AI Vulnerability Discovery

Description: Google's record-breaking fine by Russia. (How many zeroes 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?

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-999.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-999-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here with lots to talk about. Google's record-breaking fine from Russia? I don't think they plan to pay it. Firefox 132, some nice new features. A really bad exploit involving Windows .RDP files. And then Steve's going to talk about his new product, plans for his next paid product for the first time announced right here, next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 999, recorded November 5th, 2024: AI Vulnerability Discovery.

It's time for Security Now!, the Election Day edition, with Steve Gibson, where we cover all of this - oh, I've got the wrong album art up there. I'll fix that, Steve. We cover all the latest security news, privacy news, and AI news.

Steve Gibson: And look, I voted.

Leo: Lookit, he's got two stickers.

Steve: Yeah.

Leo: You voted twice.

Steve: I got Lorrie's, also. You know, vote early, vote often. That's right.

Leo: I only have the one sticker, but yeah. There's something really satisfying about participating in our democracy, I have to say. And I always enjoy it. And this time was no different, even though I did vote by mail. I didn't go in. It's just satisfying to put that in the mailbox and say, "I participated."

Steve: Yeah. And I do appreciate California making it so easy. Whether you ask for it or not, you get the ballot in the mail, thank you very much. And you get to do it three or four weeks ahead of time and drop it in the box and hope it doesn't catch fire. And then you're done.

Leo: See, it's Security Now!, everybody.

Steve: Ah, very nice.

Leo: I have fixed the album art, which means it's time for you to tell us what's ahead on the show.

Steve: Okay. So before I forget, I've been receiving some emails from people who say, hey, you're mentioning that you sent out the announcement about the podcast the day before. Yesterday it was the evening before to 12,154 people, I believe, or the week before it was the afternoon before. But people are writing saying, I didn't get it. And I got them before, but I didn't get - anyway, what's happening is - I tracked this down. Some people's email services are, in an attempt to protect them from malicious links, are link-following their email, like the links in their email. And unfortunately, my one-click instant unsubscribe really does work. And you just, like, there's no confirmation on it. You click the link, and you're out.

Leo: Bye-bye.

Steve: Bye-bye.

Leo: No warning.

Steve: And so apparently people who are using Outlook, some people using Outlook, Outlook will protect them by fetching the links in their email. Well, when they fetch the instant unsubscribe, that's it. They're not going to get anymore email from me. So everyone who's hearing this who did not get last week's or yesterday evening's email, I'm sorry, you were inadvertently unsubscribed by your overprotective email system.

Leo: Wow.

Steve: So please go back, resubscribe, and by this time next week I'm sure this will be resolved. I will have to take you to a page that asks, oh, you're sure you want to unsubscribe? And then you click yes. And now I understand why that's what everybody else does. You know, I wasn't doing that, and that was not good. It turns out you need to say "Are you sure?" Because then Outlook goes, oh look, isn't that nice, he's asking if

they're really sure. So we're not going to - he's not giving them malware or anything. Anyway, so...

Leo: So it's only Outlook that does this? I feel like Gmail has some unsubscribing [crosstalk], too.

Steve: Well, Gmail makes it very easy. So the idea is there is a standard where the one-click unsubscribe goes in the headers.

Leo: Right.

Steve: I was also putting it in the body of the mail. And so that's what Outlook is going and finding are things that users could click on that might get them into trouble. Oh, and boy, do I have a really amazing piece of news about that here today.

Leo: Oh, good, okay. Okay.

Steve: But in the headers, that's where Google will say, if you mark something, if you flag something as spam, you get a little dialog that says, oh, would you like to - or maybe it's if you just delete them. Anyway, you get an offer from Gmail to unsubscribe from this list.

Leo: Yeah, I've seen that, yeah.

Steve: The reason it knows it's a list is that in the headers, unseen by users unless they explicitly look, is an unsubscribe link which your email provider is able to use. And in fact this was really handy when I was doing the mailing to the really old email addresses. Many of the recipients, the recipient servers, would see, oh, that account hasn't been around for a decade, and so the server itself would unsubscribe them. So it was great. It sort of made, you know, it like automatically cleaned the email list for me.

So anyway, we are at Security Now! Episode 999, as you mentioned, Leo. And I should mention to our listeners, last week I noted I had not yet updated GRC's side to handle four-digit podcasts. I did that.

Leo: Yay. You had to do that.

Steve: So we are prepared to wrap into - I had to do it.

Leo: This was the last week.

Steve: This was the last time, that this was why it was going to be all over. And, boy, I didn't even realize it was going to be on Election Day when lots of things might be over.

Leo: I think you had some prescience there, saying, you know, I think my last show should be Election Day 2024.

Steve: That's right.

Leo: Wow, did you pick that one.

Steve: Okay. So we're going to talk about the interesting topic of AI being used for vulnerability discovery, which I think is going to be a big deal. And I don't want to step on my own story here, so I'm just going to leave it at that until we get there. We're going to talk about Google's record-breaking fine by Russia and wonder how many zeroes does that number have? Also Russian Television, RT's editor-in-chief admits that their hosts are AI-generated.

Leo: Oh, wow.

Steve: Yeah, probably because they sent all the actual hosts off to war. Windows 10 security updates are set to end - that's for 22H2 - set to end next October. Or are they? When a good Chrome extension goes bad. We're going to look at a real-world event that occurred. Also, Windows it turns out will launch RDP sessions, you know, Remote Desktop Protocol sessions, with a .RDP launch file which can also configure your RDP client for full zero security. And we ask what could possible go wrong with that? Actually, something has.

Leo: Oh.

Steve: Firefox 132 just received some new features. Chinese security cameras have been removed, well, more than half of them from the UK. We'll check in there. And I know our listeners would not fall for this social engineering attack we're going to look at, but I bet you that lots of people would. Also I'm going to announce GRC's next commercial software product, or at least semi-commercial software product, and talk about that a little bit. And then we're going to look at the prospect of AI, as I said, being used to analyze code to eliminate security vulnerabilities. Much as I recently suggested that AI running on the local smartphone may be the solution to allow us to preserve full end-to-end encryption by preventing bad stuff from being sent or received, I bet you that AI may be the solution to the security problem. And, oh, Leo, have we got a Picture of the Week, a goodie.

Leo: I love it. All ahead, Security Now! 999 is underway. Which would have been, if you're just joining us, the last Security Now!, until Steve changed his mind a year ago. You were ahead. You were ahead of it.

Steve: For many years people were fretting, and I was planning. But, you know, no. I'm not ready to go yet.

Leo: Congratulations. Congratulations. That's great. All right, Steve. I have prepared myself, steeled myself, if you will, for the Picture of the Week.

Steve: Just hold onto your desk.

Leo: It looks like something out of Alfred Hitchcock's "Psycho." Wow.

Steve: So I gave this one the caption, "When handrails are not optional." And I truly wonder whether you could walk down these stairs without, like, having your...

Leo: You'd have to close your eyes. The stairs are normal; right? I mean, they're not abnormal, but they sure look that way.

Steve: Correct. The stairs are completely normal. Someone put the worst imaginable pattern of carpeting on these stairs. It's all full of, like, off-axis, cross-wise - it's horizontal stripes, but they're all kittywampus is the technical term.

Leo: Yes, it is.

Steve: And, oh, I mean, you have to really focus in order to get down these things. So anyway, I've had this one for a while, and I thought it was great. You know, you could see the aisle that it goes down to is the same pattern, and that's going to be okay. But, boy, when it turns around and goes 90 degrees and goes up the stairs...

Leo: This looks like it's on a ship, too. I don't want to make it worse, but imagine you're rocking on this thing.

Steve: Wow. Not good.

Leo: Wow. Not good.

Steve: Okay. It's a shame that our favorite Russian Internet watchdog, Roskomnadzor, is not the Russian entity that's been levying these fines against Google over its management of YouTube, since it would have been fun to say that name many more times during this reporting. We only get that once. 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, that's 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 in The Moscow Times under the headline "Russia Fines Google $2.5 Decillion U.S. Dollars Over YouTube Bans."

They wrote: "The RBC news website reported Tuesday that Google has racked up some two undecillion rubles (which is the equivalent of $2.5 decillion U.S. dollars) worth of fines in Russia after years of refusing to restore the accounts of pro-Kremlin and state-

run media outlets." You know, it's like Google just said, no, we're going to kick this off YouTube. "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" - you know, Russian lawsuits - "against the company for blocking their YouTube channels. Those daily penalties" - get this - "have doubled each week." And, you know, when we're young we learn about the power of compound interest; right? So these penalties are doubling each week, "leading to the current overall fine of around two undecillion rubles."

Now, "undecillion," they explain, "is a number equal to one followed by 36 zeroes, or one trillion trillion trillion rubles. Google, whose parent company Alphabet," they report, "which reported revenue of more than $307 billion in 2023, is unlikely" - you think? - "to ever 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," they write, "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, but okay.

"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'" - it's also, well, it's filled with zeroes. "'In fact, this should be a reason for Google's management to pay attention to this and fix the situation.'" You know, Google's management doesn't care. Anyway, finally they said: "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 all advertising in Russia in order to comply with Western sanctions over the war in Ukraine." So, yeah, fine them all you want. Double it every week. You're going to run out of zeroes at some point.

And as I also noted at the top of the show, this editor-in-chief's name, Margarita Simonyan, was mentioned as one of the other 17 companies that have also filed more recent suits against Google. I had noted that she also recently admitted that many of RT's, you know, Russian Television's hosts do not exist and are entirely AI-generated, along with their fake social media accounts because I guess you've got to, you know, if you want to respond to them interactively, get all engaged, they need to have a social media account to allow you to engage with them, with their fake AI hosts. Anyway, she predicted that journalism would disappear in the near future. You know, it already has in Russia, so maybe she thinks that's going to spread. Unfortunately, she may be right. We'll see.

A recent posting to the - and this is important for all of our listeners, unlike that first one that was just a little bit of junk food. A recent posting to the 0patch blog regarding next year's end of Windows 10 security updates contained a bunch of interesting related news. This included what Microsoft plans to charge end users who would rather remain on Windows 10 come next October, or may not be a matter of rather remain, they may have no choice due to what we know are Microsoft's arbitrary minimal system requirement policies for moving to Windows 11. So 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!"

So they wrote: "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 being used securely will be to upgrade to Windows 11. Many of us don't want to, or simply can't, upgrade to Windows 11." They wrote: "We don't want to because we got used to the Windows 10 user interface, and we have no desire to search for some button where it's been moved 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" - and this is their word in the blog posting - "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, point-of-sale terminals, special-purpose devices, ATMs that run on Windows 10 and cannot be easily upgraded. And finally, our hardware may not qualify for an upgrade to Windows 11. Canalys estimates that 240 million computers worldwide" - 240 million computers worldwide - "are incompatible with Windows 11 hardware requirements, lacking Trusted Platform Module (TPM v2) 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," they say. "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," they said, "to Windows 7 in January of 2020. Today, a Windows 7 machine last updated in 2020 with no additional security patches would be really easy to compromise, as over 70" - seven zero - "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 the network.

"For everyone else, there are two options to keep Windows 10 running securely. Option 1: Microsoft's Extended Security Updates." They wrote: "If you qualify, Microsoft will happily sell you Extended Security Updates, which means another year or two or even three of security fixes for Windows 10, just like they've 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" - that is to say twice that - "for the second year, and $244" - doubling again - "for the third year of security updates, totaling $427 for every Windows 10 computer across three years." That's, you know, for the enterprise.

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 enterprise users every year. But not for end users, who can apparently maybe, it's not clear to me, maybe just pay for one year for $30 and then that's supposed to be enough of a bitter pill...

Leo: Then you're out of luck, yeah, yeah.

Steve: ...that you're pushed off to Windows 10. So they continue, 0patch says: "Opting for Extended Security Updates will keep you on the familiar monthly 'update and reboot' cycle. And if you have 10,000 computers in your enterprise network, it will only cost $4 million." They said: "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'" - their phrase - "Windows 10 v22H2, the final release of Windows, and provide critical security patches for it for at least five more years, longer if there's a demand in the market."

They wrote: "We're the only provider of unofficial security patches for Windows, and we've done this many times before. After security-adopting Windows 7 and Windows Server 2008 in January 2020, we successfully took care of six versions of Windows 10 as their official support ended, security-adopted Windows 11 21H2 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 unapplied if they're suspected of causing a problem. Again, no rebooting or application relaunching. 0patch 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," which are, they explain, "patches for vulnerabilities that have become known, and are possibly already exploited, but for which no official vendor" - that is to say Microsoft - "patches are available yet. We've fixed many such zero-days 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 become available."

Then there's "'Wontfix' patches, patches for vulnerabilities that the vendor" - again, Microsoft - "has decided not to fix for some reason. The majority of these patches currently fall into the 'NTLM'" - you know, NT LanMan - "'coerced authentication' category. NT LanMan 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."

And, finally, non-Microsoft patches. They wrote: "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 the Java Runtime, Adobe Reader, Foxit Reader, 7-Zip, WinRAR, Zoom for Windows, Dropbox app, and Nitro PDF.

"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 versions of Windows such as 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 NT LanMan 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 admin account, currently priced at 24.95 euros plus 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, et cetera, currently priced at 34.95 euros plus tax per computer for a yearly subscription." And they conclude: "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 or relevant. We know from our many years of covering 0patch, 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. So in this instance, I don't mind this sales pitch because it's easy to endorse what they're selling. Microsoft has clearly 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 to this new operating system, while older problems are still being resolved. And let's not even get started on the fact that Microsoft's Replay is an issue for Windows 11 users.

So considering 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 euros per year, which at the moment is around 27 USD per year, extending the security coverage of that beloved platform for a minimum of another five years, starting in October 2025, 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. So anyway, I just wanted to let everybody know...

Leo: I have questions.

Steve: Yes, good.

Leo: I have some questions. So first of all, 0patch, it sounds like, is patching in memory, not on the drive.

Steve: Yes. Yes. You can't patch on the drive because that would break the signature of the files.

Leo: Ah, right.

Steve: And so they would never load.

Leo: So you have something running all the time that's the 0patch tool that just loads in patches as needed.

Steve: Yes. Yes. There is a 0patch agent.

Leo: Okay.

Steve: Which is small and runs. And when we've talked about this in the past, the patches are literally 23 bytes. I mean, they're like, there are a few instructions where they just fix the problem. You know?

Leo: Yeah, they just keep it. So all of the patches are their own They are - how do they get - so Microsoft's releasing security patches, and 0patch is duplicating those patches Do they reverse-engineer them? How do they know?

Steve: Just like the bad guys do. In the same way that the bad guys do a delta on the pre- and post-patch code...

Leo: That's all you have to do, I guess, huh.

Steve: Yeah. You just find the thing that Microsoft changed.

Leo: Right. Not why.

Steve: And so basically - yeah.

Leo: What is it, yeah. Okay. That's - it's an interesting business, actually.

Steve: I think it's a great business. And I mean, they've been around for a long time. If you search GRC's transcripts for...

Leo: Oh we've been talking about them for years, yeah.

Steve: Yes, for 0patch because they often jump in before Microsoft has an update. And they don't charge you anything for an update which has not yet been officially patched. So where they're filling, I mean, just as a public service, where they're filling an emergency need that Microsoft has not filled for something being exploited in the wild,

you can get that from them for free. I mean, they're like Cloudflare in just having this feeling of being really good people.

Leo: Well, they are going to sell it down the road, which is good. That's fine. You know, they're putting a lot of work into it.

Steve: Yeah, 24 bucks for a year of protection? Many people would rather do that than be forced to use Windows 11.

Leo: Are you running it? Have you run it?

Steve: No.

Leo: No.

Steve: Because I don't believe any of this nonsense about you can't run old versions of Windows. I'm running Windows 7. I'm just fine.

Leo: Those 70 vulnerabilities don't bother you.

Steve: No. I just don't go to bad places, you know. My site doesn't have any. And I've got up-to-date browsers. Browsers are the big vector, the way stuff gets in. And oh, boy, Leo, wait till you see one of the ways, a new way that people are being tricked.

Leo: Oh, yeah.

Steve: Oh. Let's take a break, and then we're going to talk about what happens, a case in point of good extensions going bad in Chrome.

Leo: Okay. Deal.

Steve: I recommend 0patch. I think everybody who's listening should take a look at it. If the idea appeals to them, I don't see a downside.

Leo: And, mean, it keeps you running for as long as your apps continue to be secure. I mean, ultimately that's what breaks it is, you know, the browser is no longer supporting Windows 10 or something like that.

Steve: Right.

Leo: Very interesting. Steve, back to you.

Steve: Okay. So we have another example of a popular Google Chrome extension with more than 100,000 daily users suddenly becoming malicious. The extension known as Hide YouTube Shorts has been found to be performing affiliate fraud and collecting and transmitting the browsing history of every one of its users.

Leo: Find YouTube Shorts?

Steve: Hide YouTube Shorts.

Leo: Hide your shorts. Okay.

Steve: That's right.

Leo: Okay.

Steve: And apparently that's a thing. Anyway, I'll [crosstalk] in a second.

Leo: Okay.

Steve: So security researchers say that the extension appears to have turned malicious, not surprisingly, we've talked about this a lot, after it was transferred to a new developer. I went over to the Google Play Store to check it out. Now, 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 similarly suggest that they do that also. 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. He 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 didn't notice any suspicious activity. It was only after I started debugging the extension that I noted 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 and parameters, which allows reading the information in the address bar, including, for example, search history and search terms. Some users in the reviews on the extension page in the Chrome Web Store also indicated the possibility of redirecting, that is, being redirected 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 the browser extension, 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. Yeah. 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. And again, 100,000 users per day.

"The extension was originally developed," he wrote, "by a single developer who maintained the source code on GitHub; however the GitHub repository was archived on September 12th, 2023, and the plugin was acquired, or maybe sold, to another developer." He said: "I have not 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," he wrote, "I analyzed its behavior and did not see similar problems with it." So indeed this did happen downstream at some point.

He finishes: "I have no doubt 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 once again, the caution would be, you know, our takeaway from this would be to attempt to minimize the use of browser extensions. We know that by, you know, by far for the most part, extensions developers are well meaning and acting aboveboard. 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 might be malicious.

And I just haven't had any time to dig into uBlock Origin further, but I've got this nagging sense that, for example, if you wanted to block YouTube shorts, uBlock Origin would just do that by turning on, by using the dropper and clicking on, like, something in YouTube shorts, and they would just go away. I've had anecdotal reports of that in feedback from our listeners. So you probably don't even need more special purpose extensions. You probably just need to better utilize uBlock Origin. At some point I'm going to make time to do that for us because...

Leo: It's just a css div probably that you could, you know, if you knew the name of it, you could just block it automatically.

Steve: Exactly that.

Leo: Yeah, yeah.

Steve: And in fact that little - the little dropper thing finds that for you.

Leo: And fix the div, yeah.

Steve: It just, yes, exactly, and just does that, and creates a rule.

Leo: Yeah.

Steve: So anyway, the fewer the better when it comes to extensions. Okay. This is one. Oh, boy. We all know the trouble Windows has had, over and over and over, over something as simple as .LNK link files.

Leo: Oh, yeah.

Steve: I mean, that, Leo, you were covering these before the Security Now! podcast on your weekend show.

Leo: Anything you double-click that does something is always risky; right?

Steve: Uh-huh. So the exploits of those have been epic, you know, and we've lost count of the number of times they've been "fixed," in air quotes, only to rear up again. You know, some design concepts are just bad and are notoriously prone to abuse. And Leo, you just summed it up. Anything you can double-click, that's a problem. So 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 never learned anything from the past. And as we know, those who do not learn from the past are destined to repeat it.

Okay. So the generic tech press reporting on this just said: "Microsoft says that a notorious Russian cyberespionage group is using a clever" - okay, 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 onto the compromised device."

Leo: But it's convenient.

Steve: It's so simple.

Leo: Yeah, so simple.

Steve: "Microsoft has attributed the operation to Midnight Blizzard." Remember they're the people who got into their email also. They don't like the Midnight Blizzard people.

Leo: No, they don't.

Steve: "A cyber unit inside Russia's SVR Foreign Intelligence Service. The group has used the new technique since October 22nd and has targeted individuals in government, academia, defense, and NGOs across the U.S. and Europe. This is the same campaign that was spotted by AWS and CERT-UA."

Okay. Now, since the inherent insecurity of this entire design was just too much to believe, I went to the source, where Microsoft themselves explain. They said: "On October 22nd, 2024, Microsoft identified a spear-phishing campaign in which Midnight Blizzard sent phishing emails to thousands of users in over 100 organizations. These emails were highly targeted, using social engineering lures relating to Microsoft, Amazon

Web Services, and the concept of Zero Trust. The emails contained a Remote Desktop Protocol (RDP) configuration file signed with a Let's Encrypt certificate." Because you can get those for free.

Leo: Why not, yeah.

Steve: "RDP configuration (.RDP) files," they wrote, "summarize automatic settings and resource mappings that are established when a successful connection to an RDP server occurs." Imagine that. Let's make that easy. Let's make it one click. "These configurations extend features and resources of the local system to a remote server, controlled by the actor." Where we insert, what could possibly go wrong?

Leo: I'm sorry, I missed my cue.

Steve: It's okay. We'll have a few more by the time we're done here.

Leo: Oh, good.

Steve: "In this campaign, the malicious .RDP attachment contained several sensitive settings that would lead" - yeah, like let's map the C drive - "that would lead to significant information exposure. Once the target system was compromised, it connected to the actor-controlled server." Oh, and by the way, where they say "was compromised" they're being quite kind. By that they mean when the user received the email containing the .RDP extension and clicked it. That now qualifies as you have just compromised your computer, baby.

Leo: Oh, geez.

Steve: Because you've clicked on a file that your email wasn't trained to block. Notice that you can't send EXEs anymore. Those die an immediate death if you try to email someone an EXE. There's just no hope. But RDP, yeah.

Leo: I would submit that your computer was compromised the minute you enabled RDP, that that's...

Steve: Well, it's enabled by default, and that's another one of those, here we go, what could possibly go wrong?

Leo: I didn't miss that one.

Steve: Okay. So as they say, "Once the target system was compromised" - meaning the user clicked on something in email, which is all it takes to compromise Windows these days - "it connected to the actor-controlled server and bidirectionally mapped" - this is Microsoft - "and bidirectionally mapped the targeted user's local device's resources" -

meaning hard drives - "to the server." Bidirectionally mapped means not only can, you know...

Leo: It can read it and write it.

Steve: That's right.

Leo: Wow.

Steve: "Resources sent to the server may include, but are not limited to" - this is Microsoft saying this - "all logical hard disks, clipboard contents, printers, connected peripheral devices, audio, and authentication features and facilities of the Windows operating system, including smart cards." Basically you've just given them access to your entire system.

Leo: Everything. Everything. Yeah.

Steve: And Microsoft wrote: "This access could enable the threat actor" - okay, the only way it wouldn't is if they were literally asleep when this mapping occurred, otherwise, oh - "could enable the threat actor to install malware on the target's local drives" - actually, it's probably automated, and so they can be asleep, and it'll happen in their sleep - "and mapped network shares, particularly in Auto Start folders." Oh, so they have those, too. "Or install additional tools such as remote access trojans 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 signed-in user to the target system." This, again, Microsoft writing.

"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." Right? Protected by Recall. Don't worry. You're safe. Oh, right. Windows Hello, not safe. "Passkeys or security keys; clipboard data; point of service, also known as point of sale or POS devices." And they go on and on and on.

In their blog posting, Microsoft 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? Which is, I think, Leo, the first thing you suggested upon hearing this.

Leo: Yeah, there you go. Yeah.

Steve: 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 to a maximally insecure state and connect to a previously unknown remote server?

Leo: Well, it's there for - it's for, like, remote support; right? Yeah.

Steve: Well, I use RDP extensively. And, yes, RDP saves its connection profile settings into individual .RDP files, and that can be 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're 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 of in individual RDP files, and then this problem would not exist.

Handy as it inarguably is, there's just no safe way to send somebody, anybody, 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. There's no safe way to do that. You know, at the very least this facility should be firmly disabled by default for everyone, and then only those few people who actually need to do this should then be forced to jump through some hoops to enable it on their machine only, and even then possibly only for some self-limiting time. And if that were the case, Russia would have never bothered to create this because it would be off for 99.999999% of the people in the world.

I hope everyone knows to never click on anything received in an 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 unfortunately clever bad guys keep finding new ways to abuse this useful capability.

Leo: But, man, an RDP link is so powerful. Now, I don't allow port 139 on my router. Most people probably don't. But I guess because it's an outbound request your firewall's not going to stop it.

Steve: Yeah, doesn't matter. And it runs on 6800, or it runs on a high port number.

Leo: Oh, okay.

Steve: As I recall, also.

Leo: But it doesn't matter because you're outgoing, saying, hey, Russian server, come on in.

Steve: Yeah. And you can bet that Russia has their port wide open and listening.

Leo: They're open.

Steve: For anybody to connect.

Leo: Oh, man.

Steve: And Leo, this started on October 22nd, meaning that - and thousands of emails went out to hundreds of companies, highly targeted, looking legitimate. People clicked on them, and they got themselves immediately compromised. That's how bad guys then get a foothold inside an enterprise. And talk about a foothold. I mean, this is...

Leo: You've got everything. You've owned it.

Steve: This is a body hold.

Leo: Yeah. You own it.

Steve: Yes.

Leo: Wow.

Steve: And speaking of owning it, Leo, let's give our listeners a chance to own something. And then we will continue.

Leo: You're not anxious to get to some other...

Steve: [Panting]

Leo: I have the TV on here, Steve. You're not missing anything.

Steve: That's not fair. Okay.

Leo: There's nothing going on until...

Steve: No polls are closing on the East Coast.

Leo: You've got at least an hour before Georgia closes, so you're good.

Steve: Yeah.

Leo: This is the fastest-paced show we've ever done. I can't keep up. Whew. Okay. We're going to have some more great stuff coming from Steve, as always. Steve's amazing with the quality of the information you get here. Steve?

Steve: Okay. We've got a new Firefox. We're now at 132. It 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 they also block favicons if they're loaded via

HTTP. Back when we were looking at Firefox's third-party cookie handling, there was a great deal of confusion since Firefox's UI - we talked about it at the time on the podcast - Firefox's UI and its behavior, its actual demonstrable behavior appeared to be at odds with one another.

So among the improvements that we got in 132, I was pleased to see the sentence: "Firefox now blocks third-party cookie access when Enhanced Tracking Protection's Strict mode is enabled." So that's what everyone thought it was doing, but we saw that it wasn't. It is now. So as we suspected, you know, GRC's cookie forensics system showed what was happening, and that's been fixed in Firefox 132, which everybody probably has.

As I mentioned at the top of the show, under the sad but understandable category of "we don't trust camera-equipped black boxes made in China," we have the news...

Leo: Really.

Steve: Yeah.

Leo: Okay.

Steve: We have the news that the - we talked about DJI drones in one example of camera-equipped black boxes. We have the news that the UK government now says that over 50% of all Chinese-made security cameras have been removed from sensitive sites, such as government buildings and military bases. The government says it expects removal to be completed by April of next year, 2025, despite the fact that the removal was initially ordered well back in November of 2022, as we covered at the time. And I was thinking, wow, you know, it took them until now to get rid of half of them?

But then I thought, okay, there's probably a long procurement cycle for such things, so it took some time to get the replacement cameras in the pipeline. And as we know, UK officials ordered all sensitive sites in the UK to remove all Chinese-made cameras, citing national security concerns because anything is possible. And basically that's it; right? No evidence, but anything's possible. So, yeah, I think certainly for sensitive installations that makes sense.

Leo: I'm not sure I would announce that, oh, we've removed half of them.

Steve: Yeah. Start using the other half before they...

Leo: Hey, good news, half of them are gone.

Steve: That's right. Okay, now, Leo.

Leo: Yes.

Steve: Okay. And I know that our listeners are savvy.

Leo: Yeah.

Steve: I was first 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 and horrific form...

Leo: I think I would fall for this. I hate to say it.

Steve: Again, I can see people, like, I know lots of people who would, definitely. A very clever form of social engineering attack, and I think it might ensnare many non-suckers. So it's not the sucker born every minute, it's that, you know, maybe it's a little more than do you have a pulse, but still, not much. Okay. It leverages the fact, the true 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, right, they just follow the instructions and do what's asked of them and hope for the best. And that's why I can understand why this new and rather blatantly obvious to techies exploit is actually succeeding out in the wild. And it's horrifying to contemplate.

Okay. It begins with a faked CAPTCHA pop-up which, of course, we're all seeing now. So it starts...

Leo: See them everywhere.

Steve: You get something you expect to see; right? Like, okay, I'm going to have to prove that I'm not a robot.

Leo: It even says ReCAPTCHA, which is legit.

Steve: Right, right. So in this case someone - in this case it was used where somebody wishes to watch a video. They need to click on the CAPTCHA button to start authenticating that they are human. Okay. But this click that the user makes actually runs, it's created by JavaScript, and it runs a bit of JavaScript which places a dangerous Powershell executable string onto their Windows clipboard.

Leo: Oh, my god.

Steve: JavaScript is able to read and write the clipboard. So when you click on this, it puts this Powershell script onto your clipboard, and it uses an encrypted command tail that Powershell will decrypt. So it just looks like gobbledygook, like okay, whatever. Okay. After pasting this trojan-invoking Powershell script onto their clipboard, it then displays the remaining instructions they must follow to ostensibly prove their humanity. Okay, well, they are definitely about to prove their humanity, but not in a way that they intend. Get this. The pop-up reads "Verification steps: Press Windows Button," and then it shows you that little Windows, you know, four Window pane icon, plus R.

Leo: Oh. I wouldn't fall for this part.

Steve: I know. Again, okay, but we know people who would; right?

Leo: Sure, because most people don't know what Windows+R and CTRL+V and ENTER do.

Steve: Don't have any clue what any of this is about.

Leo: Right.

Steve: Step number two, press CTRL+V. Step number three, press ENTER.

Leo: Step number four, what could possibly - wow, wow.

Steve: Okay. So Windows+R brings up the Windows Run dialog with its, you know, "what would you like me to run" field highlighted.

Leo: Right.

Steve: CTRL+V pastes this horrendous Powershell EXE command into the system's clipboard, well, from the system's clipboard into that Run field so that the Run field now contains the executable Powershell script to download and install and run trojan malware on their computer. And then this all culminates when they follow the final instruction of pressing ENTER to, as Picard would say, make it so.

Leo: Oh.

Steve: Again, as I observed, none of us would do this. But again, most people don't know what any of this is.

Leo: Right.

Steve: So they're just following the steps because they want to see the video; you know? They want the carrot. And so, wow.

Leo: Fortunately, Windows+R does nothing on a Macintosh, so I'm safe.

Steve: You're safe. Oh, you in the minority.

Leo: The minority's growing. And it's because of things like this I'm convinced. But okay, go on.

Steve: Wow, yeah. So anyway, I don't know what to tell our listeners. I know none of our listeners would fall for this, but I know they know people who would.

Leo: Oh, yeah. Oh, yeah.

Steve: So, you know, wow. It's bad enough to be forced to click things, like forced to click things in your browser when it could be a spoofed window. Our browsers are designed to try to minimize the damage.

Leo: Right.

Steve: But it's possible for JavaScript to put something on our clipboard. And then these instructions basically say, oh, thank you, here's what we want you to do now. And it involves getting that thing to run, which those keystrokes will do. Wow.

Okay. I said last week that I wanted to announce the next big thing I'm working on.

Leo: Oh, boy.

Steve: I recently finished the work on GRC's email system. And actually I have a caveat to that now, as I said, because it turns out that Outlook is doing link following to protect people from malicious links, and in the process unsubscribing people from their mailing list. So I'll fix that in the next day. And then it's on to what comes next. Oh, and I forgot to mention last week, one of the system's, the email system's originally missing features was the capability to allow its users to easily update and migrate their email addresses at any time they may want to. 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 their old account under their old email and then create another one under their new email. So not really a need to explicitly rename their existing account.

But after I saw very high spam complaint rates when initially mailing to SpinRite's owners from 20 years ago, who were like, what the heck is this, I migrated SpinRite's purchase data into the email system, which allowed me to send email which opened with the line, "Back in 2005, someone named Joe Schmo at this email address purchased SpinRite." And as I mentioned at the time, that had a profound effect upon the spam complaint rates. Suddenly everyone was like, oh, yeah, I remember that. Anyway, now the email system is able to handle updates.

The email system knows about SpinRite owners, so there is more actual data contained in an account, and I'd like to keep it there. So I've added a simple "rename" field to the email management page which any of our listeners will see next time they go there, like to resubscribe to the Security Now! podcast, which they were just mistakenly unsubscribed from. So I wanted to let everyone know that, since they last visited the email management page, editing has been added.

Once that was done, I was then able to address the final remaining loose end of the SpinRite 6.1 documentation offering, which was to create a video walkthrough demonstration showing 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 now exists. I posted it on my YouTube channel. I posted it over on GRC. So it's hard not to find it. And if anyone is curious, there you go.

And that brings me to the announcement that I teased last week. As I've mentioned a number of times, GRC's number one by far, I mean, far, 9.3 million downloads so far, most popular software of all time is the DNS Benchmark. I have been astounded by its popularity. When I was putting the show notes together, I guess it was Sunday, it had been downloaded 9,313,642 times, at around 1,600 downloads per day. The Benchmark pages have a page that solicits feedback, and I am constantly receiving requests for new features. Mostly people are wondering how the speed of encrypted and privacy-protecting DNS using encryption - DoH, DoT or DNSCrypt - compares with regular plaintext DNS. Is it slower? Is it faster? What?

And despite the glacial progress of IPv6, as we talked about last week, many people are requesting that I add support for IPv6 to the Benchmark. And actually I think that makes sense because, when IPv6 is available, our systems use it preferentially. So you may be using an IPv6 DNS server which the benchmark won't benchmark. So other great ideas have been to allow the Benchmark to verify the domain filtering being done by services like NextDNS, and others have been wishing to avoid local domain name blackouts where the DNS services they're using don't let them access sites they want to, so the Benchmark could be used to help them locate servers that would allow them to get access to those sites.

So anyway, the other thing I hear more generically is that people would like to have a way of supporting my continuing work here on all things GRC. You know, newsgroups, forums, ShieldsUp!, DNS Spoofability tests, all the freeware that I write and am able to offer, and everything else.

So I've decided that my next project, before I create "Beyond Recall" for super-fast, super-secure data deletion, which will precede the development of SpinRite 7 for Windows, will be to revisit the DNS Benchmark and to give it a major version 2.0 update. There will still and always be a free release available, like it is now. But I would like it to be able to support itself, if it can. And I think it should be able to, based upon its observed popularity. So I plan to offer all those new features for $9.95 in a "Plus" edition; and also, for the real DNS pro guys, a "Pro" edition for $19.95, which will do a whole bunch more, run as a service, background logging, lots of long-term charting, and a bunch of other stuff.

Leo: That sounds great.

Steve: So anyway, that's the plan.

Leo: Count me in. When is it available? I'll buy it now.

Steve: Well, and that's my hope is that I'm going to, because it's an update to an existing product, it's not going to be a long time coming.

Leo: Right.

Steve: 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 its future of that edition forever without ever any additional cost. So if it succeeds, as it might, it would create a revenue stream that would justify its ongoing improvement over time and continuing development, you know, as new DNS-related technologies arise.

So anyway, I will have a substantial new - a pair of, you know, an upgrade to the freeware that'll still be available, and then for people who want more, you know, for less than 10 bucks - well, not much less, 9.95 - you can get that and own it forever and its entire future. So that's my [crosstalk].

Leo: Smart to have the 9.95 and then the next one up because I know that everybody looking at that's going to go, well, for 10 bucks I can get Pro, but I want the super-duper edition for 20 bucks because 20 bucks is not...

Steve: Yeah. And actually I got that thought from John Dvorak, who - he and I talked, like, just sort of - yeah, oh, he wrote to me, and then we ended up having a couple hour conversation because he wanted to know what email system I was using because he was leaving monkey mail, whatever that thing is called.

Leo: Chimp Mail.

Steve: Anyway, and the point he made was he said, you know, don't put a cap on what people can pay you because they might want to pay more.

Leo: He's done very well with that, I might add. Good. All right.

Steve: Okay. So let's take our last break.

Leo: Yes.

Steve: And then we're going to talk about AI's application in security vulnerability discovery. And I have an Episode 999 sort of editorial to lead in on that with.

Leo: Oh, good. All right.

Steve: So good stuff.

Leo: The good news is, 999's not the last.

Steve: Indeed not.

Leo: Next week for Episode 1,000.

Steve: 1,000.

Leo: Or are you going to do it in hex? I don't know what he's going to do. What would that be? I don't even know. Okay, Steve. Vulnerabilities.

Steve: On the 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, I'm sure our listeners have 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, frankly, alive to watch this happen. Seriously. You know? My parents...

Leo: It is very science fiction futurism; isn't it. I mean...

Steve: It is, and it's happening.

Leo: Yeah.

Steve: You know, and 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, I think, because I believe this is going to be that big. I believe AI is going to be something that changes the entire world.

Leo: Wow.

Steve: Like most of those in the baby boomer generation, during my lifetime and my awareness, I've watched vacuum tubes give way to transistors, and transistors give way to many generations of integrated circuits. Digital memory moved from relays, and then to 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's 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, as I said, 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. That's what I want to share. Here's what they wrote. They said: "Google said it discovered a zero-day vulnerability in the SQLite open-source database engine using its large language model-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 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 Hacker News said: "The vulnerability in question is a stack buffer overflow 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 a 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." And I'll also note that that made it, you know, it was a newly introduced bug that this thing immediately found.

They said: "Project Naptime was first detailed by Google in June of 2024 as a technical framework to improve automated vulnerability discovery approaches. It has since developed into Big Sleep, as part of a broader collaboration between Google Project Zero [yay] 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.'"

And The Hacker News finishes: "The company, however, also emphasized that these are still experimental results, adding that '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.'"

Okay. So while this may be just the first time AI has been deployed for this, my own intuition is screaming that AI-driven code verification and vulnerability detection is going to be huge. To me, it feels as though this is 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 and apparently intractable 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."

Leo: Now you're just fantasizing.

Steve: That would be something, yeah.

Leo: Wouldn't it be something?

Steve: Yeah. And it really, to me, it's impossible for us to reach if we don't do something like this.

Leo: Yes.

Steve: With AI, it does not seem that farfetched. It may be that today's large language model training style doesn't really apply for this. That's my feeling. I don't think that's the way to attack this. But I'm not nearly close enough to AI to know. But I'm sure there are people who are.

Of course, you know, this won't solve all of our problems since there will always be people who are opening dangerous service ports to the Internet, or following instructions in a believable-looking CAPTCHA, telling them to just bend over.

Leo: Just copy this, yeah, paste it in.

Steve: Yes. And, you know, even when their UI's AI cautions them not to do that. 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 dumb things. I think that was Pournelle, something like that; right? He was famous for citing that. But code, code is pure. It's why I love it so. It's just combinatorial math, and it's fully deterministic. So it really seems to me as though code verification would be a natural habitat for AI. And lord knows we need it.

If I were a younger man, that might be where I might aim my own focus. 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 is 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 think about that. If you're young and, you know, full of future, and you're looking for something to sink your teeth into, I have no idea how you would do it. But I guarantee you that in a decade, and I'll still be here watching this stuff happening, I will guarantee you this is going to change. AI, I think, is going to be what solves our end-to-end encryption problem, as I said last week, because it's going to give governments the warm and fuzzies that abuse of children can no longer get past the AI monitoring their device locally.

Leo: Oh, interesting, yeah.

Steve: And I think AI is going to be the thing that solves, like, our endless software vulnerability problems. It's a big problem; but, you know, what fun.

Leo: Hey, if it can do that, there's probably a lot of other things that AI will be up to, as well.

Steve: Oh, it's going to revolutionize medicine, Leo.

Leo: Yeah.

Steve: It's going to revolutionize drug discovery. And it's, I mean, it is going to change the world.

Leo: Yeah. And by the way, this is - I loved how you started because I think this is exactly what you and I, who have watched many changes in our lifetime, are hoping for one last big one.

Steve: Yes. This is it.

Leo: And this could be the big one. This could be the one that changes humanity and launches us into an entirely new realm. I kind of agree with you. So I'm excited, too. That's Steve Gibson, GRC.com. He's got a new product coming. Now, timeframe? You don't like to do that.

Steve: I can't guess. A couple months probably. I'm hoping a couple months.

Leo: Put me down for one of those $20...

Steve: Thank you. I will.

Leo: ...subscriptions because I'll be the first in line to get it, I can see [crosstalk].

Steve: I can't wait to find out how encrypted DNS compares to un. I have no idea.

Leo: Yes. Yeah, you'll have fun with this. Or IPv6 or what OpenDNS, what NextDNS is doing, things like that. This will be really useful.

Steve: Yeah. And because the Pro version - so there's Plus at 9.95, and it has all the features, except the Pro can run as a service.

Leo: In the background.

Steve: Because it's all written in assembler, it's a couple of hundred K. It's not these ridiculous hundreds of megs sitting in your machine.

Leo: Oh, yeah, I'll need a Windows machine, won't I. Oh, shoot.

Steve: But to be able to look at graphs and charts of long-term DNS server performance, I think it's going to be very cool.

Leo: It's going to be very, very interesting. And that's what we hope for.

Steve: Oh, I forgot, built-in spoofability testing, too. So you can check the spoofability of the servers without having to do it generically over at GRC.

Leo: Nice.

Steve: So, yeah, lots of stuff.

Leo: Yeah. I run a network analysis program in the background almost all the time to keep an eye on, you know, our bandwidth and so forth, Fing. And I think this will be equally useful running in the background. I definitely look forward to it.