A Large Language Model in Every Pot

Description: 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 10 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.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-977.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-977-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. Have you ever been pwned? Well, here's a way to know and whether you should worry about it. What certificate authority just lost their TLS server business? We'll talk about that, the end of ICQ, and Microsoft's new Recall feature that's coming to all Copilot+ PCs. Steve explains why it is not as secure as Microsoft has said, why it's in fact a real danger. He also has a theory, and I agree with it 100%, of why Microsoft is doing this. So a very interesting play for your information, next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 977, recorded Tuesday, June 4th, 2024: A Large Language Model in Every Pot.

It's time for Security Now!. Yes, adjust your spectacles and put your beanie on straight because this guy, Steve Gibson, is going to challenge you, he's going to excite you, he's going to thrill you, he's going to make you a geek just by proximity. Hello, Steve Gibson.

Steve Gibson: I think if you've survived more than a couple of these podcasts...

Leo: You qualify.

Steve: ...your geek status has been already established.

Leo: You more than qualify.

Steve: If you haven't gone running for the hills. It's like, agh. I got a piece of mail from one listener who said, okay, so I think I understand about 5% of what you're talking about.

Leo: That's pretty good. You're doing well.

Steve: But I do come away with something useful every week, so I keep coming back for more abuse. No, for more edification.

Leo: Well, and it's like lifting a heifer. Like when a cow, a baby cow is first born, you can lift it. If you lift it every day, you'll be able to lift a full-grown cow.

Steve: This is the analogy you've come up? We're lifting - lifting cows.

Leo: Keep listening every week, and in a year or two you'll be able to lift a cow. How about that?

Steve: And maybe you'll be able to throw a honeypot.

Leo: There you go.

Steve: Oh, actually there's a really interesting piece that Microsoft just revealed, the details of a honeypot they had been running for a long time. Anyway, I may be talking about that next week, if nothing more interesting comes along. But speaking of pots, today's title is "A Large Language Model in Every Pot." And we're going to go back and talk about Recall again. Well, okay. I'm stepping on my own sequence here. So we've got a lot of things to talk about.

When is a simpler application better than something complex? How did the first week of GRC's new email system turn out? 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? Whoops. And remember that early messaging service ICQ? Whatever became of it?

Finally, after I share a tip about what I consider to be 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 10 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. I think I know where they're headed. And it's very exciting, if I'm right. And it's also very troubling. And it's really a shame that they've been screwing around with Windows, adding features nobody wanted instead of making it more secure because they really can't do what they want to do. So we're going to have fun today.

Leo: Very interesting.

Steve: Unlike all of the other 976 podcasts that came before.

Leo: So boring. So boring.

Steve: Yeah.

Leo: No. We're going to have fun today, I promise you.

Steve: And we do have a great Picture of the Week.

Leo: Oh, haven't read it, I just know the caption. All right. All right. I am ready for the Picture of the Week, Mr. Gibson.

Steve: So I gave this picture the title, "But Officer..."

Leo: Okay. Does it need no explanation?

Steve: It really doesn't, once you see the picture.

Leo: All right. It's going to take me a minute to get it up on this computer. Here it comes. All right. I'm prepared. Are you ready? I'm going to scroll up. We shall enjoy it together. But Officer... There's a one-way street sign, a stop sign, and a no right-turn sign. What the - what am I supposed to do?

Steve: You know, Leo, you just have to wonder, like...

Leo: What the heck?

Steve: I know. Okay. So for people who aren't seeing this, we have a picture where a public street has come up to a T intersection. So you have to turn left or right. Well, there's a stop sign, so you certainly need to consider your options, thus stopping. The problem is that the street that you are intersecting with has been labeled as one-way, where all the traffic is moving from left to right. But below the stop sign, it's also very clearly marked that you must not turn right. There's the right-turn arrow with a big red slash through it. So I don't know. Do you back up? You know, like backing up would be the only thing you can do.

Leo: I think it's all you can do. But notice there's no outlet. You're in a cul-de-sac. So you're really dead in the water.

Steve: So you're right. Is that what the yellow sign says?

Leo: Yeah, it says "No Outlet."

Steve: It says "No Outlet"? I thought, yeah, I thought so. So now that...

Leo: So this is the worst.

Steve: That's something that would be seen by people going down the street waving at you because you're stuck, and you can't go anywhere.

Leo: I think this is a prank being played on self-driving cars. Whoever lives on this street added that sign knowing that a self-driving vehicle would then be complete stuck.

Steve: It would just explode, Leo.

Leo: It can't do anything.

Steve: It would just say, okay, I quit.

Leo: I can't do anything. I'm stuck. Oh, my god, that's hysterical.

Steve: Welcome to America. Okay. So I wanted to thank all of our listeners who correctly recalled that the random notes DOS app we were trying to remember last week was "Tornado Notes."

Leo: And I don't even remember that one, so I wouldn't have gotten it.

Steve: Yeah, it was not well - Leo, it was DOS. But you used DOS back in the day.

Leo: Oh, I used Sidekick. I used a lot of [crosstalk] DOS, yeah.

Steve: Yeah, yeah, yeah. So it was not Phil Katz of PKZIP fame. It was a guy named Jim Lewis of Micro Logic Corporation. And when I first encountered Tornado Notes from a company named Micro Logic Corporation of Hackensack, New Jersey, I wondered, why is that name so familiar? And it turned out it was because the same guy had created one of the most useful sets of 8.5 x 11" double-sided plastic sheet processor instruction reference cards the world had ever encountered. I have a picture of them in the show notes.

Now, 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...

Leo: Bury them with you.

Steve: ...buried alongside me.

Leo: This is a 6502, a Z80, and an 8086. It's all in there.

Steve: And there is a 68000, as well. I cannot begin to express how important they were back when I was writing assembly code...

Leo: Wow, look at this.

Steve: ...first for Apple's and later Atari's 6502-based machines.

Leo: This is so cool.

Steve: And Leo, I've got links on the next page to the PDFs of them. I mean, these things were significant to so many people.

Leo: Yeah.

Steve: I ran across someone over on Reddit who commented that it was a good thing these were 100% plastic or he would have worn his out. You know, they were indispensable. And I don't know where mine are. I'm sure they're here somewhere because I would have never thrown them out. They were just perfect. Now, you have on the screen now the 6502 card. And notice all the blank boxes. Those are missing opcodes. So that was important. You had to know, you know, what was available and what wasn't. And one of the reasons the 6502 microprocessor was so well used - Apple chose it, Atari chose it, Commodore chose it - was because it was so inexpensive. And the reason it was inexpensive...

Leo: It didn't do diddly.

Steve: It didn't do much, exactly. It, you know, it transferred all the burden to the programmer, and like most of those opcodes are empty in there.

Leo: Wow.

Steve: But it did just enough in order to get the job done. But this was just - it was so - so this guy named Jim Lewis, who later gave us Tornado Notes for DOS, a TSR, you know, the reason I knew his name when Tornado Notes came along is like, wait a minute, I've got these Instruction Reference Cards that I've been using forever. But anyway. Tornado Notes for DOS 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 - and this was the other thing - instantaneously. It began as a

massively overwhelming disorganized pile of rectangular notes. Didn't matter, you could just put anything, just random text in, didn't matter what shape or size they were. But then, as you typed successive characters of a string, all those notes that did not contain the substring that had been entered thus far would instantly disappear. 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 contained that same substring, which was often surprisingly useful at times.

Unfortunately, Jim, for all his brilliance, did not understand that Tornado Notes succeeded due to the constraints imposed upon it by its DOS environment. So when he created its successor, which was 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. And, you know, 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 could. But the very thing that was so beautiful about Tornado Notes was everything it did not do. So as it turned out in retrospect, you know, the thing, I mean, it being so minimal was what made it so compelling and useful.

And I'm mentioning this because there's a larger lesson here. One of the things the original designers of Unix also got exactly right was the idea of creating many simple commands that took some input, did something to it, and then produced some output. And then 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 kind of incrementally build up...

Leo: Exactly, yes.

Steve: ...a much more complex solution that's assembled from many smaller pieces interacting.

Leo: They call it "composable" because you compose a larger program out of pieces of smaller programs. And to my mind it makes it so much easier because you can bite off a little bite, figure out how it works, and because it's basically functional, you know, it's always going to give you the same result with the same input. You could slow put those together and build something out of it.

Steve: Yeah.

Leo: It feels to me like woodworking almost, like assembling a machine. It's great.

Steve: Like crafting a solution.

Leo: Crafting, exactly, yes.

Steve: So anyway, the point I hope to make here is that more is not always better. And, you know, for example, this is a lesson that the people who design the remote controls for A/V equipment appear to have never learned. Oh, my goodness, it's a joke that those things are so crazy. And I did notice that, you know, when I was thinking about this, that my freeware all just does one thing. You know? I create a little program. It just does one thing. If you want that one thing, that's the program you use. It's, you know, 23K. It does its job. And then you're done. And actually through the years people have been asking for many, many more features from SpinRite, and I've just said no, you know, SpinRite does what it's supposed to do. And that's what it's for.

So anyway, I just - I want to thank all of our listeners who said "I think you guys were thinking about Tornado Notes." And sure enough. And I wouldn't be surprised, I mean, there are DOS boxes around that could run Tornado Notes. I haven't run across a copy of it, but I probably have one on a hard disk around here somewhere. Anyway, 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, that is, if you don't listen to them all and don't know about it yet, you could go to our old GRC.com/feedback page which we've been talking about for 20 years which explains a bit about the nature of web form spam, which unfortunately is a thing, and it contains a pointer over to our new page, GRC.com/mail.

Anyway, the only post-announcement glitch we encountered was from users mostly using Gmail, but also a few other ISPs, I think Virgin Media was one, that use their own domains backed by those services, like Gmail. But 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 the securitynow@grc.com mailbox 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 that's, you know, shown in the email From: header of their email.

So some people were going over to the GRC.com/mail page and putting in their account name and their own domain, even though it's a front for Gmail. It turns out that the mail that they send actually comes from Gmail, so that was not an account that we'd ever seen before and so their mail was bouncing. As soon as I understood what was going on, I added a little comment on the form just to say, you know, for Gmail people, that was like by far the majority of users who were having a bounce problem, that that was what they had to do, and that problem went away. So people are paying attention to that.

Oh, also, anyone using an anonymizing email service will have a problem. I received an email from a listener who was using the SimpleLogin email anonymizing service by Proton, which by the way appears to be a very nice service. When that listener sent email to GRC, the sender's email was this bizarre long one-time 54-character random account name in front of the @simplelogin.com domain name. So again, GRC's filter had never seen that before, probably will never see it again. And it bounced that mail back. So we're not compatible, our approach is not compatible with email anonymizing services.

And I didn't mention it last week, but I actually have at the GRC.com/mail page what I call "The Prime Directive," which is nobody will ever get mail from us that they don't want. I mean, and I'm serious about that. We will also never divulge anyone's email address. Since sending email is a pain, you know, please unsubscribe if you're ever not happy and so forth. Anyway, to make a long story short, our listeners love the simple solution. You just register one time. You optionally subscribe to whatever announcement lists, if any, you may wish. And then from then on you can simply send email to securitynow@grc.com. I have been overwhelmed with notes of thanks and congratulations from listeners, and people I've never heard from before who were never going to sign up to Twitter just to maybe send me a note.

You know, and in fairness, Twitter is about so much more than that. You know, it's about building a community and a following, and following people, and networking. I had been just using it as a point-to-point instant messaging service, which after all is exactly what email is. So anyway, needless to say, as I said, I will never share anyone's email address. Oh, and I did want to say, if somebody writes to me, I will never share your email address when I share your feedback. And anyone requesting anonymity for their name, of course I will honor that.

Now, I should mention, and Leo I remember you mentioning this, too, when we first talked about it, one of the nice things about GRC's now-retired web form was that it solicited our listeners' location. And it was nice being able to include that when sharing feedback, you know, since it made the email feel a little bit more personal. So if you happen to think of it, let me know where you're writing from when you send me a note, and I'll just sort of toss that in when I share your feedback.

Leo: I'm wondering, you said it has to have the same domain as the server. So most email clients will let you choose a personality that says - so, for instance, I might be running on Gmail, but my email, I would like it to be leo@leoville.com. I can choose leo@leoville.com as my personality in Gmail. And even though it's originating from the Gmail server, it should look to you, to your server, like leo@leoville. You don't look at the underlying outbound server; do you? Maybe you do.

Steve: Yes. I actually do.

Leo: You do.

Steve: Yes, because...

Leo: Instead of just the email address, the reply-to address, in other words.

Steve: Yeah. The problem is the reply-to address is trivially spoofable.

Leo: Of course.

Steve: And so I wanted something that is a little less spoofable.

Leo: Okay.

Steve: I have a thread that I've not yet caught up in over in the newsgroups to do some brainstorming about whether I ought to change that because it would be easier if I just use the From: address. And I'm not sure that it really matters because any spammer could certainly be spoofing the Receipt-To address, as well.

Leo: Right, right.

Steve: So I may rethink that and change that, just to make it a little bit easier for...

Leo: That's a good advisory. You have to use the email address that your service provides.

Steve: Correct.

Leo: As opposed to any personality, any identity that you use.

Steve: Correct. And we ran across that with Gmail people and also, as I mentioned, SimpleLogin people. It's an anonymizing service from Proton. They also had to do that. But really, after I explained it, we stopped having anymore problems with signup. So my current work, this moment, you know, this evening, is to finish up automating and catching real-time email bounces. So I could immediately inform someone when GRC is able to 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. You've got to be careful. And so we'll ramp up from there.

So anyway, I wanted to thank everybody for their support. You know, everyone's interest is the reason I became convinced that we need to keep this going past 999. And, you know, here we are, already at 977, with our 20th birthday coming up in August.

Leo: Yeah, see? Yeah.

Steve: Yeah.

Leo: Old doesn't mean in the way.

Steve: While I was writing the note above yesterday, 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." Okay. The breach occurred one week ago on May 28th. In the breached data - get this, Leo - 361,468,099 email accounts were found. And HIBP (Have I Been Pwned) sent this email because 16 of those 361 plus million belonged to GRC.com.

The description of the breach that Troy included said: "In May 2024, two billion rows of data with 361 million unique email addresses were collated from malicious Telegram channels. The data contained 122GB across 1,700 files with email addresses, usernames, passwords and in many cases, the website they were entered into."

Leo: Does Troy email every one of those addresses? How did you get that? He must; right? Or do you sign up for some?

Steve: No, no. Yes. So I subscribed to a domain-wide free - it doesn't cost anybody, so I would recommend this. It's domain-wide. So you would, you know, do Leoville.com and TWiT.tv. And then you have to prove ownership of the domain. And once you do, anytime

Troy gets a hold of any new breach data, he'll scan the email addresses in the breach content and then notify you of any hits which may be one of your active email addresses having just been disclosed.

Okay. So he said: "In this case the data contained from this Telegram Combolists data breach, 122GB across 17,000 files, with email addresses, usernames, passwords, and in many cases the website they were entered into." He said: "The data appears to have been sourced from a combination of existing combolists and info-stealer malware." And we'll be hearing a little bit more about info-stealer malware because that comes up when we're talking about Recall again.

Okay. So naturally I went over, after I received this email from him, to see whether any of those 16 addresses which HIBP reported were of concern. Okay. The short version is none were. The longer version is 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 first names, just, you know, abigail@grc.com, amantha@grc.com, and so forth, A through Z, hoping to get lucky. Immediately after that we retired our original and, you know, oh-so-very-innocent use of our first names for email. That just became impractical.

The wonderful open source 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. There really is nothing comparable. I know lots of people run, you know, Sendmail and Postfix and so forth over on Linux. And I get that. Those are certainly mature platforms, too. But Windows hMailServer. It's another of those rare software creations that has no bugs. Just like John Dvorak gets no spam, this thing 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 just last week after many years of trouble-free service only to obtain support for TLS 1.3, which I did not have in my previous instance. And remember, 1.2 appears to be fine. You know, 1.3 exists. It's real. People should support it. But 1.2 ain't going away anytime soon because it's still, what is it, 86% of connections or something like that.

Anyway, 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, in my case at GRC. I just checked the server when I was writing this yesterday. I currently have the blocklist expiration set for two hours. And at the moment I checked, 473 individual IP addresses were currently being blocked. So within the previous two hours, 473 different spamming SMTP servers had connected to GRC and attempted to send spam. Not to actually, you know, not even to any valid email address, but just to throw crap at the wall, hoping to get lucky.

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

Leo: Yeah. I think it just reflects humanity. That's the problem.

Steve: Yeah.

Leo: As we go along, it's more and more like the people who make it.

Steve: Yes, yes. The trifecta of the Internet being anonymous, global, and free, those three things, 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. You know, we're not going to let any of that in.

Leo: No.

Steve: Okay. Okay. So I want to talk about a new supply chain attack vector. But let's take a break first, and then we will get into some security news of the week.

Leo: All right. I think you should write a manifesto, Steve. We're mad as hell, and we're not going to take it anymore.

Steve: Well, we're going to stamp on it. We're going to hide behind our NAT routers and hope that all that junk out there - I mean, come on, 473 servers just hooking up to GRC in the course of two hours, spewing junk at it.

Leo: It's mindboggling; isn't it.

Steve: It's just...

Leo: It's just amazing, yeah.

Steve: It's really sad.

Leo: Yeah. It's the world we live in, I'm sorry to say. Well, you know, one good thing about doing this show is because you focus on all this stuff, we have the best sponsors when it comes to security; right? They flock to us. In fact, we talk to people all the time that say, hey, can I be on Security Now!? And most of the time I'm happy to say we have to say, no, it's sold out for the next quarter. All right, Mr. G. On we go with the show.

Steve: And speaking of what a sad mess the greater Internet has become...

Leo: Yes.

Steve: ...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' [in air quotes] Stack Overflow users to push malware." Okay, now, 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, which as far

as I know no one has ever done before, 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 (who's also 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," they write, "are usually promoted using names similar to other popular packages," you know, a process we've talked about before 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," Bleeping Computer writes, "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 Stack Overflow account' - it had a nonsense name of EstAYA G - 'created roughly two 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" - again in air quotes - 'to their issue, even though the "solution" is unrelated to the questions being posed by developers.'

"In this case, the pytoileur package contains a setup.py" - you know, Python - "file that pads a Base64 encoded command which executes with spaces, so that unless you enable word wrapping in your IDE, you know, your Integrated Development Environment, or text file editor, this Base64 blob will be pushed all the way out past the right margin and offscreen so you'll never see it. When that blob of Base64 is deobfuscated, the command will download an executable named 'runtime.exe' from a remote site and run it."

They write: "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 the users' web browsers. It also appears to search through documents for specific phrases and, if found, steals the data in them, as well. All of this information is then sent back to the attacker, who can sell it on the dark web markets or use it to breach further accounts that are owned by the victim."

They said: "While malicious PyPI packages and information-stealers are nothing new, the cybercriminals' strategy now to pose as helpful contributors on Stack Overflow is an interesting new 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" - and they said "with word wrap enabled" - "for unusual or obfuscated commands which will be executed."

I have a picture in the show notes of the window. And you can where there is a Python class named "install command," and then a definition of run which is going to print something, and then you can see a big bunch of white space. Well, that's all spaces that will push this huge green blob of Base64 encoded code far off to the right so that, if someone did not have word wrap enabled, they'd never see this. They would look at it

and go, huh. Well, okay. I don't quite get what it's doing, but looks fine. Nothing bad there. When in fact there's a big blob of badness which the exec function will deobfuscate and then run.

So anyway, 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 on the same machine. That means that Recall's SQLite database is 100% vulnerable to exactly this sort of info-stealing malware. So it's not like Microsoft has created some miracle that is going to protect this database. And we'll be talking about more of that in a minute.

So in other news, we have another certificate authority 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, Google will be using a recently added new feature that allows it to manage certificate trust based on certificate issue dates. So Chrome will not be trusting any new certificates issued by GlobalTrust after the end of this month, June 30th.

Now, through the nearly 20 years of this podcast we've seen and discussed 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 that 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 - well, "established" is an interesting choice of my words - a multi-year history of misconduct, and they've lost the trust of the industry. Google will be enforcing a ban retroactively on all Chrome versions down to 124. So lots of previous Chrome versions. I don't know who would not be keeping their version of Chrome up to date, but okay. And the other browser makers have not yet announced a similar decision, although Mozilla appears to be aware of the problems with GlobalTrust and is concerned.

On the other hand, 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 business of selling web server certificates. In other words, whether or not Apple and Mozilla should choose to follow, GlobalTrust is done for now, at least on the TLS web server certificate business. They may be selling lots of certificates for other purposes, but not for any Chrome browsers in the future.

Those of us who have been around since the dawn of the Internet will likely remember the first successful instant messaging app known as ICQ. It was meant to be short for "I seek you." The system was originally developed back in 1996 by an Israeli company named Mirabilis. I practiced pronouncing it earlier, and now I can't do it. Mirabilis.

Leo: No, Mirabilis.

Steve: Mirabilis. I thought - okay, right. Mirabilis.

Leo: Mirabilis.

Steve: Two years after it was created - ICQ was created 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 sort of thought of it like through the years, wondering whatever became of it. At its peak around 2001, it had more than 100 million accounts registered.

Leo: Wow.

Steve: 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 had dropped to around 11 million monthly users.

And finally, the reason the subject came up is that a week and a half ago, on May 24th, the website of ICQ.com announced that the service would be shut down about three weeks from now, on June 26th, 2024. So it had a pretty good 28-year run for an instant messaging service that was largely passed by when smartphones and other major social media service got into the game. But it was there from the beginning and kind of cool.

Okay. Now, completely off topic, but this has been something that I've been wanting to just make sure everybody knew about for a while. 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 "Dj Vu" which stars...

Leo: I feel like I've seen it before. I don't...

Steve: You probably have, Leo. It's not new. And yes, I get your...

Leo: Okay, just checking. Actually, I don't feel like I've ever seen it.

Steve: Oh, no kidding?

Leo: I don't usually think of Val Kilmer and Denzel Washington as being sci-fi stalwarts.

Steve: Oh, Leo.

Leo: Oh, all right.

Steve: Okay. So listen to - okay. So Denzel Washington, Val Kilmer, and some other recognizable actors from Hollywood's inventory. As I was watching it for maybe the fourth time, I kept thinking over and over, you know, it is - as I was watching this perfectly and often leisurely paced two-hour movie unfold 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's 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. You know, that's part of what's so gratifying about this movie. Now, as I said, it's not a new movie. It was released 18 years ago, back in 2006. But it stands up, and it feels 100% contemporary.

I realized that since this podcast is closing in on its 20th birthday, every time I've seen this movie I'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, you know, 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 Leo, I guess you don't...

Leo: Lisa said she's seen it. So it'll be dj vu for her, but it will be whatever it is, premiere view for me.

Steve: It is just so good. I just...

Leo: I'm watching it tonight. I need something to watch.

Steve: It is wonderful sci-fi.

Leo: I love Denzel, of course.

Steve: And, yes, I do, too. And it will not disappoint you.

Leo: Okay. Thank you. Finally, something to watch tonight.

Steve: So our listener Jeff Price, he wrote and said: "Leo touched on this, but Fastmail allows you to create these unique random email addresses. What most people forget is Apple lets you create these, as well. They call it Hide My Email." So I just 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, you know, as third-party cookies and as Google tries to promote their sandbox anti-tracking technologies.

Kirk Sexton wrote: "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." And he's talking about Recall. He says: "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 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?" And he finished: "To 999 and beyond! All the best, Kirk Sexton."

So Kirk raised a great point, I think. 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 beyond what they've announced. But everything we know now suggests that the Recall data are just SQLite files stored under the user's AppData directory in a new folder called "CoreAIPlatform." 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 they can later be 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. So, you know, there just seems to be so many things that have not been well thought through here.

Okay. And then just one piece of feedback. I'm way far behind, just so everybody knows. The first week of listener feedback email was intense, with many listeners, you know, wanting to say hi, to express their happiness there's now a way to send me thoughts without engaging in social media. So, yeah, as I said, I'm way behind. But I figured I'd share one piece of feedback that's primarily about a SpinRite owner's experience, first with SpinRite 6.0, or by comparison with SpinRite 6, and then with 6.1.

Our listener Mark Jones sent email with the subject "Wow! SpinRite 6.1 is amazing." He wrote: "Dear Steve. Long-time listener, occasional source of feedback." He says: "(I was @mjphd on Twitter.) I'm so happy to be using email. I only kept my X account for Security Now! feedback." He said: "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 1TB spinner and a 250GB SSD that seemed to have slowed. The results are nothing short of remarkable on both drives. In only four hours, the 1TB was rejuvenated. That would have taken days using SpinRite 6. 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 at 19 MB/s at the front and middle, and 80 MB/s at the end." So 19 front and middle, 80 MB 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's mind-blowingly 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."

Okay. So let's talk about Recall again because we have additional information. And Leo, I'll find a point to pause here.

Leo: Sure.

Steve: For our final.

Leo: Yeah.

Steve: Okay. So I think that a data-driven theory about Microsoft's future plans for this technology emerged after I read a recent posting by a well-known and well-informed security researcher named Kevin Beaumont. Since last week's episode, which I titled, as

we know, "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, immediately 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 PCs' screens echoes my own. It's immediately obvious to anyone who's been around the block a few times that this is, indeed, a 50GB 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. We'll get to that.

Ever since his immediate posting in reaction to the announcement of Recall, Kevin has been playing with it. After reading what Kevin wrote, a light bulb went off for me. So I'm first going to share Kevin's follow-up piece which further describes Recall in much more detail. 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."

Okay. Now, before switching into Q&A mode, which he does later, Kevin began his newly informed discussions of Recall by writing this. He said: "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 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 thing 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 wouldn't have. This was like watching Microsoft become an Apple Mac marketing department. At a surface level, it is great if you're a manager at a company with 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," Kevin wrote.

He said: "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."

Okay, now, 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 work subject, but that for the common user - the rest of us - the idea that our machines are watching and recording everything we do, even if it would only be for our own later access, is mostly just creepy. You know, 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're 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 like to now detail why." He said: "Strap in. This is crazy. I'm going to structure this as a Q&A with myself now, sourced from comments I've seen online, as it's really interesting seeing how some people hand-wave the issues away."

Okay. So now Kevin switches into Q&A format. He asks himself a question. So the question is someone saying, "Well, the data is processed entirely locally on your laptop; right?" Answer: "Yes. They made some smart decisions here. There's a whole subsystem of Azure AI, et cetera, code that processes on the device." Okay, question: "Cool, so hackers and malware can't access it; right?" And he says, "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's not what criminal hackers do. For example, info-stealer 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." And then he has a snippet from the journalist that says: "That is what Microsoft told me, that attackers would have to get physical access to your laptop and sign into it to get hold of the screenshots." Kevin says: "Not true." The questioner says: "Microsoft say only that user can access the data." Kevin: "That is not true. I can demonstrate another user account on the same device accessing the database."

Okay. The question: "So how does this work?" Kevin answers: "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 plaintext. OCR is a process of looking an image and extracting the letters." Question: "What does the database look like?" And Kevin shows some screenshots like those that we saw last week. Just looking like, you know, a SQLite database with rows and columns, recognizable filenames.

Question: "How do you obtain the database files?" Answer: "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 seconds of video of two Microsoft engineers accessing the folder." And then Kevin 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 is 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."

Question: "But normal users don't run as admins." Answer: "According to Microsoft's own website, in their Recall rollout page, they do." And then he has a snippet from Microsoft.com where it says: "Making admin users more secure: Most people," says Microsoft, "run as full admins on their devices, which means..."

So Kevin says: "In fact, you don't even need to be an admin to read the database. More on that in a later blog." Question: "But a UAC prompt appeared in that video, that's a

security boundary." Kevin replies: "According to Microsoft's own website and MSRC, UAC is not a security boundary." And he quotes Microsoft saying: "More important, Same-desktop Elevation in UAC is not a security boundary. It can be hijacked by unprivileged software that runs on the same desktop. Same-desktop Elevation should be considered a convenience feature." So now Microsoft is saying, oh, well, you know, that's just for convenience.

So the questioner asks: "So where's the security here?" Answer: "They've 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." Question: "Does it automatically not screenshot and OCR things like financial information?" A: "No." We know that it does. Q: "How large is the database?" Kevin says - and here was one of the first ahas that hit me. Kevin says: "It compresses well. Several days working is around 90KB," nine zero kilobytes for several days of work. He says: "You can exfiltrate several months of documents and key presses in the space of a few seconds with an average broadband connection."

Question: "How fast is search?" He says: "On device is really fast." Question: "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." He said: "I actually have a whole bunch of things to show, and I 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."

So question is "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," he says, for example, "Microsoft Edge InPrivate mode is excluded, but Google Chrome isn't." He said: "Every user interaction, for example minimizing a window. There is an API for user activity, and third-party apps can plug in to enrich data and also view stored data." Well, that's news, and interesting. He says: "It also stores all websites you visit, even if third party."

Question: "If I delete an email/WhatsApp/Signal/Teams message, is it deleted from Recall?" A: "No, it stays in the database indefinitely." Question: "Are auto-deleting messages in messaging apps removed from Recall?" A: "No, they are scraped by Recall and available." Q: "But if a hacker gains access to run code on your PC, it's already game over." Kevin says: "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 info-stealer," he said, "I used Microsoft Defender for Endpoint, which detected the off-the-shelf info-stealer. But by the time the automated remediation kicked in, which took over 10 minutes," he notes, "my Recall data was already long gone."

Question: "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" - he means that the people offering the service have seen, he said - "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'll be able to buy stolen customer data from insurance

companies, et cetera, because all code required to do this has been pre-installed and enabled on Windows by Microsoft."

Q: "So did Microsoft mislead the BBC about the security of Copilot?" A: "Yes." Q: "Have Microsoft misled customers about the security of Copilot?" A: "Yes. For example," he says, "they describe it as an optional experience, but it is enabled by default, and people can optionally disable it. That's," Kevin says, "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 of course I need to note here that the tyranny of the default will be at work. We know that whatever is the default setting is what 99.99% of all Windows users will leave active. I don't know if any of you have seen people using Windows computers, but for some reason they always leave those stickers all over the keyboard. And I just, I can't believe it. It's like you realize the computer will still work if you peel those stickers off the keyboard. You don't need to be, you know, advertising the crap that came from the manufacturer. But anyway, the tyranny of the default.

So question: "Recall only applies to one hardware device." Kevin replies: "That's not true. There are currently 10 Copilot+ devices available to order right now from every major manufacturer. Additionally, Microsoft's website says they're 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?" Kevin replies: "I'm not a privacy person or a legal person. I will say that privacy people I have 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."

The question: "Are Microsoft a big, evil company?" Kevin: "No."

> **Leo:** Hell, yes. Oh.

Steve: "That's insanely reductive." He says: "They're super smart people, and sometimes super smart people make mistakes. What matters is what they do with knowledge of mistakes." So the question: "Aren't you the former employee who hates Microsoft?" Kevin says: "No. I just wrote a blog this month praising them. It was 'Breaking down Microsoft's pivot to placing cybersecurity as a top priority.' My thoughts on Microsoft's 'last chance saloon' moment on security."

So we have a couple, just two more. Question: "Is this really as harmful as you think?" Answer: "Go to your parents' house, your grandparents' house, et cetera, and look at their Windows PC. Look at the installed software in the past year. Try to use their device. Run some AV 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?" Answer: "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."

He says: "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.'" He said: "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.

Leo: Seems like a great piece. I mean, I've read it, and I was very impressed.

Steve: Yup.

Leo: And he makes a strong case. The one thing that's a question mark, a lot of the things he describes sounds like you had to be on the physical PC. But he says you don't. So malware would be able to escalate the UAC and do all those things, look across accounts.

Steve: Yup.

Leo: All of that stuff. Okay. So the real issue is, if malware gets in your system, they've got access to everything you've done.

Steve: Right. There is now much more that it has access to. Let's take our final break, and then I'm going to talk about what I think is really going on.

Leo: Yeah. Why would Microsoft do all this?

Steve: Yup.

Leo: What's the plan here? Hmm.

Steve: I think there is one.

Leo: All right. Steve. You've set us up well. Obviously this is a bad idea.

Steve: Why would they do it?

Leo: But Microsoft's going full speed ahead with it. Why?

Steve: Okay. So 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. Now, I have to say I have some personal experience with endeavoring, and failing, to get Microsoft to change its plans.

Leo: Can anybody say "raw sockets"?

Steve: Uh-huh. 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 MSBlast worm would have taken them off the Internet had it 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 have never been present in a consumer OS. Raw sockets never came back because they just beg to be abused.

Okay, now, I learned my lesson from that experience. I have no interest in lobbying Microsoft to change its behavior. You know, 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 they're not. The entire security industry is jumping up and down, waving their arms, and saying "Don't do it," exactly as I once did before with XP. Yet Microsoft is certain that they know better.

Now, 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 clear-type 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 Kevin's "several days" to mean two, then that's around 45K of storage required per day. That means that 50GB 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 "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 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 you had ever 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 of 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 large language model.

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 but to ignore it. 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. 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're planning, then this really is 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 why 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 AIs.

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 users' 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 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.

Leo: Wow. I think you're exactly right. I mean, that's almost what they're proposing anyway is you can always query the machine about everything you've done.

Steve: Well, they're saying "timeline," that you can query a timeline. But if this thing, if they're capturing text from the screen, Leo, and Kevin saw 90KB was stored after several days of use, that means that that 50GB that they want to set aside, this is not a 90-day rolling window which I thought last week. They're going to store everything you ever do for your entire life of your use of that machine.

Leo: Right, right.

Steve: And in fact you're going to want that to be portable to the next machine you move to so that you're able to take that accrued data with you from one, you know, three years from now when you need to buy a new Windows 13 machine.

Leo: It could be secured; right? You could do this right; couldn't you?

Steve: Yes. And what they're doing - I think you could. I think you could, I mean, you would need new hardware because you need some sort of the equivalent of an HSM. Basically you'd want this super Jeeves to be in its own enclave that could not be exfiltrated from.

Leo: Yes, that's right.

Steve: Where data goes in, and nothing comes out. And then, I mean, but it would - imagine that, Leo. It would be compelling to be able to ask your computer anything that you ever did with it.

Leo: I'm well aware of that. That's the...

Steve: It's perfect recall.

Leo: Yeah, the endgame for all of this. I've even referred back...

Steve: And you've been talking about your own local smaller corpuses or corpi and how useful that is.

Leo: Right. And I've talked about...

Steve: This would be that.

Leo: ...the founder of DEC, not the founder, one of the designers of DEC just passed away recently.

Steve: Gordon Bell.

Leo: Gordon Bell, who had the same idea. He had a camera around his neck. He wanted to record everything he ever did. This is ever before we had these powerful LLMs.

Steve: And the storage capacity to record our life.

Leo: Right, right. Well, the issue always was, and with Gordon's database, is well, okay, I've got it. What do I do with it?

Steve: Right.

Leo: I can't in a reasonable way parse it. Well, now we can.

Steve: Yes, yes.

Leo: And so I'm very interested. I ordered the Limitless Pin which records all our conversations, the idea, same thing, being to allow you to query that. You know, what did I say to Steve? I think this is the single most useful persuasive use of AI is as an assistant that knows everything about you. But, boy, that poses some big problems. It's almost as if we need an initiative to create a way. It also solves other problems because data privacy's a huge issue. We need a way, something that you can - Stacey Higginbotham used to call it The Blob, a place where you could securely, in Secure Enclave, store all your data for your own personal use, not so that other people could invade your privacy, but for your own personal use. And this is the best possible use.

So I think we're on the right track. I think this Microsoft implementation could kill it in its tracks. It could actually have a - this is what worries me is people are moving so fast, with so little regard for safety, that they could have the opposite effect. They could get people so scared about their security and privacy that they give up entirely on AI.

Steve: Well, and they're frankly lying...

Leo: Yeah.

Steve: ...about the security.

Leo: They're misrepresenting it, yes.

Steve: Yes. I mean, all this is is some files under the user's app directory.

Leo: Right.

Steve: This is not some hocus pocus. And so everybody knows how to exfiltrate files. Kevin did it. There's now a GitHub project that is able to display all your Recall data.

Leo: Well, I'm glad that he published this paper. I'm glad you did this show because up to now the press, not knowing any better, and I include myself, we've parroted Microsoft's assertions that, well, it's all on device. It's all local. It's all safe. It's encrypted. It's only available to you. I always - I have pointed out in the past that it's only encrypted as long as you don't log in. This is the second part of that. Once you're logged in, it's decrypted, and then available to any malware on your system. Yeah, I think people will - I hope the press will start to come around and say, hey, wait a minute, this isn't as secure as you said it was.

Steve: Well, our listeners are preemptively protected; right? I mean, they're going to turn this off.

Leo: Like that.

Steve: But unfortunately, there's no reach. Well, there's minimal reach. But, you know, there's a bazillion Windows 10 or Windows 11 users, and they're going to think, hey, this is cool. I can scroll back in history. And this is Microsoft getting ready for something that comes next.

Leo: Yeah, I agree. You know, Apple has a solution called Timeline. It's a backup solution that keeps everything you do in a timeline database, a vault. Hard links to every version of every document. So they're kind of doing something similar. Nobody's ever questioned the usefulness or the security of it. I don't know how different it is. But, yeah, this is a problem. This really is a problem.

Steve's done it again, hasn't he, kids. This is why we wait for Tuesday with bated breath. Steve is the man in charge of GRC.com, the Gibson Research Corporation dot com. And it is the place you can email him. Now, what should they do again? They email...

Steve: So first you need to register. Otherwise your email will not get through. So just go to GRC.com/mail.

Leo: Okay, there you go.