Security Now! #999 - 11-05-24

AI Vulnerability Discovery

This week on Security Now!

Google's record-breaking fine by Russia. (How many 0's is that?) RT's editor-in-chief admits that their TV hosts are AI-generated. Windows 10 security updates set to end next October... or are they? When a good Chrome extension goes bad. Windows .RDP launch config files. What could possibly go wrong? Firefox 132 just received some new features. Chinese security cameras being removed from the UK. I know YOU wouldn't fall for this social engineering attack. What's GRC's next semi-commercial product going to be? And what's the prospect for AI being used to analyze code to eliminate security vulnerabilities?

When handrails are not optional

Security News

Google's Record-breaking Fine

It's a shame that our favorite Russian Internet watchdog, Roskomnadzor, is not the Russian entity that's been levying fines against Google over its management of YouTube, since it would have been fun to say that name many more times during this reporting. But nevertheless, this bit of news was too fun – and bizarre – to pass up. It seems that by Russia's accounting Google currently owes some large Russian media outlets a rather significant sum in fines.

We noted last week that the few millions of dollars that the U.S. SEC had levied in fines against four publicly traded U.S. companies would be unlikely to change those companies' behavior because the fines fell far short of being significant for them. However, this is not the case here with Google and these Russian media companies – quite the reverse, in fact.

Here's the story as it was recently reported by *"The Moscow Times"* under the headline *"Russia Fines Google $2.5 Decillion US Dollars Over YouTube Bans."*

The RBC news website reported Tuesday that [scofflaw] Google has racked up some 2 undecillion rubles (which is the equivalent of $2.5 decillion US dollars) worth of fines in Russia after years of refusing to restore the accounts of pro-Kremlin and state-run media outlets. RBC cited an anonymous source familiar with court rulings against Google.

According to RBC's sources, Google began accumulating daily penalties of 100,000 rubles in 2020 after the pro-government media outlets Tsargrad and RIA FAN won lawsuits against the company [of course, Russian lawsuits] for blocking their YouTube channels. Those daily penalties have doubled each week, leading to the current overall fine of around 2 undecillion rubles.

Undecillion [they explain] is a number equal to 1 followed by 36 zeros, or one trillion trillion trillion rubles. Google, whose parent company Alphabet reported revenue of more than $307 billion in 2023, is unlikely ever to pay the incredibly high fine as it far exceeds the total amount of money on earth.

A total of 17 Russian TV channels have filed legal claims against Google, according to one of RBC's sources. Among them are the state-run Channel One, the military-affiliated Zvezda broadcaster and a company representing RT editor-in-chief Margarita Simonyan.

YouTube, which is owned by Google, blocked several Russian state-run media outlets over their support of the full-scale invasion of Ukraine. Authorities in Moscow retaliated with these fines but stopped short of blocking YouTube outright. On Thursday, the Kremlin called the fine against Google "symbolic." [I'd be inclined to call it embarrassing.]

Kremlin spokesman Dmitry Peskov told reporters at a daily briefing: "Although it is a concretely formulated sum, I cannot even pronounce this number. Rather it is filled with symbolism. In fact, this should be a reason for Google's management to pay attention to this and fix the situation."

This seems unlikely given that Google's Russian subsidiary filed for bankruptcy in the summer of 2022 and was officially declared bankrupt last fall. And Google had earlier halted advertising in Russia to comply with Western sanctions over the war in Ukraine.

RT Margarita Simonyan: Our hosts are fake
Just a quick note, since I saw the editor-in-chief's name Margarita Simonyan mentioned in that piece. I had noted that she also recently admitted that many of RT's Television hosts do not exist and are entirely AI-generated, along with their fake social media accounts. She predicted that journalism would disappear in the near future.

Windows 10 has one year to go (or does it?)
A recent posting to the 0patch (zero patch) blog regarding next year's end of Windows 10 security updates contained a bunch of interesting news. This included what Microsoft plans to charge users who would rather remain on Windows 10 – or who have no choice due to Microsoft's arbitrary minimal system requirements policies. Here's what the folks at 0patch recently wrote. Their blog post headline was *"Long Live Windows 10... With 0patch"* and their subhead was: *"End of Windows 10 Support Looming? Don't Worry, 0patch Will Keep You Secure For Years To Come!"*

October 2025 will be a bad month for many Windows users. That's when Windows 10 will receive their last free security update from Microsoft, and the only "free" way to keep Windows using securely will be to upgrade to Windows 11. Many of us don't want to, or simply can't, upgrade to Windows 11.

We don't want to because we got used to the Windows 10 user interface and we have no desire to search where some button has been moved to and why the app that we were using every day is no longer there, while the system we have is already doing everything we need.

We don't want to because of increasing enshittification including bloatware, Start Menu ads, and serious privacy issues. We don't want to have an automated integrated screenshot- and key-logging feature constantly recording our activity on the computer.

We may have applications that don't work on Windows 11.

We may have medical devices, manufacturing devices, POS terminals, special-purpose devices, ATMs that run on Windows 10 and can't be easily upgraded.

And finally, our hardware may not qualify for an upgrade to Windows 11: Canalys estimates that 240 million computers worldwide are incompatible with Windows 11 hardware requirements, lacking Trusted Platform Module (TPM) v2.0, supported CPU, 4GB RAM, UEFI firmware with Secure Boot capability, or supported GPU.

So what's going to happen in October 2025?

Nothing spectacular, really. Windows 10 computers will receive their last free updates and will, without some additional activity, start a slow decline into an increasingly vulnerable state as new vulnerabilities are discovered, published and exploited that remain indefinitely present on these computers. The risk of compromise will slowly grow over time, and the amount of luck required to remain unharmed will grow accordingly.

The same thing happened to Windows 7 in January 2020; today, a Windows 7 machine last updated in 2020 with no additional security patches would be really easy to compromise, as over 70 publicly known critical vulnerabilities affecting Windows 7 have been discovered since.

Leaving a Windows 10 computer unpatched after October 2025 will likely open it up to the first critical vulnerability within the first month, and to more and more in the following months. If you plan to do this, at least make sure to make the computer difficult to access physically and via network.

For everyone else, there are two options to keep Windows 10 running securely.

Option 1: Microsoft's Extended Security Updates

If you qualify, Microsoft will happily sell you Extended Security Updates (ESU) , which means another year, two or even three of security fixes for Windows 10 - just like they have done before with Windows 7, Server 2008 and Server 2012.

Extended Security Updates will be available to consumers for one year only (until October 2026) for the price of $30. Educational organizations will have it cheap - just $7 for three years, while commercial organizations are looking at spending some serious money: $61 for the first year, $122 for the second year and $244 for the third year of security updates, totaling in $427 for every Windows 10 computer across three years.

In other words – to interject here for just a moment – the cost to have Microsoft repair the mistakes that it has previously made in the design and operation of their own Windows software will double for their users every year.

Opting for Extended Security Updates will keep you on the familiar monthly "update + reboot" cycle and if you have 10k computers in your enterprise network it will only cost $4 million

If only there was a way to get more for less... [Oh, wait… There is!]

Option 2: 0patch

With October 2025, 0patch will "security-adopt" Windows 10 v22H2 – the final release of Windows – and provide critical security patches for it for at least 5 more years – longer if there's demand in the market.

We're the only provider of unofficial security patches for Windows, and we have done this many times before: after security-adopting Windows 7 and Windows Server 2008 in January 2020, we successively took care of 6 versions of Windows 10 as their official support ended, security-adopted Windows 11 v21H2 to keep users who got stuck there secure, took care of Windows Server 2012 in October 2023 and adopted two popular Office versions - 2010 and 2013 - when they were abandoned by Microsoft. We're still providing security patches for ***all*** *of these.*

With 0patch, you will be receiving security "micropatches" for critical, likely-to-be-exploited vulnerabilities that get discovered after October 14, 2025. These patches will be really small, typically just a couple of CPU instructions (hence the name), and will be applied to running processes in memory without modifying a single byte of original Microsoft binary files.

There will be no rebooting the computer after a patch is downloaded because applying the patch in memory is done by briefly pausing the application, patching it, and then allowing it to resume. Users won't even notice that their computer was patched while they were writing a document, in the same way that servers protected by 0patch get patched without any

downtime at all.

And just as quickly and easily, our micropatches can be un-applied if they're suspected of causing problems. Again, no rebooting or application re-launching.

0patch also brings "0day", "Wontfix" and non-Microsoft security patches. With 0patch, you won't only get patches for known vulnerabilities that are getting patched on still-supported Windows versions. You will also get:

- *"0day" patches - patches for vulnerabilities that have become known, and are possibly already exploited, but for which no official vendor patches are available yet. We've fixed many such 0days in the past, for example "Follina" (13 days before Microsoft), "DogWalk" (63 days before Microsoft), Microsoft Access Forced Authentication (66 days before Microsoft) and "EventLogCrasher" (more than 100 days before Microsoft). On average, our 0day patches become available 49 days before official vendor patches for the same vulnerability do.*

- *"Wontfix" patches - patches for vulnerabilities that the vendor has decided not to fix for some reason. The majority of these patches currently fall into the "NTLM coerced authentication" category: NTLM protocol is more prone to abuse than Kerberos and Microsoft has decided that any security issues related to NTLM should be fixed by organizations abandoning their use of NTLM. Microsoft therefore doesn't patch these types of vulnerabilities, but many Windows networks can't just give up on NTLM for various reasons, and our "Wontfix" patches are there to prevent known attacks in this category. At this time, our "Wontfix" patches are available for the following known NTLM coerced authentication vulnerabilities: DFSCoerce, PrinterBug/SpoolSample and PetitPotam.*

- *Non-Microsoft patches - while most of our patches are for Microsoft's code, occasionally a vulnerability in a non-Microsoft product also needs to be patched when some vulnerable version is widely used, or the vendor doesn't produce a patch in a timely manner. Patched products include Java runtime, Adobe Reader, Foxit Reader, 7-Zip, WinRAR, Zoom for Windows, Dropbox app, and NitroPDF.*

Though you're probably reading this article because you're interested in keeping Windows 10 secure, you should know that these patches are also available for supported Windows versions such as Windows 11 and Windows Server 2022, and we keep updating them as needed. Currently, about 40% of our customers are using 0patch on supported Windows versions as an additional layer of defense or for preventing known NTLM attacks that Microsoft doesn't have patches for.

So what about the cost? Our Windows 10 patches will be included in two paid plans:

- *0patch PRO: suitable for small businesses and individuals, management on the computer only, single administrator account - currently priced at 24.95 EUR + tax per computer for a yearly subscription.*

- *0patch Enterprise: suitable for medium and large organizations, includes central management, multiple users and roles, computer groups and group-based patching policies, single sign-on etc. - currently priced at 34.95 EUR + tax per computer for a yearly subscription.*

The prices may be adjusted in the future but if/when that happens anyone having an active

subscription on current prices will be able to keep these prices on existing subscriptions for two more years.

Okay. So this was obviously a sales pitch. But that doesn't make this any less true and relevant.

We know from our many years of covering 0patch that these guys are the real deal, and that they really do present a viable alternative to Microsoft's doubling-every-year extortion for the enterprise.

In this instance, I don't mind this sales pitch since it's easy to endorse what they're selling. Microsoft has made a strategic gamble to deliberately abandon its users to its buggy and vulnerability-ridden software as a clear means of scaring them into migrating to a fully supported operating system that most users would rather avoid... even when what that really means is that there will still be a constant flow of new vulnerabilities always being introduced while older problems are being resolved. And let's not even get started on the fact that Microsoft's Replay is an issue for Windows 11 users.

So, consider that remaining on a platform that works and that you love, into which Microsoft will no longer be continually introducing new vulnerabilities and which will, nevertheless, continue receiving updates for any newly discovered critical security vulnerabilities. This is the niche 0patch has decided to fill, and I think that for just 25 EUR per year – which is around $27 USD per year, extending the security coverage of that beloved platform for a minimum of another 5 years, makes a great deal of sense. And to top it all off, their on-the-fly RAM-based code patching system is significantly more user-friendly than Microsoft's nagging reboot and wait system.

Windows 10 users still have a year to go before that final Windows 10 v22H2 will need either third party or extended Microsoft update help. This podcast will be somewhere around episode 1045 at that point, and among other things, we should know a lot more about Recall by then.

Good Chrome extension goes Bad

We have another example of a popular Google Chrome extension with more than 100,000 installs suddenly becoming malicious. The extension known as "Hide YouTube Shorts" has been found to be performing affiliate fraud and collecting the browsing history of all its users. Security researchers say the extension appears to have turned malicious after it was transferred to a new developer.

I went over to the Google Play store to check it out. It's unclear to me why someone would want or need to "Hide YouTube Shorts" but it's clearly "a thing" since there were many other similar extensions listed as alternatives whose names suggest that they also do that. But in any event, in response to questions, the extension's new owner defends the overreach of the extension's privileges by saying that "in the future" there might be the need for more latitude. The brief write-up from the researcher who took the time to dig into this was interesting. They wrote:

What initially piqued my suspicions were the strange search suggestions on YouTube. Completely unrelated and disconnected from the context of my searches, sometimes in foreign languages. However, after analyzing the traffic in the browser tab and developer console, I did not notice any suspicious activity. It was only after I started debugging the extension that I noticed suspicious network activity and requests being sent to an unknown external service, containing the addresses of all visited sites and unique identifiers.

The extension does what it says it will do, but in the background it collects and sends information about all visited pages to an external server hosted on AWS. The information that the extension collects and sends includes an unique user identification number, installation number, authentication token, language, timestamp and full URL with path and arguments/ parameters, which allows reading the information in the address bar, including e.g. search history. Some users in the reviews on the extension page in the Chrome Web Store also indicated the possibility of redirecting to phishing pages. Due to the malicious nature of this extension, I do not know what other information it could have collected before, but due to the wide permissions of browser extensions, it should be assumed that it could also read information transmitted in forms, including credentials, logins, passwords, personal and sensitive data. Such data can be used for a wide range of attacks, so anyone who has used such an extension should assume that ALL data viewed and transmitted via the browser has been compromised and take immediate precautions.

The extension was originally developed by a single developer who maintained the source code on github, however the github repository was archived (LINK) on Sep 12, 2023 and the plugin was acquired (or maybe sold) to another developer. I haven't analyzed everything to the extent I would like, especially earlier versions to find out when the malicious change was made, although it seems that the first developer for some reason decided to use the all-pages reading model. When the extension was just entering the Google Web Store I analyzed its behavior and did not see similar problems with it.

I have no doubts about the intentional nature of the current developer's actions, as his responses to comments about the extension's permissions being too broad clearly demonstrate his intent.

So, again, the caution would be to attempt to minimize the use of browser extensions. We know that by far for the most part, their developers are well meaning and above board. But we also have incontrovertible evidence that there are also malicious actors swimming in these waters. Without the ability to fully analyze and vet every extension it becomes a numbers game where, statistically, the greater number of extensions being used the greater the chance that one of them will be malicious.

Russia's "Midnight Blizzard" leverages RDP config files

We know of all the trouble Windows has had over something as simple as .LNK link files. The exploits of those have been epic and we've lost count of the number of times they have been "fixed", only to rear up again. Some designs are just bad and are notoriously prone to abuse.

That's what I was put in mind of when I read that it's possible for a Windows .RDP file to preconfigure and launch a remote desktop session. It's like Microsoft has never learned anything from the past. And as we know, those who do not learn from the past are destined to repeat it.

The generic tech-press reporting on this explains:

Microsoft says that a notorious Russian cyber-espionage group is using a clever new technique to compromise victims and deploy malware on their systems. The technique involves sending malicious RDP configuration files to victims via email. If executed, the files connect a victim's PC to a remote RDP server. The connection allows the Russian group to steal data and deploy malware on the compromised device. Microsoft has attributed the operation to Midnight Blizzard, a cyber unit inside Russia's SVR Foreign Intelligence Service. The group has used the

new technique since October 22 and has targeted individuals in government, academia, defense, and NGOs across the US and Europe. This is the same campaign also spotted by AWS and CERT-UA.

Since the inherent insecurity of this entire design was just too much to believe, I went to the source, where Microsoft explains:

On October 22, 2024, Microsoft identified a spear-phishing campaign in which Midnight Blizzard sent phishing emails to thousands of users in over 100 organizations. The emails were highly targeted, using social engineering lures relating to Microsoft, Amazon Web Services (AWS), and the concept of Zero Trust. The emails contained a Remote Desktop Protocol (RDP) configuration file signed with a LetsEncrypt certificate. RDP configuration (.RDP) files summarize automatic settings and resource mappings that are established when a successful connection to an RDP server occurs. These configurations extend features and resources of the local system to a remote server, controlled by the actor.

In this campaign, the malicious .RDP attachment contained several sensitive settings that would lead to significant information exposure. Once the target system was compromised, it connected to the actor-controlled server and bidirectionally mapped the targeted user's local device's resources to the server. Resources sent to the server may include, but are not limited to, all logical hard disks, clipboard contents, printers, connected peripheral devices, audio, and authentication features and facilities of the Windows operating system, including smart cards. This access could enable the threat actor to install malware on the target's local drive(s) and mapped network share(s), particularly in AutoStart folders, or install additional tools such as remote access trojans (RATs) to maintain access when the RDP session is closed. The process of establishing an RDP connection to the actor-controlled system may also expose the credentials of the user signed in to the target system.

When the target user opened the .RDP attachment, an RDP connection was established to an actor-controlled system. The configuration of the RDP connection then allowed the actor-controlled system to discover and use information about the target system, including:

- *Files and directories*
- *Connected network drives*
- *Connected peripherals, including smart cards, printers, and microphones*
- *Web authentication using Windows Hello, passkeys, or security keys*
- *Clipboard data*
- *Point of Service (also known as Point of Sale or POS) devices*

In their blog posting, Microsoft then goes into detail about the attacks and provides pages and pages of IoCs – Indications of Compromise. Under their "Mitigation" section they have pages of things that can be done to keep this from happening.

I have an idea! How about never building this inherently incredibly dangerous and abuse-prone facility into Windows in the first place? If it's not there there's nothing to abuse. Seriously. Is it necessary to have an .RDP file type that causes a machine to configure a maximally insecure connection to a previously unknown remote server? I use RDP extensively and yes, RDP saves its connection settings into individual .RDP files and that's useful. But when those files are given the capability to initiate a connection on their own, this becomes an extremely dangerous design pattern.

If they are going to exist at all, such files should be tightly bound to the machine that created them – not something that can be received in the mail and then clicked-on by an unwitting user. Microsoft loves storing things in the registry, so RDP settings for the local machine could be retained there instead and this problem would not exist.

Handy as it inarguably is, there's just no safe way to send someone, anyone, a file that when executed causes their machine to connect to any foreign unknown machine with all of its local resources shared. There just isn't.

At the very least this facility should be firmly disabled by default for everyone, and only those few people who actually need to do this should then jump through some hoops to enable it on their machine, and possibly only for a self-limiting time.

I hope everyone knows to never click on anything received in email – even if it appears to have been sent from someone you know and trust. We can now add another to the long and growing list of email-based exploits. Emailed attachments are too useful to ban outright, and clever bad guys keep finding new ways to abuse this useful capability.

Mozilla has released Firefox 132

We're now at Firefox 132 which adds some new features and security fixes. The biggest new feature in 132 is support for a post-quantum key exchange mechanism under TLS 1.3 and blocking favicons if they are loaded via HTTP. Back when we were looking at Firefox's 3rd-party cookie handling there was a great deal of confusion since Firefox's UI and its behavior appeared to be at odds. So among the improvements, I was pleased to see the sentence: "Firefox now blocks third-party cookie access when Enhanced Tracking Protection's Strict mode is enabled." As we suspected and as GRC's cookie forensics system was revealing, that wasn't actually working the way its UI said it was. That's fixed by number 132.

The UK has removed about 50% of all Chinese-made security cameras

Under the sad but understandable category of "we don't trust camera-equipped black boxes made in China" we have the news that the UK government now says that over 50% of all Chinese-made security cameras have been from sensitive sites, such as government buildings and military bases. The government says it expects removal to be completed by April of 2025 despite the fact that the removal was initially ordered back in November of 2022, as we covered at the time. Presumably, there's a long procurement cycle for such things, so it took some time to get the replacement cameras into the pipeline. As we know, UK officials ordered all sensitive sites in the UK to remove all Chinese-made cameras, citing national security risks.

A surprisingly effective social engineering attack

At first I was tempted to call this the "there's a sucker born every minute attack" in honor of PT Barnum. But upon further reflection I think that would be too harsh because this is actually a rather clever form of social engineering attack that I suspect might ensnare many non-suckers.

It leverages the fact that most people, who are using the Internet and PCs today, have never really been and probably never will be completely certain or confident about how any of this magical hocus pocus stuff works. Mostly they just follow the instructions and do what's asked of them. And that's why I can understand why this new and rather blatantly obvious exploit is succeeding out in the world.

It begins with a faked CAPTCHA pop-up that instructs someone, in this case someone who wishes to watch a video, that they need to click on the "CAPTCHA" button to start authenticating that they are human:

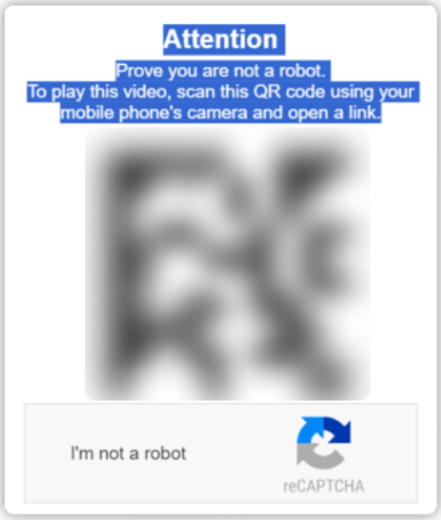

But this click that the user makes actually runs a bit of JavaScript which places a dangerous Powershell executable string onto their Windows clipboard:

```
powershell.exe -eC bQBzAGgAdABhA<...>MAIgA=
```

After pasting the Trojan-invoking powershell script onto their clipboard, it then displays the remaining instructions they must follow to ostensibly prove their humanity. Well, they are definitely about to prove their humanity, but not in a way that they intend. The pop-up reads:

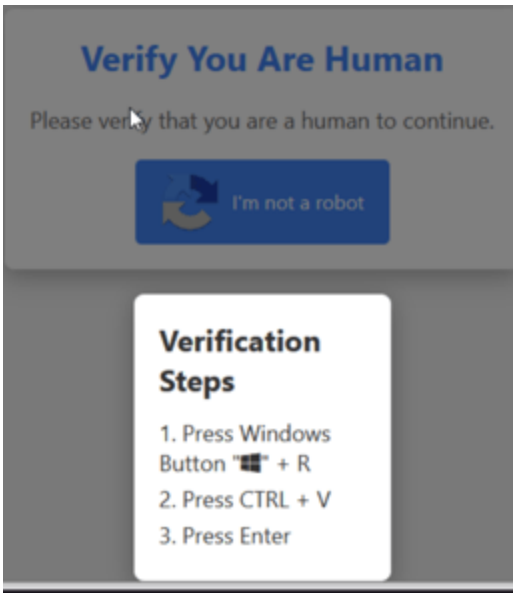

Verification steps:

- Press the "Windows Button" + R
- Press CTRL + V
- Press Enter.

The Windows + R brings up the Windows "Run" dialog with its "what would you like me to run" field highlighted. The Ctrl + V pastes the mysterious contents of the system's clipboard into that

RUN field. So that Run field now contains the executable powershell script to download, install and run Trojan malware on their computer. And this all culminates when they follow the final instruction of pressing "Enter" to, as Picard would say… "make it so."

As I observed, none of us would do this, but most people have no context for judging the consequences of any of their actions because they really don't understand what most of the buttons on their computer do.

My Next Project

DNS Benchmark — Version 2

I recently finished the work on GRC's email system. One of the system's originally-missing features was the capability to allow its users to update and migrate their email addresses at any time. My original thought was that since an email account didn't have anything other than zero, one or two subscriptions associated with it, anyone could simply delete an old account under an old email and create another under their new email. But after I saw very high spam complaint rates when mailing to SpinRite's owners from 20 years ago, I migrated SpinRite's purchase data into the email system which allowed me to send email which opened with "Back in 2005, someone named "whatever" at this email address purchased SpinRite. That had nothing short of a profound effect upon the spam complaint rates. So the email system now knows about SpinRite owners. And while I have also interconnected our e-commerce and e-mail systems, it would be better to have people able to simply edit their email address any time they wish.

So, I wanted to let everyone here know that since they last visited GRC's email management page it now has a simple editable email address field to allow anyone to easily change their email address at any time.

And once that was done, I was able to address the last remaining loose end of the SpinRite 6.1 offering, which was to create a video walkthrough demonstration of SpinRite in action. Since booting DOS and using a textual user interface is becoming increasingly foreign, I wanted a way to allow someone who might be considering whether to purchase SpinRite to get a quick and clear sense for what it looks like when it's running. So that also now exists...

Which brings me to the announcement I teased last week. As I've mentioned a number of times, GRC's #1 by far most popular software of all time is the DNS Benchmark. I've been astounded by its popularity. As I'm assembling these show notes it's been downloaded 9,313,642 times at 1600 downloads per day. The Benchmark pages have a page that solicits feedback and I'm constantly receiving retests for new features. Mostly, people are wondering how the speed of encrypted and privacy-protecting DNS using DoH, DoT or DNSCrypt compares with regular plaintext DNS. And despite the glacial progress of IPv6, many people are requesting that I add support for IPv6 to the Benchmark. Other great ideas have been to allow the Benchmark to verify the domain filtering being done to services like NextDNS, and others who are wishing to avoid local domain name blackouts would like to use the Benchmark to locate servers that are not blacked out.

The other thing I hear more generically is that people would like to have a way of supporting my continuing work on all things GRC. Our newsgroups, forums, ShieldsUp!, the DNS Spoofability test, all the other freeware that I write and we offer, and everything else.

So I've decided that my next project, before I create "Beyond Recall" for super-fast and super-secure data deletion, which will precede the development of SpinRite 7, will be to revisit the DNS Benchmark and give it a major version update to v2.0. There will still and always be a free release available. But I would like it to support itself, and I think it should be able to. So I plan to offer all of those new features for $9.95 in a "Plus" edition and for the real DNS pro, I'll also offer a "Pro" edition for $19.95. In addition to everything the "Plus" edition will have, the Pro edition will be able to run in the background as a Windows service, continually profiling DNS servers, able to generate graphs and charts on demand, to log and dump past historical information, and even automatically move the system's DNS server use to the fastest available servers if the servers being used become slower than others.

And since I hate the model of subscription software with a passion, despite the fact that the rest of the world appears to be going that way, the agreement I'll be making with the purchasers of the Benchmark is that they only ever pay once and they own it and the future of that edition forever without additional cost. If it succeeds as it might, it would create a revenue stream that would justify its ongoing improvement and continuing development as new DNS-related technologies arise.

So substantial new editions of GRC's DNS Benchmark will be my next focus as this podcast moves into 4-digit podcast episode numbering.

AI Vulnerability Discovery

On this occasion of episode #999 of this Security Now! Podcast, I want to take a minute, before we talk about something Google recently announced where AI was used to discover an important vulnerability in a widely used piece of software, to put AI into a broader context.

By now, our listeners have probably correctly determined that I'm one of those in the camp who is overall quite bullish on AI. All of the evidence I've seen and witnessed firsthand informs me that we are, indeed, on the verge of something truly transformative. And I'm very glad I'm still alive to watch this happen. Seriously. My parents and a bunch of my close friends who would have been fascinated by this are no longer here to see this happen. And that's a shame because I believe this is going to be that big. I believe AI is going to be something that changes the entire world.

Like most of those in the "baby boomer" generation, during my lifetime and my awareness, I watched vacuum tubes give way to transistors and transistors to give way to many generations of integrated circuits. Digital memory moved from relays and then magnetic cores to insanely dense electromagnetic and electrostatic storage. Computers evolved from what was essentially an automated calculator many times more expensive than people's homes at the time to incredibly powerful devices that we now discard without a second thought. And the Internet happened during the second half of baby boomers' lifetimes. We've had the privilege of watching this incredible global network interlink the computers we all now casually carry around in our pockets. We are truly living through what was science fiction near the start of our lives.

And now, those of us who are still here, are going to have the privilege of watching AI happen. Given everything I've already watched unfold during my nearly 70 years on this planet, and given what I've seen of it so far, I believe that AI's impact upon our lives is destined to be bigger than anything that has preceded it. More significant than everything that has come before.

For the longest time, the technologies that appeared to have the most impact were those that facilitated communication. The printing press changed the world. And that was followed by the telegraph, which was followed by radio and the telephone which were similarly transformative. The reason the Internet has changed everything again is that it, too, is about communication. It could be argued that automotive transportation is also a form of communication.

Communication has been so universally transformative because it has been about linking the thoughts and intentions of people. By comparison, I believe that AI is going to utterly eclipse the transformative power of communication because it **IS** the thoughts and intentions of people. AI **is** the currency of people.

And, sure, it's easy for cynics and skeptics to find fault. There's always fault to find in the beginning of anything new where big claims about the future are being made. That's just the nature of "new". "New" is the start of the journey, not the end. Personal computers were initially a joke. As were the first luggable laptops. But no one's laughing now. Back at the start of Bitcoin and the invention of cryptocurrency there were many skeptics. But I sure wish I had not installed Windows over my 50 bitcoin. My point is, what AI is today is not what it's going to be tomorrow

– it never is. And I believe we're only at the start of what is going to be more significant than the invention of anything that has come before – because AI is potentially the currency of people, and there's never been anything like that before. I'm glad we're all going to be here to witness it together.

Okay, so what happened with AI and Google? Google has a long posting in their Project Zero blog, but The Hacker News assembled a very nice summary. Here's that they wrote:

Google said it discovered a zero-day vulnerability in the SQLite open-source database engine using its large language model (LLM) assisted framework called Big Sleep (formerly Project Naptime). The tech giant described the development as the "first real-world vulnerability" uncovered using the artificial intelligence (AI) agent.

The Big Sleep team said in a blog post: "We believe this is the first public example of an AI agent finding a previously unknown exploitable memory-safety issue in widely used real-world software."

The vulnerability in question is a stack buffer underflow in SQLite, which occurs when a piece of software references a memory location prior to the beginning of the memory buffer, thereby resulting in a crash or arbitrary code execution. This typically occurs when a pointer or its index is decremented to a position before the buffer, when pointer arithmetic results in a position before the beginning of the valid memory location, or when a negative index is used.

Following responsible disclosure, the shortcoming was addressed in early October 2024. It's worth noting that the flaw was discovered in a development branch of the library, meaning it was flagged before it made it into an official release.

Project Naptime was first detailed by Google in June 2024 as a technical framework to improve automated vulnerability discovery approaches. It has since evolved into Big Sleep, as part of a broader collaboration between Google Project Zero and Google DeepMind.

With Big Sleep, the idea is to leverage an AI agent to simulate human behavior when identifying and demonstrating security vulnerabilities by taking advantage of a large language model's code comprehension and reasoning abilities. This entails using a suite of specialized tools that allow the agent to navigate through the target codebase, run Python scripts in a sandboxed environment to generate inputs for fuzzing, debug the program and observe results.

Google said: "We think that this work has tremendous defensive potential. Finding vulnerabilities in software before it's released means that there's no scope for attackers to compete: the vulnerabilities are fixed before attackers have a chance to use them."

The company, however, also emphasized that these are still experimental results, adding "the position of the Big Sleep team is that at present, it's likely that a target-specific fuzzer would be at least as effective (at finding vulnerabilities)."

While this may be just the first time AI has been deployed for this, my intuition is screaming that AI-driven code verification and vulnerability detection is going to be huge.

It feels to me as though this is right dead center in AI's bailiwick and that it may be that AI is what finally comes to our rescue in the seemingly never-ending fight against both the continuous introduction of new vulnerabilities and the discovery and eradication of old ones. Microsoft must be hard at work figuring out how to use AI in this way. Imagine a day when Patch Tuesday is: *"Sorry, nothing to fix here. No new known vulnerabilities have been found, reported, or known to be under exploitation."* That would be something, and it doesn't seem far fetched. It may be that today's large language model training style doesn't really apply for this. I'm not nearly close enough to AI to know. But I'm sure there are people who are.

Of course, this won't solve all of our problems since there will always be people who are opening dangerous service ports to the Internet – even though the UI AI cautions them not to do so. So I'm not worried that AI is going to put this podcast out of business anytime soon. As always, there are users, and users can always be counted on to do something dumb. But code is just combinatorial math and it's fully deterministic. So it really seems as though code verification would be a natural habitat for AI. And lord knows we really need it.

If I were a younger man that might be where I would aim my research. And I'm serious about this. We often get listeners who are just starting out and who are looking for and asking for some direction. So here's some: It feels to me as though AI could have incredible traction in the field of code behavior verification and software vulnerability discovery. And these days it's possible to borrow big compute resources from cloud providers, which makes basement or garage development not only possible but practical. And if such technology were created it feels like the sort of thing that would be snapped up by any of the big tech giants in a heartbeat.

So... here we are.

We've reached the end of the long-dreaded and then long-awaited podcast #999. I could not be more pleased that we'll be continuing next week into the realm of 4 digits with podcast 1000!