears to have been sourced from a combination of existing combolists and info stealer malware.

So, naturally, I went over to see whether any of those 16 addresses which HIBP reported were of concern. None were. The only two that were **ever** valid were greg@grc.com and offices@grc.com, neither of which we have used for decades. I once watched a spammer's server connect to GRC's email server and just run down a list of firstnames, hoping to get lucky. Immediately after that we retired our original and oh-so-very-innocent use of our first names for email.

The wonderful open source windows email server I've been using for years is known as hMailServer. Anyone looking for an utterly solid, feature packed, no nonsense, free, windows-hosted email server should look no further. It's another of those rare software creations that has no bugs. The only time it's been updated for years is to keep up with improvements in the OpenSSL library which it uses to make its TLS client and server connections. And, in fact, I updated it last week after many years of trouble free service to obtain support for TLS v1.3.

hMailServer has a dynamic blocklist feature that will block, for a configurable period of time, any remote server, by IP address, that attempts to deliver email to any nonexistent address at GRC. I just checked the server. I currently have the blocklist expiration set for 2 hours and at the moment I checked, 473 individual IP addresses were currently being blocked. So within the previous 120 minutes, 473 different spamming SMTP servers had connected to GRC and attempted to send spam, not to any actual valid email addresses here, but just to throw crap at the wall hoping to get lucky.

GRC has been around for a long time. The domain is well known. But we're certainly not particularly high profile. It so saddens me to see what a sewer our beloved Internet has become. I'm unsure what it teaches us about humanity, but I am pretty sure I don't want to know.

The trifecta of the Internet being anonymous, global and free enables every last miscreant on Earth to attempt to have their way with everyone else. Fortunately, the rest of us are far from powerless and we have this podcast to help us stay ahead of the tidal wave of incoming crap that's out there pounding on the door, trying to get in. We're not going to let any of that in.

Security News

A new "supply chain" attack vector

And speaking of what a sad mess the greater Internet has become, and of not letting any of that mess into our lives, one of our listeners, Terence Kam, pointed me to a recent piece in BleepingComputer titled *"Cybercriminals pose as 'helpful' Stack Overflow users to push malware"*.

For those who have never encountered it, Stack Overflow is a forum community of developers of widely ranging skill. It's essentially a place where coders can help one another. When I've been struggling with a programming problem, such as when I was working to get server-side on-the-fly code signing to work remotely with a certificate stored in an HSM, the Stack Overflow site would often be listed among Google's search results. And I'm a member there, since I've enjoyed answering questions and giving back when I can. So BleepingComputer writes:

Cybercriminals are abusing Stack Overflow in an interesting approach to spreading malware—answering users' questions by promoting a malicious PyPi package that installs Windows information-stealing malware.

Sonatype researcher Ax Sharma (and a writer at BleepingComputer) discovered this new PyPi package is part of a previously known 'Cool package' campaign, named after a string in the package's metadata, that targeted Windows users last year.

This PyPi package is named 'pytoileur' and was uploaded by threat actors to the PyPi repository over the weekend, claiming to be an API management tool. Malicious packages like this are usually promoted using names similar to other popular packages, a process known as Typo-squatting. However, with this package, the threat actors took a more novel approach by answering questions on Stack Overflow and promoting the package as a solution.

As Stack Overflow is a widely used platform for developers of all skill sets to ask and answer questions, it provides a perfect environment to spread malware disguised as programming interfaces and libraries.

Sonatype's Ax Sharma said in their report: "We further noticed that a StackOverflow account [with the nonsense name] "EstAYA G" created roughly 2 days ago, is now exploiting the platform's community members who are seeking debugging. It's directing them to install this malicious package as a "solution" to their issue even though the "solution" is unrelated to the questions posted by developers."

In this case, the pytoileur package contains a 'setup.py' file that pads a base64 encoded command to execute with ample space-padding, so it's hidden unless you enable word wrap in your IDE or text file editor. [In other words, it deliberately spaces the code far to the right so that it winds up off-screen.] When deobfuscated, this command will download an executable named 'runtime.exe' from a remote site and execute it.

This executable is a Python program converted into an .exe that acts as an information-stealing malware to harvest cookies, passwords, browser history, credit cards, and other data from web browsers. It also appears to search through documents for specific phrases and, if found, steal the data as well.

While malicious PyPi packages and information-stealers are nothing new, the cybercriminals' strategy to pose as helpful contributors on Stack Overflow is an interesting approach as it allows them to exploit the site's trust and authority within the coding community.

This approach serves as a reminder of the constantly changing tactics of cybercriminals and, unfortunately, illustrates why you can never blindly trust what someone shares online.

Instead, developers must verify the source of all packages they add to their projects, and even if it feels trustworthy, check the code (with word wrap enabled) for unusual or obfuscated commands that will be executed:

```

1  from setuptools import setup, find_packages
2  from setuptools.command.install import install
3  from pathlib import Path
4  import os
5
6
7  VERSION = '1.0.2'
8  DESCRIPTION = 'Cool package.'
9  this_directory = Path(__file__).parent
10 long_description = (this_directory / "README.md").read_text()
11
12
13 class InstallCommand(install):
14
15     def run(self):
16         try:
17             print("")
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100

```

I have a snapshot of the original code in the show notes. If you ever see a big blob of base64 encoded text anywhere in a package that's supposed to be human readable, like Python source code should be, run away.

And I'll just note that before the end of today's podcast, the security researcher Kevin Beaumont is going to show us, despite Microsoft's claims to the contrary, that the database underlying Microsoft's new Recall system can, in fact, be exfiltrated remotely, does not require SYSTEM privilege and can be accessed by any other user of the same machine. That means that Recall's SQLite database is 100% vulnerable to exactly this sort of info-stealer malware. More on this in a bit.

Another CA in the DogHouse

Google has announced that it will be removing its trust of all new TLS certificates issued by the Austrian certificate authority, GlobalTrust. Rather than yanking GlobalTrust's root certificate, which would invalidate all previously-issued GlobalTrust certs, Chrome will be using a recently added new feature that allows it to manage certificate trust based upon their issue dates. So Chrome will not be trusting any new certificates issued by GlobalTrust after the end of this month, June 30th.

Through the nearly 20 years of this podcast we've seen a range of misbehavior on the part of those who have been given the privilege of essentially printing money. Certificate authorities charge their customers hundreds of dollars in return for encrypting a hash of a small block of bits the customer presents. But in return for this money-printing privilege, the CA must abide by a significant code of conduct. When that code is broken, and only after bending over backwards with more than ample warnings, the industry can and has summarily withdrawn its trust from the signatures of those CAs on the grounds that if the CA cannot be trusted, neither can anything they have signed.

In this case GlobalTrust has established a multi-year history of misconduct and they have lost the trust of the industry. Google will be enforcing a ban retroactively on all Chrome versions down to v124 and the other browser makers have not announced a similar decision though Mozilla appears to also be aware and concerned. Still, since no customer would purchase a certificate for a web server which anyone visiting with Chrome would be unable to connect to securely, this immediately puts GlobalTrust out of the web server certificate business. In other words, whether or not Apple and Mozilla choose to follow, GlobalTrust is done for now.

ICQ to shutter its service

Those of us who have been around since the dawn of the Internet will likely remember the first successful desktop instant messaging app known as ICQ – which was meant to be short for “I seek you.” The system was originally developed back in 1996 by an Israeli company named Mirabilis. Two years later it was acquired by AOL in 1998, and then by the Russian Mail.Ru Group in 2010. It had a neat kind of funky flower petal logo and I've wondered what became of it.

At its peak around 2001 it had more than 100 million accounts registered. And nine years later when AOL sold it to Mail.Ru it had around 42 million daily users. And it has been puttering along in the background ever since. Two years ago it was at around 11 million monthly users. And finally, the reason the subject came up is that a week and a half ago, on May 24, 2024, the website of ICQ.com announced that the service would be shut down about three weeks from now, on June 26, 2024. So a pretty good 28-year run for an instant messaging service that was largely passed by when Smartphones and other major social networking services got into the game.

Sci-Fi

“Déjà vu”

My wife recently agreed to join me in watching one of my favorite science fiction movies of all time. We know I'm a pushover for science fiction. But unfortunately, far more horrible science

fiction movies have been made, than good ones; and even more rare is the perfect science fiction movie. So we settled down to watch "Déjà vu" which stars Denzel Washington, Val Kilmer and some other recognizable actors from Hollywood's inventory. As I was watching it, for maybe the 4th time, I kept thinking over and over, as this perfectly and often leisurely paced two hour movie unfolded scene by scene, and everything was happening exactly the way it should, that I was sitting here watching one of the all too rare perfect movies. This movie offers convincing acting that is not distracting, a brand new and perfect concept, a perfect script, and a plot that's both surprising and where what happens is better than someone steeped in science could have ever hoped for. The writers enlisted the help of Brian Greene, a Cornell and Columbia University physicist, to get the science right – and boy did they! That's part of what's so gratifying about this movie. It's not a new movie, it was released 18 years ago, back in 2006. But it stands up and it still feels 100% contemporary.

I realized that since this podcast is closing in on its 20th birthday, every time I've seen this movie we've done this podcast a few days later, yet somehow I've never thought to mention it. I searched our transcripts and there was no mention of it. So that's my 'bad' and that's fixed now. I know quite well that not everyone's taste is the same. Not everyone will feel as I do about this. But if you don't already know this movie and you've been looking for something to watch, the movie Déjà vu gets my highest recommendation.

Closing The Loop

Jeff Price

Leo touched on this, but Fastmail allows you to create these unique random email addresses. What most people forget is Apple lets you create theses as well. They call it Hide My Email.

I wanted to share Jeff's note since I have the feeling email aliasing services are going to become increasingly popular as websites turn to collecting and sharing whatever they can about their visitors as a means of increasing their advertising revenue.

Kirk Sexton

Hi Steve, Great work on the new email system! I never miss a show – I listen on my morning runs and in the car on my way to work. Sometimes I have to run a little further or sit in my car for a few minutes longer after arriving so I don't interrupt a point before hitting pause.

I may have missed this point, but I don't recall hearing anything about those users who sync their accounts on Microsoft OneDrive, or for that matter use other cloud based backup services. Backing up files is one thing; it would be expected that anything committed to local storage will be backed up to the subscribed cloud storage. However, temporary information that is used just for the moment, will now be stored locally (think passwords, credit cards, or other sensitive information) within the screen grabs. Microsoft has said that it will only be stored locally, but what about cloud-syncing with OneDrive or other services. I see it as the problem just mushrooming into multiple attack vectors. Am I missing something?

To 999 and beyond! All the best, Kirk Sexton

You raise a great point. We're about to spend the rest of the podcast looking at what one security researcher found and also about what may be Microsoft's significantly greater plan. But everything we know now suggests that the Recall data are just SQLite files stored under the user's standard AppData directory in a new "CoreAIPlatform" folder. Microsoft has indicated that BitLocker will be used to encrypt the data at rest. But online backups are made of live unencrypted data so that it can be later retrieved and there's nothing we know so far that would prevent anything that was backing up a user's machine from also backing up their machine's Recall history.

SpinRite

The first week of listener feedback email into GRC's new system was, let's just say, intense with many listeners wanting to say Hi! and to express happiness that there was now a way to send me thoughts without engaging social media. At the moment I am way behind and far from caught up, but I figured I'd share one piece of feedback that's primarily about a SpinRite owner's experience with SpinRite 6.0 and then, finally, with v6.1.

Mark Jones sent email with the subject: "Wow! SpinRite 6.1 is amazing", writing:

Dear Steve, Long time listener, occasional source of feedback (I was @mjphd on twitter). I'm so happy to be using email. I only kept my X account for Security Now feedback.

I've listened to you discuss both the speed of 6.1 and the magic it does on an SSD. Ever the experimentalist, I thought I would put it through its paces. I have two drives, a 1 TB spinner and a 250 GB SSD that seemed to have slowed. The results are nothing short of remarkable on both drives. In only 4 hours, the 1 TB was rejuvenated. That would have taken days using 6.0. The boot into Windows 10 is now seconds instead of minutes and the random slowdowns that were plaguing the system are gone.

The real miracle was on the SSD. The new drive test showed I was a 19 MB/s at the front and middle, with 80 MB/s at the end. The whole drive is now over 546 MB/s after a level 3 scan. Saying computer performance has returned feels inadequate. It is mindblowingly fast compared to yesterday.

Truly amazing! Thanks for the great work and I'm happy there will be a future past 999.

*Regards
Mark Jones*

A Large Language Model in Every Pot

A data-driven theory about Microsoft's future plans for Recall emerged after I read a recent posting by the well known and well informed security researcher, Kevin Beaumont. Since last week's episode, which I titled "The 50 Gigabyte Privacy Bomb", Kevin whom we often quote and refer to, has again weighed in on Microsoft's new Recall facility. His first posting on the subject, which he made on May 21st, the day following Microsoft's announcement, was titled "*How the new Microsoft Recall feature fundamentally undermines Windows security.*" As a mature, seasoned and experienced security researcher, his immediate "what could possibly go wrong?" reaction to the idea of having Windows continually recording and storing our PC's screens echoes my own. It's immediately obvious to anyone who's been around the block a few times that this is, indeed, a 50 gigabyte privacy bomb. What wasn't clear to me until just yesterday, was **why** Microsoft may be doing this and what they probably have planned for the future.

Ever since his immediate posting in reaction to the announcement of Recall, Kevin's been playing with Recall. After reading what Kevin wrote, a lightbulb went off for me. So I'm first going to share Kevin's follow-up piece which further describes Recall. Then I'll share what I think it really means. Kevin titled his follow-up piece, which he posted four days ago, after spending a week and a half with Recall: "*Stealing everything you've ever typed or viewed on your own Windows PC is now possible with two lines of code — inside the Copilot+ Recall disaster.*" Before switching into Q&A mode, which he does later, Kevin began his newly informed discussions of Recall by writing:

I wrote a piece recently about Copilot+ Recall, a new Microsoft Windows 11 feature which — in the words of Microsoft CEO Satya Nadella — takes "screenshots" of your PC constantly, and makes it into an instantly searchable database of everything you've ever seen. As he says, it is a photographic memory of your PC life.

I got a hold of the Copilot+ software and got it working on a system without an NPU about a week ago, and I've been exploring how this works in practice, so we'll have a look into that shortly. First, I want to look at how this feature was received as I think it is important to understand the context.

The overwhelmingly negative reaction has probably taken Microsoft leadership by surprise. For almost everybody else, it won't have. This was like watching Microsoft become an Apple Mac marketing department. At a surface level, it is great if you are a manager at a company with too much to do and too little time as you can instantly search what you were doing about a subject a month ago.

In practice, that audience's needs are a very small (tiny, in fact) portion of Windows overall user base — and frankly, talking about screenshotting the things people in the real world, not executive world, are doing, is basically like punching customers in the face. The echo chamber effect inside Microsoft is real here, and oh boy... just oh boy. It's a rare misfire, I think.

*I think Recall is an interesting, entirely optional feature with a niche initial user base that would require incredibly careful communication, cybersecurity, engineering and implementation. **Copilot+ Recall does not have any of these.***

The work has clearly not been done to properly package it together.

A lot of Windows users just want their PCs so they can play games, watch porn, and live their lives as human beings who make mistakes that they don't always want to remember, and the idea other people with access to the device could see a photographic memory is.. very scary to a great many people on a deeply personal level. Windows is a personal experience. This shatters that belief.

I thought Kevin's take on this was interesting. His observation that Microsoft appears to be oblivious to the fact that not all users of PCs are even close to being the same. That a manager in a corporate environment might indeed find it useful to be able to look a month back for some specific subject, but that for the common user – the rest of us – the idea that our machines are are watching and recording everything we do, even if it would only be for our own later access, is mostly just creepy.

We don't know the future. We don't know what's going to happen a month or two from now. But Recall would make what's happening on our machines now, available to that unknown future. Anyway, Kevin finishes his lead-in by writing:

I think they are probably going to set fire to the entire Copilot brand due to how poorly this has been implemented and rolled out. It's an act of self harm at Microsoft in the name of AI, and by proxy... real customer harm. More importantly, as I pointed out at the time, this fundamentally breaks the promise of security in Windows.

I'd now like to detail why.

Strap in — this is crazy. I'm going to structure this as a Q&A with myself now, sourced from comments online, as it's really interesting seeing how some people handwave the issues away.

Okay, so we go with Kevin's Q&A format disclosure of what he found:

Q. *The data is processed entirely locally on your laptop, right?*

A. *Yes! They made some smart decisions here, there's a whole subsystem of Azure AI etc code that processes on the device.*

Q. *Cool, so hackers and malware can't access it, right?*

A. *No, they can.*

Q. *But it's encrypted.*

A. *When you're logged into a PC and run software, things are decrypted for you. Encryption at rest only helps if somebody comes to your house and physically steals your laptop — that is not what criminal hackers do.*

For example, InfoStealer trojans, which automatically steal usernames and passwords, have

been a major problem for well over a decade — now these can be easily modified to support Recall.

Q. *But the BBC said data cannot be accessed remotely by hackers.*

A. *They were quoting Microsoft, but this is wrong. Data can be accessed remotely. This is what the journalist was told for some reason:*



Q. *Microsoft say only that user can access the data.*

A. *This is not true, I can demonstrate another user account on the same device accessing the database.*

Q. *So how does it work?*

A. *Every few seconds, screenshots are taken. These are automatically OCR'd by Azure AI, running on your device, and written into an SQLite database in the user's folder. This database file has a record of everything you've ever viewed on your PC in plain text. OCR is a process of looking at an image, and extracting the letters.*

Q. *What does the database look like?*

A: [And Kevin shows some screen shots like we saw last week.]

Q. *How do you obtain the database files?*

A. *They're just files in AppData, in the new CoreAIPlatform folder.*

Q. *But it's highly encrypted and nobody can access them, right?!*

A. *Here's a few second video of two Microsoft engineers accessing the folder.*

[Kevin then quotes an earlier Mastodon post of his at cyberplace.social where he notes that the Risky Business episode on Recall is good, but with one small correction – Recall does not need SYSTEM rights. He notes that since it's just a SQLite database, it's trivial to access.

And he finishes by saying: "I'm not being hyperbolic when I say this is the dumbest cybersecurity move in a decade. Good luck to my parents safely using their PC.]

Q. ...But, normal users don't run as admins!

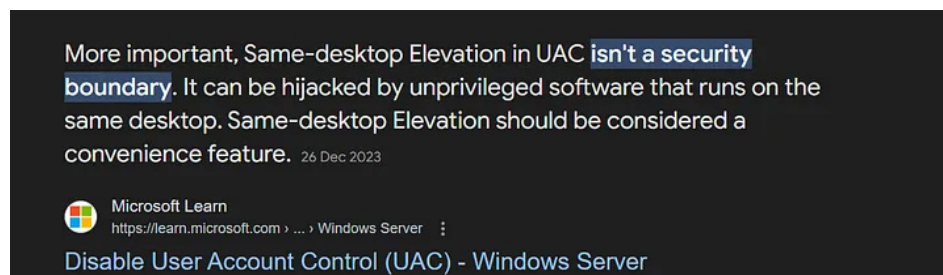
A. According to Microsoft's own website, in their Recall rollout page, they do:



In fact, you don't even need to be an admin to read the database — more on that in a later blog.

Q. But a UAC prompt appeared in that video, that's a security boundary.

A. According to Microsoft's own website (and MSRC), UAC is not a security boundary:



Q. So... where is the security here?

A. They have tried to do a bunch of things but none of it actually works properly in the real world due to gaps you can fly a plane through.

Q. Does it automatically not screenshot and OCR things like financial information?

A. No:

Q. How large is the database?

A. It compresses well, several days working is around ~90kb. You can exfiltrate several months of documents and key presses in the space of a few seconds with an average broadband connection.

Q. How fast is search?

A. On device is really fast.

Q. Have you exfiltrated your own Recall database?

A. Yes. I have automated exfiltration, and made a website where you can upload a database and instantly search it. I am deliberately holding back technical details until Microsoft ship the feature as I want to give them time to do something.

I actually have a whole bunch of things to show and think the wider cyber community will have so much fun with this once it's generally available.. but I also think that's really sad, as real world harm will ensue.

Q. *What kind of things are in the database?*

A. *Everything a user has ever seen, organized by application. Every bit of text the user has seen, with some minor exceptions (e.g. Microsoft Edge InPrivate mode is excluded, but Google Chrome isn't). Every user interaction, e.g. minimizing a window. There is an API for user activity, and third party apps can plug in to enrich data and also view store data.*

It also stores all websites you visit, even if third party.

Q. *If I delete an email/WhatsApp/Signal/Teams message, is it deleted from Recall?*

A. *No, it stays in the database indefinitely.*

Q. *Are auto deleting messages in messaging apps removed from Recall?*

A. *No, they're scraped by Recall and available.*

Q. *But if a hacker gains access to run code on your PC, it's already game over!*

A. *If you run something like an info stealer, at present they will automatically scrape things like credential stores. At scale, hackers scrape rather than touch every victim (because there are so many) and resell them in online marketplaces.*

Recall enables threat actors to automate scraping everything you've ever looked at within seconds.

While testing this with an off the shelf infostealer, I used Microsoft Defender for Endpoint — which detected the off the shelf infostealer — but by the time the automated remediation kicked in (which took over ten minutes) my Recall data was already long gone.

Q. *Does this enable mass data breaches of website?*

A. *Yes. The next time you see a major data breach where customer data is clearly visible in the breach, you're going to presume the company who processes the data is at fault, right?*

But if people have used a Windows device with Recall to access the service/app/whatever, hackers can see everything and assemble data dumps without the company who runs the service even being aware. The data is already consistently structured in the Recall database for attackers.

So prepare for AI powered super breaches. Currently credential marketplaces exist where you can buy stolen passwords — soon, you will be able to buy stolen customer data from insurance companies etc because all code required to do this has been pre-installed and enabled on Windows by Microsoft.

Q. *Did Microsoft mislead the BBC about the security of Copilot?*

A. Yes.

Q. *Have Microsoft mislead customers about the security of Copilot?*

A. Yes. For example, they describe it as an optional experience — but it is enabled by default and people can optionally disable it. That's wordsmithing.

Microsoft's CEO referred to "screenshots" in an interview about the product, but the product itself only refers to "snapshots" — a snapshot is actually a screenshot. It's again wordsmithing for whatever reason. Microsoft just need to be super clear about what this is, so customers can make an informed choice.

And need I note that "the tyranny of the default" will be at work, here. **We know** that whatever is the default setting is what 99.99% of all Windows users will have active.

Q. *Recall only applies to 1 hardware device!*

A. That is not true. There are currently 10 Copilot+ devices available to order right now from every major manufacturer <https://www.microsoft.com/en-gb/windows/copilot-plus-pcs#shop> Additionally, Microsoft's website says they are working on support for AMD and Intel chipsets. Recall is coming to Windows 11.

Q. *How do I disable Recall?*

A. In initial device setup for compatible Copilot+ devices out of the box, you have to click through options to disable Recall. In enterprise, you have to turn off Recall as it is enabled by default.

Q. *What are the privacy implications? Isn't this against GDPR?*

A. I am not a privacy person or a legal person. I will say that privacy people I've talked to are extremely worried about the impacts on households in domestic abuse situations and such. Obviously, from a corporate point of view organizations should absolutely consider the risk of processing customer data like this — Microsoft won't be held responsible as the data processor, as it is done at the edge on your devices — you are responsible here.

Q. *Are Microsoft a big, evil company?*

A. No, that's insanely reductive. They're super smart people, and sometimes super smart people make mistakes. What matters is what they do with knowledge of mistakes.

Q. *Aren't you the former employee who hates Microsoft?*

A. No. I just wrote a blog this month praising them: Breaking down Microsoft's pivot to placing cybersecurity as a top priority. My thoughts on Microsoft's last chance saloon moment on security.

Q. *Is this really as harmful as you think?*

A. *Go to your parent's house, your grandparent's house, etc and look at their Windows PC, look at the installed software in the past year, and try to use the device. Run some antivirus scans. There's no way this implementation does not end in tears — there's a reason there's a trillion dollar security industry, and that most problems revolve around malware and endpoints.*

Q. *What should Microsoft do?*

A. *In my opinion — they should recall Recall and rework it to be the feature it deserves to be, delivered at a later date. They also need to review the internal decision making that led to this situation, as this kind of thing should not happen.*

Earlier this month, Microsoft's CEO emailed all their staff saying "If you're faced with the tradeoff between security and another priority, your answer is clear: Do security."

We will find out if he was serious about that email. They need to eat some humble pie and just take the hit now, or risk customer trust in their Copilot and security brands.

Frankly, few if any customers are going to cry about Recall not being immediately available — but they are absolutely going to be seriously concerned if Microsoft's reaction is to do nothing, ship the product, slightly tinker or try to wordsmith around the problem in the media.

Okay. So what do we learn from this?

We now know that Microsoft currently plans to enable this whole PC history recording by default. They also know that unless Windows ships with it enabled and running, no one will use it. So they want to blow everyone's mind by AI-enabling Windows PCs somehow, and this is what they've come up with. I doubt there's an informed security-minded technologist anywhere who doesn't think this is a very bad idea. Yet, until we learn otherwise, this is exactly what Microsoft intends to do.

I have some personal experience with endeavoring, and failing, to get Microsoft to change its plans. Before their release of Windows XP, which grew out of Windows 2000, I tried to keep Microsoft from shipping XP with the totally unnecessary access to raw sockets available to the operating system's client software. They ignored me until the MS Blast worm would have taken them off the Internet if it had not been targeted at the wrong domain. After that near death brush with being attacked by an entirely unnecessary feature of their own operating system, XP's service pack 3 removed unprivileged access to raw sockets... and no one cared. The fact that no one cared demonstrated that the unnecessary feature should never have been present in a consumer OS. Raw sockets never came back because they just beg to be abused.

I learned my lesson from that experience. I have no interest in lobbying Microsoft to change its behavior. Microsoft is like Godzilla – it does whatever it wants to do – all anyone can do is stay out of its way. But what's so odd about this moment where we find ourselves, is that they have just made all this noise about how security is now job number one and Kevin quoted Satya Nadella saying: *"If you're faced with the tradeoff between security and another priority, your answer is clear: Do security."* Except that they're not. The entire security industry is jumping up and down, waving their arms and saying "don't do it" — exactly as I did once before with XP — yet Microsoft is certain that they know better.

It's interesting that Kevin believes that the screen is being "OCR'd". I strongly doubt that's actually the case, at least not unless an actual JPEG or PNG-style graphic image is being displayed – in which case OCRing the image would be the only choice. As I noted last week, hooking into the Windows API that paints text onto the screen would be **far** more efficient. Behind every character glyph, what we see on the screen is a 16-bit UNICODE character which was rendered through a chosen font and turned into cleartype colorized pixel text. There's just no reason to look at the pixels of a screen that was just rendered from UNICODE and try to determine which characters they are. So my assumption would be that the textual output graphic API is being hooked and intercepted by Recall.

It was also very interesting to learn how economical Recall's storage is. This makes sense if it's storing and compressing text, since we know how much redundancy exists in linguistic text. But Kevin said that several days worth of work compresses to around 90KB of database storage. If we take "several days" to mean two, then that's around 45K per day. 50 gigabytes of storage allocation, consumed at the rate of 45K per day, would yield **3,042 years** worth of storage. I'm sure we'll learn more going forward, but I don't think Recall will be storing the past 90 days of a PC's use. It appears that it will always be recording the PC's entire life of use. That's why the title of Kevin's second post makes far more sense. His title began with: *"Stealing everything you've ever typed or viewed on your own Windows PC ..."* and I think that's exactly what Microsoft is actually planning to do. If they're able to capture and compress all of the text displayed on Windows 11 screens, and given the explosion in local mass storage capacity, and the efficiency of text compression, they clearly have the storage capacity to capture **everything for all time**.

And this brings us to the title I gave today's podcast: *"A Large Language Model in Every Pot"*.


Why would Microsoft want to be capturing every single thing a user types and views on their own PC throughout its entire lifetime of use? **I have a theory:** Microsoft wants to make a big splash in AI. So how about using all of that data to train an entirely personal local large language model? What if a future local large language model was not just used to index and search your PC's history timeline, but was continually being trained across your entire corpus of personal data so that it would be possible to conversationally interact with your own personal AI that has grown to know you **intimately** because it has been watching and learning everything you've been doing – for years? It would "know" (and I have "know" in air quotes) everything that had ever been entered into its keyboard and displayed on its screen. The entire history of that machine's use would become an ever-growing corpus that is continually training the model.

That would completely and profoundly forever alter a user's interactive experience with their PC. It would be a true game changer. It would be transformative for the PC experience. And if Microsoft has that up its sleeve, I can see how and why they would be super excited about Recall, even though Recall would be just the beginning. Even if the local large language model technology is not yet ready for delivery, the time to begin capturing all of a user's use of their machine **is as soon as possible**. That begins creating the corpus that will be used to train a future personal local LLM.

If this view of the future is correct, there's one large and glaring problem with this, which Kevin highlights and which Microsoft is conveniently ignoring – because they have no choice. What Microsoft **MUST** ignore is that the **actual** security of today's Windows is a catastrophe. Microsoft has not been paying more than begrudging and passing attention to security while they've been busily adding trivial new feature after new feature and never getting ahead of the game.

Last month's Patch Tuesday saw Microsoft patching 61 newly recognized vulnerabilities, 47 of them in Windows and another 25 for anyone paying for extended security updates.

May 2024 Risk Analysis



(<https://www.crowdstrike.com/blog/patch-tuesday-analysis-may-2024/>)

44% of those were remote code execution, 11% were information disclosure and 28% were elevation of privilege – none of which suggests that Windows would be a safe place to store the data that will be used to drive an entity that can be queried about nearly any aspect of you and your life which it has observed throughout the entire history of your use of that machine.

If this is, indeed, what Microsoft is planning – and having voiced it now it's difficult to imagine that it's not exactly what they are planning – then this is really a double-edged sword. The world stumbled upon the startling power of large language models, which Microsoft just so happens to own a big chunk of, and someone inside Microsoft realized that by leveraging the power of next-generation neural processing units, it would be possible to train a local model on the user's entire usage history of their computer. And that would create a personal assistant of unprecedented scope and power.

I would wager that today, the smarter people within Microsoft are wishing more than anything else, that instead of screwing around with endless unnecessary features and new unwanted versions of Windows, they had been taking the security of their existing system seriously. Because if they had, they would own a secure foundation and would stand a far greater chance of successfully protecting the crown jewels of a user's computer usage legacy. Instead, what they have today is a Swiss cheese operating system that is secure only so long as no one really cares what its user has stored. Depending upon who the user is, the data that **will** be accumulated by Recall will represent a treasure that is certain to dramatically increase the pressure to penetrate Windows. The entire professional security community understands this, which is the reason it's going batshit over Recall, while Microsoft has no choice other than to deny the problem because they're desperate to begin the data aggregation of their users so that it can be used to train tomorrow's personal PC assistant AI's.

So Microsoft will declare as they always do, that Windows is more secure than it's ever been, even though history always shows us afterward, that's never been true. Microsoft is going to have Recall installed, running and collecting its user's data in all forthcoming qualifying CoPilot+ Windows 11 PCs.

And don't get me wrong, the idea of being able to ask a built-in autonomous personal AI assistant about absolutely anything that we've ever typed into or seen on our computer is intoxicatingly powerful. For many of us who live much of our lives through our computers it would be like having a neural-link extension of our brain with flawless perfect recall. But it also represents a security and privacy threat the likes of which has never existed before.

When you consider the amount of digital storage that anyone can now easily own, it seems pretty obvious that this is going to happen sooner or later. Unfortunately, Microsoft has not proven itself to be a trustworthy caretaker of such information.

