Security Now! #934 - 08-08-23
Revisiting Global Privacy Control

This week on Security Now!

What was it that also just, last week, happened with Voyager 2? What did Tenable's CEO Amit Yoran have to say about Microsoft's security practices? And what did Bruce Schneier have to say about the recent attack on Azure by Chinese hackers? There's more to AI than ChatGPT. What did some academic researchers in the UK accomplish by adding new deep learning modeling to a classic and previously weak attack? And after discussing some interesting listener feedback from the prior week, we're going to revisit a topic we covered when it was young because it's beginning to show signs that it might have a life of its own and may not be destined to fall by the wayside, as all brokers of personal information would hope.

Where there's a will…

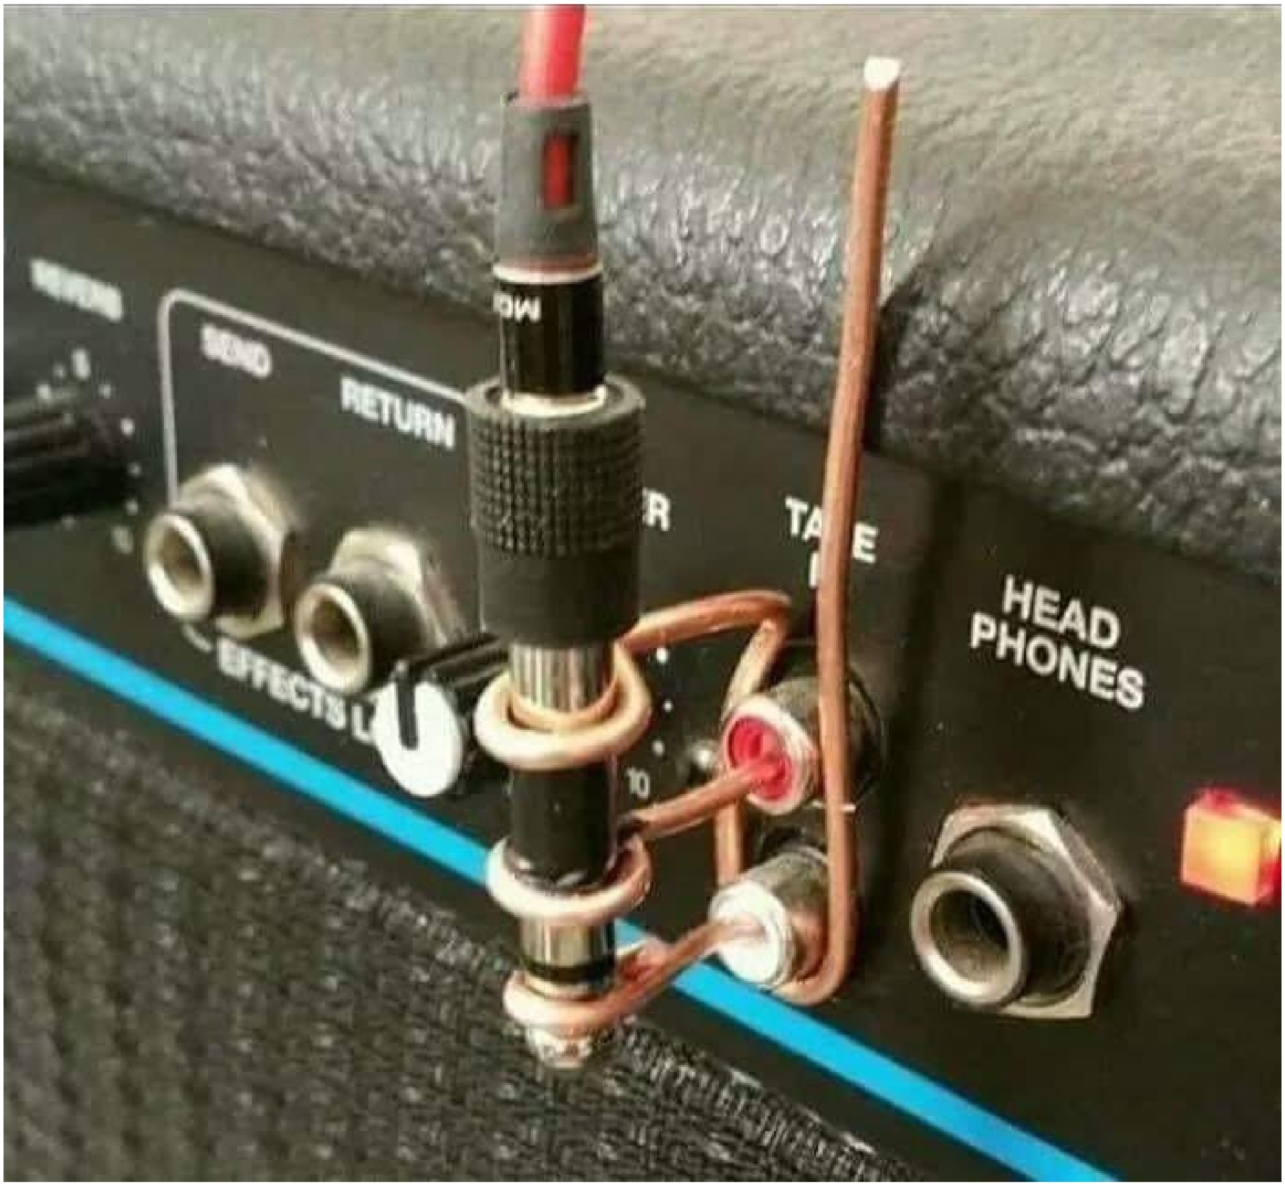

Security News

NASA “shouted” at Voyager

So when we last left the Voyager 2 space probe it had received a series of mistaken commands from ground control which caused it to turn 2 degrees away from earth. At its present distance of 12.3 billion miles, 2 degrees might as well be 90. This meant that no more data could be received nor could any corrective commands be sent. The good news is that as long as all is going well, Voyager has a fail-safe that was expected to perform an automatic reorientation this coming October. But that’s still three months away. NASA wrote:

> *Voyager 2 is programmed to reset its orientation multiple times each year to keep its antenna pointing at Earth; the next reset will occur on Oct. 15, which should enable communication to resume. The mission team expects Voyager 2 to remain on its planned trajectory during the quiet period.*

But then last week we received an update from NASA:

> *UPDATE, Aug. 1, 2023: Using multiple antennas, NASA’s Deep Space Network (DSN) was able to detect a carrier signal from Voyager 2. A carrier signal is what the spacecraft uses to send data back to Earth. The signal is too faint for data to be extracted, but the detection confirms that the spacecraft is still operating. The spacecraft also continues on its expected trajectory. Although the mission expects the spacecraft to point its antenna at Earth in mid-October, the team will attempt to command Voyager sooner, while its antenna is still pointed away from Earth. To do this, a DSN antenna will be used to “shout” the command to Voyager to turn its antenna. This intermediary attempt may not work, in which case the team will wait for the spacecraft to automatically reset its orientation in October. Either way, once the spacecraft’s antenna is realigned with Earth, communications should resume.*

Then, last Friday we received the good news:

> *UPDATE, Aug. 4, 2023: NASA has reestablished full communications with Voyager 2.*
> *The agency’s Deep Space Network facility in Canberra, Australia, sent the equivalent of an interstellar “shout” more than 12.3 billion miles (19.9 billion kilometers) to Voyager 2, instructing the spacecraft to reorient itself and turn its antenna back to Earth. With a one-way speed of light delay of 18.5 hours for the command to reach Voyager, it took 37 hours for mission controllers to learn whether the command worked. At 12:29 a.m. EDT on Aug. 4, the spacecraft began returning science and telemetry data, indicating it is operating normally and that it remains on its expected trajectory.*

You can just imagine the breath-holding that was going on during those 37 hours. But the entire project is an incredible engineering accomplishment.

Tenable CEO accuses Microsoft of negligence in addressing security flaw
https://cyberscoop.com/tenable-microsoft-negligence-security-flaw/

Another view of Microsoft
Everyone who listens to this podcast knows that I often become upset with Microsoft's behavior and with their performance. I sometimes feel odd since I can imagine someone reasonably saying: "If you have so much trouble with Microsoft why don't you just switch to Mac or Linux?" And it's true that I do love Windows and I have very little trouble with it myself. But due to their size and their global dominance, Microsoft's behavior matters and affects the world, regardless of what desktop platform I've personally chosen. And since this podcast covers security it also needs to explore Microsoft's many behaviors relating to security.

Last Wednesday, August 2nd, someone else weighed in on Microsoft's security practices from their own perspective and significant experience. Since I sometimes feel a bit self conscious tearing into Microsoft over and over I wanted to share this additional viewpoint. But for what this individual wrote to have any weight and bearing you need to know something about the posting's author, Amit Yoran. Wikipedia informs us:

> ***Amit Yoran*** *is chairman and chief executive officer of Tenable, a position held since January 3, 2017. Previously, Yoran was president of computer and network security company RSA.*
>
> *Yoran joined RSA during his tenure as CEO of NetWitness Corp., which was acquired by RSA's parent company, EMC, in April 2011. Prior to his time at NetWitness, Yoran was the National Cyber Security Division director within the United States Department of Homeland Security. He took up the post in September 2003 and served as the initial director of the US-CERT (that's, of course, the US Department of Defense's Computer Emergency Response Team.) He resigned from his position at US-CERT in October 2004.*
>
> *Earlier in his career, Yoran was a co-founder and CEO of Riptech, which was acquired by Symantec in August 2002. He has also served on the board of directors of Cyota (acquired by RSA), Guardium (acquired by IBM), Guidance Software, and other internet security technology companies.*
>
> *Yoran is a graduate of the United States Military Academy and served as one of the founding members of the US Department of Defense's Computer Emergency Response Team. He has a master's degree in computer science.*

Okay. So this guy has earned some street cred by being in the middle of computer security for many years. His LinkedIn posting last Wednesday is titled: "Microsoft…The Truth Is Even Worse Than You Think"
https://www.linkedin.com/pulse/microsoftthe-truth-even-worse-than-you-think-amit-yoran

Here's what Amit wrote and posted publicly on LinkedIn, a platform Microsoft purchased:

> *Last week, Senator Ron Wyden sent a letter to the Cybersecurity and Infrastructure Security Agency (CISA), the Department of Justice and the Federal Trade Commission (FTC) asking that they hold Microsoft accountable for a repeated pattern of negligent cybersecurity practices, which has enabled Chinese espionage against the United States government.*
>
> *According to data from Google Project Zero, Microsoft products have accounted for an aggregate 42.5% of all zero days discovered since 2014.*

> *Microsoft’s lack of transparency applies to breaches, irresponsible security practices, and to vulnerabilities, all of which expose their customers to risks they are deliberately kept in the dark about.*
>
> *In March 2023, a member of Tenable’s Research team was investigating Microsoft’s Azure platform and related services. The researcher discovered an issue which would enable an unauthenticated attacker to access cross-tenant applications and sensitive data, such as authentication secrets. To give you an idea of how bad this is, our team very quickly discovered authentication secrets to a bank. They were so concerned about the seriousness and the ethics of the issue that we immediately notified Microsoft.*
>
> *Did Microsoft quickly fix the issue that could effectively lead to the breach of multiple customers' networks and services? Of course not. They took more than **90 days** to implement a partial fix – and only for new applications loaded in the service.*
>
> *That means that as of today, the bank I referenced above is still vulnerable, more than 120 days since we reported the issue, as are all of the other organizations that had launched the service prior to the fix. And, to the best of our knowledge, they still have no idea they are at risk and therefore can’t make an informed decision about compensating controls and other risk mitigating actions. Microsoft claims that they will fix the issue by the end of September, four months after we notified them. That’s grossly irresponsible, if not blatantly negligent. We know about the issue, Microsoft knows about the issue, and hopefully threat actors don’t.*
>
> *Cloud providers have long espoused the shared responsibility model. That model is irretrievably broken if your cloud vendor doesn’t notify you of issues as they arise and apply fixes openly.*
>
> *What you hear from Microsoft is “just trust us,” but what you get back is very little transparency and **a culture of toxic obfuscation**. How can a CISO, board of directors or executive team believe that Microsoft will do the right thing given the fact patterns and current behaviors? Microsoft’s track record puts us all at risk. And it’s even worse than we thought.*

“A culture of toxic obfuscation.”

By looking at the facts through the years, we’ve documented a great many instances where Microsoft’s behavior, whether apparently deliberate or inadvertent – yet either way quite difficult to see as anything other than “we’re so big we don’t need to care and you can’t make us” – has clearly damaged their own customers; even significantly. But their enterprise and government customers are as captive as I am. I’m held captive by my decades long investment in Windows. And by the fact that there is no viable alternative to Windows for the things I want to do. I depend upon many tools that are only hosted on Windows. And Microsoft’s big enterprise customers have invested massively in their own solutions which are not portable to any other platform. The incredible power this position gives Microsoft should not be underestimated. It leaves the entire world asking, “Please, sir... may I have some more soup?”

Amit Yoran’s posting on LinkedIn prompted an interview by CyberScoop. They, in turn, wrote:

> *Veteran cybersecurity executive Amit Yoran accused Microsoft on Wednesday of dragging its feet on fixing a critical vulnerability affecting its Azure platform and said the tech giant’s slow*

response illustrates a negligent approach to security.

His harsh public critique of Microsoft — a relatively rare event for a high-profile corporate figure in cybersecurity — follows criticism from lawmakers and researchers alike after a recent cyberattack affecting U.S. government officials resulted from a Microsoft security lapse.

As the CEO of Tenable, a firm that helps companies understand and mitigate their cybersecurity vulnerabilities, Yoran said he works with hundreds of companies every year to disclose and patch vulnerabilities. ***Microsoft, he said, consistently fails to proactively and professionally address vulnerabilities in their products.***

Yoran told CyberScoop in an interview: "In Microsoft's case you have a culture which denies the criticality of vulnerabilities."

According to a timeline in a limited blog published to Tenable's website, Microsoft acknowledged the issue the same day it was disclosed on March 30, and confirmed it four days later. Tenable asked for an update June 27 (90 days later) and was told on July 6 that it was fixed, but Tenable says it was merely a partial fix.

[Gee, where have we heard that before? How many times have we noted here that someone at Microsoft, who was shown a serious vulnerability and even given the fix for it by a security researcher, apparently didn't take the time, or care, to actually understand the underlying problem, and so only half patched it to resolve one of the problem's symptoms? Anyway...]

On July 21, Microsoft told Tenable that it would take until Sept. 28 for a complete fix. Tenable agreed to withhold technical details and proofs-of-concept until Sept. 28.

In his blog post, Yoran described Microsoft's approach to addressing the issue as ***"grossly irresponsible, if not blatantly negligent."*** *Yoran wrote that "More than 120 days since the vulnerability was reported, the bank in question remains vulnerable", adding that many vulnerable organizations "still have no idea they are at risk and therefore can't make an informed decision about compensating controls and other risk mitigating actions."*

[And then we heard from Microsoft, I love this:]

A spokesperson for Microsoft said that the company appreciates "the collaboration with the security community to responsibly disclose product issues" and that security updates are ultimately "a delicate balance between timeliness and quality, while ensuring maximized customer protection with minimized customer disruption."

Microsoft said Friday in a blog post that the issue has "been fully addressed for all customers," no customer remediation action is required and that all affected customers were notified via eMail starting Friday. Microsoft said its investigation "identified anomalous access only by the security researcher that reported the incident, and no other actors."

Yoran's broadside against Microsoft comes amid growing scrutiny of Microsoft in Washington after one of the company's products was abused by hackers based in China to steal the email messages of senior U.S. officials. In that incident, hackers based in China were able to steal an encryption key that they could then use to forge authentication tokens, and security researchers have sharply criticized the company for not only allowing an encryption key to be stolen but for building a computing architecture in which tokens could be forged in this way at all.

> *The incident spurred Oregon’s senator Ron Wyden to call Microsoft “negligent” in its security practices and request that the Justice Department investigate whether Microsoft’s actions in the incident broke the law.*
>
> [Good luck with that, Washington. A LONG time ago when Microsoft was much smaller and far less powerful it was nearly impossible to hold its behavior to account. There’s just no possibility of doing so any longer.]
>
> *While Microsoft has insisted that the Chinese operation was highly targeted, research by the cloud security company Wiz suggests the incident may have been more broad than first understood — a claim Microsoft has dismissed as speculative.*
>
> [Right. Microsoft dismissed this as speculative because, after all, a delicate balance is required between timeliness and quality, while ensuring maximized customer protection.]
>
> *The vulnerability discovered by Tenable allowed “an unauthenticated attacker to access cross-tenant applications and sensitive data, such as authentication secrets,” according to Yoran’s blog post. It appears that vulnerability does not exploit the same types of authentication flaws seen in the recent incident involving Chinese hackers, but may add pressure on Microsoft to improve its security practices.*
>
> *Industry professionals and government officials pointed out that the Chinese operation was only detected because a government agency was paying additional money for more sensitive logging capabilities. Microsoft later reversed that policy and expanded logging visibility and retention for certain customers.*
>
> *Yoran, who has grown increasingly critical of Microsoft in recent years, told CyberScoop that the company’s dominant position in the technology ecosystem makes many computer security researchers hesitant to speak up about its security practices but that doing so is especially important given the ubiquity of its products.*
>
> [Exactly.]
>
> *“Microsoft is a pretty strategic problem in the security space given the pervasiveness of their software, of their infrastructure,” Yoran said. “I also think they have to be part of the solution.”*

I’m not a fan of complaining about problems that no one has any power to resolve. As an engineer and technologist I most enjoy discovering and sharing solutions to problems. But ignoring truly important issues in a podcast that’s focused upon security seems negligent, too. So we’ll just keep perspective, discuss problems, and celebrate those companies who **do** act quickly and responsibly in the best interests of the users of their products.

But there is the issue of that recent serious attack by Chinese hackers...

What about this Chinese attack?

Several weeks ago, while working on a previous episode of this podcast, I saw this news that's referred to in the CyberScoop piece. I suppose I let it slide past because, well, what's that expression about beating a dead horse? At some point I'm sure that we all get tired of complaints about Microsoft. Sort of like how many ransomware attacks are we going to detail? At some point, what's the point? But saturation shouldn't keep us from covering important security events and this Chinese attack was very important and quite significant.

The best way to deal with covering it now is to refer to a well-known industry expert who very nicely framed what happened. He is, Bruce Schneier, and Bruce posted under the title "Microsoft Signing Key Stolen by Chinese"...

> *A bunch of networks, including US Government networks, have been hacked by the Chinese.*
>
> *The hackers used forged authentication tokens to access user email, using a stolen Microsoft Azure account consumer signing key. Congress wants answers. The phrase "negligent security practices" is being tossed about—and with good reason.* ***Master signing keys*** *are not supposed to be left around, waiting to be stolen.*
>
> *Actually, two things went badly wrong here. The first is that Azure accepted an expired signing key, implying a vulnerability in whatever is supposed to check key validity. The second is that this key was supposed to remain in the system's Hardware Security Module—and not be in software. This implies a really serious breach of good security practice. The fact that Microsoft has not been forthcoming about the details of what happened tells me that the details are really bad.*
>
> *I believe this all traces back to SolarWinds. In addition to Russia inserting malware into a SolarWinds update, China used a different SolarWinds vulnerability to break into networks. We know that Russia accessed Microsoft source code in that attack. I have heard from informed government officials that China used their SolarWinds vulnerability to break into Microsoft and access source code, including Azure's.*
>
> *I think we are grossly underestimating the long-term results of the SolarWinds attacks. That backdoored update was downloaded by over 14,000 networks worldwide. Organizations patched their networks, but not before Russia—and others—used the vulnerability to enter those networks. And once someone is in a network, it's really hard to be sure that you've kicked them out.*
>
> *Sophisticated threat actors are realizing that stealing source code of infrastructure providers, and then combing that code for vulnerabilities, is an excellent way to break into organizations who use those infrastructure providers. Attackers like Russia and China—and presumably the US as well—are prioritizing going after those providers.*

So, Bruce nicely and succinctly explained what happened with the Microsoft Azure mess. In short, they first deeply screwed up then they failed to take responsibility for their screw-up. And only now Washington is starting to wonder how Microsoft became this powerful. *News Flash!*

I also thought of Bruce Schneier recently in another context because I love to quote one of his pithy observations *"Attacks always get better; they never get worse."*

While that’s kind of obvious, reminding ourselves of its truth serves as a nice reality check. And in this case it explains what recently happened with the classic “attack” of listening to someone typing on a keyboard:

AI meets Keyboard Acoustic Side-Channel attacks

Although significant controversy surrounds questions regarding the current and future impact of ChatGPT-style conversational AI models, a huge amount of far less glamorous yet nonetheless important work is being done by applying some of these newly emerging AI’ish techniques to previously explored domains. We’ve talked before about the concept of having a smartphone resting on a desk surface with its microphone passively listening to the keystrokes being typed nearby. If this were practical, it would represent acoustic side-channel leakage from the keyboard. And since confidential information might be entered through that keyboard – and since in general, no one wants or expects to have their keystrokes surreptitiously monitored and recorded – it would represent an attack. And speaking of attacks… they always get better, they never get worse.

Last Thursday, August 3rd, a trio of researchers from three different universities in the UK, published a paper for the “2023 IEEE European Symposium on Security and Privacy Workshops” conference. Their paper is titled: *“A Practical Deep Learning-Based Acoustic Side Channel Attack on Keyboards.”* Here’s what they described from their research and of its success:

> *With recent developments in deep learning, the ubiquity of microphones and the rise in online services via personal devices, acoustic side channel attacks present a greater threat to keyboards than ever. This paper presents a practical implementation of a state-of-the-art deep learning model in order to classify laptop keystrokes, using a smartphone integrated microphone. When trained on keystrokes recorded by a nearby phone, the classifier achieved **an accuracy of 95%**, the highest accuracy seen without the use of a language model. When trained on keystrokes recorded using the video-conferencing software Zoom, an accuracy of 93% was achieved, a new best for the medium. Our results prove the practicality of these side channel attacks via off-the-shelf equipment and algorithms. We discuss a series of mitigation methods to protect users against these series of attacks.*

https://arxiv.org/pdf/2308.01074.pdf

So, this is a phenomenal level of recognition for an outbound external microphone that’s simply listening to keystrokes. And to only lose 2% accuracy when significantly compressing the audio through Zoom is further astonishing. Imagine being able to process the recorded sounds of someone typing after the fact to obtain a near-perfect rendition of what they originally keyed. This is achieved essentially by utilizing far more of the total available information than any previous efforts have managed.

For anyone who wants all of the details, I’ve included a link to the entire 21-page research report. But I think we already have the gist of the idea. And there’s an important lesson here for us. Regardless of the outcome of the debate over the true longer term value of ChatGPT-style interaction, I think it’s very clear that something has happened recently and that the world has been changed. We’re still not sure of the “what” and “how” of these changes; and I’m also certain that they're still underway. Research like this demonstrates that applications of the new deep learning models have only begun to be explored. I expect we're going to be seeing some very significant discoveries in the future once these relatively new capabilities become more widely available. Lord only knows what those side-channel attack masters at the Ben-Gurion University of the Negev in Israel are going to come up with once they add deep learning modeling to their many bags of tricks.

Closing the Loop

Rusty / @rusty0101

With another take on the "in the cloud or on the ground" discussion…

> *Listening to this week's SN, with the discussion of running things in the cloud. I'd note that more and more people are running their own power stations, either with solar, wind, or water-wheel systems. Including Amazon for at least one of their AWS sites. I think that's becoming less and less of a useful argument. Additionally there have been recent cloud providers who've decommissioned equipment that was providing cloud services right up until it was shut down, and apparently end users didn't get the word for some reason, some of whom have lost significant functionality as a result. Perhaps that's not going to be an issue for some of the larger providers, but if you are trying to work within a budget, there may be storm clouds on their way.*

I think there's no question that there's a real and vital role for cloud-based services. I'm not intending to suggest otherwise. But there can also be a bit of a gold rush mentality of imagining that the only reason there's still anything that's **not** "in the cloud" is inertia and that eventually eventually will be. I think the reality is there's probably a place for both. That's the point that I intended to make.

Alan C. Bonnici / @chribonn

> *Hi Steve, I heard you speak about Authy and decided to give it a try. I reset the 2FA code on a gmail account to generate a new Code. What is strange is that the TOTP in my Password Manager is different from Authy. I managed to log in with both. Could it be clock differences between my desktop and my phone? If yes why would both work? Fan of the show.*

Alan followed-up a bit later to confirm that it was, indeed, a clock difference. As to why both of the different codes would work, many authentication receivers will continue to accept a recent if not 100% current code. When they receive a code and the present one doesn't work, they may try the next one that's about to come up, or they might try the previous code, and maybe even the one before that. The point is that for the system to work, both endpoints need to share not only the same secret key which keys the pseudo-random sequence, but they must also agree upon the time of day. In today's Internet-connected world, it's easy for devices to be within very clock time agreement, but it's also reasonable to make some allowances for them not being.

Joe LaGreca / @lagreca

I'm finally ready to leave the Google Chrome browser.......Which browser do you use or recommend? Firefox?

Firefox is where both Leo and I are and I'm completely happy with the choice. Some time ago I tried using Bing, but I was stunned when I encountered some sites that it would not render. And I had been using Chrome too just to see how it compared to my longtime use of Firefox. Today, having satisfied my curiosity, I'm back to Firefox and many things about it are right for me. I need to use an add-on to get my tabs to run down the left side of the browser. But there's a slick session manager that allows me to save entire browser sessions. I use it with I'm working on the podcast and need to change locations.

Anyway, yes. 100% Firefox. And when we get to talking about today's topic you'll learn another reason why it continues to be my choice.

Seven / @dergit73

Apologies in advance if this is a topic you've covered ad nauseam. I listened to SN religiously from episode 1 through several hundred, but I had to take a few years off from extra curricular listening. I've since subscribed to Club TWiT and returned to attending weekly services. I don't know if my question will be simple enough to address in a DM, but perhaps with at least with a suggestion of where to start.

After receiving a notification that one of my accounts was compromised INCLUDING the password, I have come to fully realize that NO passwords are safe. Period. I use 2FA wherever possible, but of course 2FA support isn't consistent across all services. Is there a best way to simplify the process of not relying on passwords alone?

Is there a simple answer to the question: WWSGD?

It took me a moment to parse WWSGD, but it's clearly meant to mean: "What Would Steve Gibson Do?" So "Seven" is addressing the question of remote authentication over the Internet. I'll expand a bit on Seven's question a bit by answering "What does Steve think about the current and probable future state of identity authentication over the Internet?"

One way to view our current security environment (and I'll discuss a second way after this) is to see that what's developing is a spreading spectrum of options. This is always what we get when new and better solutions at last start being adopted. The reason we wind up with a spectrum spread is that the appearance of new and better solutions doesn't kill off the older and less secure solutions.

Despite the fact that 2FA has been widely available for many years **most** sites still don't offer it as an option. Partly that's due to inertia and partly due to a lack of perceived need, and partly because making logon more difficult increases support overhead. And now we have passkeys which represents another step forward. But will passkeys kill off two-factor authentication and passwords? Of course not. Over time, more sites offering passkey support will appear.

Eventually, support for 2FA and Passkeys will be baked into servers and servers may take more responsibility for authentication. But we also know that many sites still won't care. They'll feel that identifying their visitors with an eMail address and a password is sufficient. And for many sites they're probably correct. More and more often, Internet users are being asked to "create an account" as a requirement to get in the front door. Why? Probably because it forces its visitors to turn over an eMail address for the receipt of follow-up spam. It allows a site you may never choose to visit again to continue to plague you into the future. And it may also be that sites will then be able to further monetize your existence by selling whatever information they managed to accumulate about you. This is one place where today's topic "Global Privacy Control" may be relevant.

But the other fact is that eMail, and one's control over an eMail address, remains the ultimate fallback when anyone is unable to remember how to logon. I've joked here in the past that the "I forgot my password" link appearing underneath every password prompt makes a strong case for not bothering with ever remembering any passwords. Just bang on the keyboard for a while when you create an account, then click the "I forgot my password" link whenever you want to come back. And I actually think that people would do that if it weren't actually quicker and easier to have a password manager remember and then fill-in the answer for you. But what does that mean about the actual security being delivered?

And from an actual security standpoint, I have to say that what's really infuriating and even somewhat confounding is to see a two-factor authentication prompt followed by a link saying: *"I'm unable to use my authenticator right now."* **What?!?!** What's the point of "requiring" one if you can just say *"my dog ate it?"* and then be allowed to logon without it? Wikipedia now has an entry on the topic of *"Security Theatre"*, defining it as: *"Security theater is the practice of taking security measures that are considered to provide the **feeling** of improved security while doing little or nothing to achieve it"* and Wikipedia references Bruce Schneier as the originator of that perfect term.

And we know that passkeys are going to be just the same, right? Since no one ever wants increased authentication security to actually prevent anyone from authenticating, there will always be the get out of jail free card of *"just click the link we sent to the eMail address you have on file with us so we know it's really you."*

So, to Seven's original question I would reply: If you want the most security possible, the only thing you can do is to take advantage of the most secure authentication option available on a site-by-site basis. Use a password manager to remember your random and long secrets. Use an authenticator app to generate your one-time passwords. As Passkeys become available, use them wherever possible. But also keep in mind something that has not received enough attention anywhere by anyone – and I sincerely hope that it does not receive more attention by the bad guys – which is that because security is about the lowest common denominator, and due to the ubiquitous role that eMail continues to play as the ultimate "my dog ate my authentication" authentication recovery – which might be better termed "total authentication bypass" – any entity who is able to obtain access, even transient access, to your eMail flow obtains unfettered access to your entire online life.

Revisiting Global Privacy Control

Today's podcast adventure was triggered when I followed a news link yesterday over to TechCrunch. The screen darkened with an overlay, as screens do these days when a site wants to bring something to its visitor's attention. And I was left staring at an interesting notice from TechCrunch:

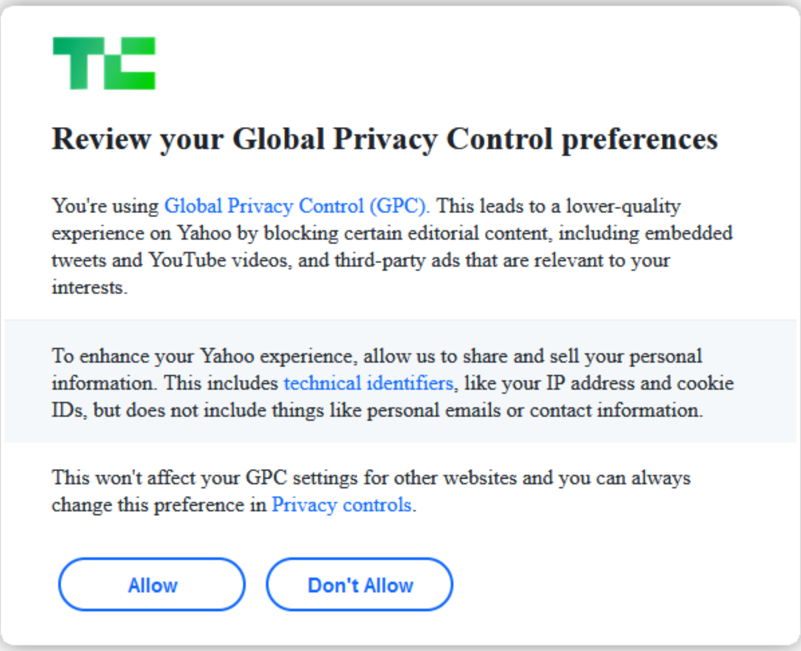

Review your Global Privacy Control preferences

You're using Global Privacy Control (GPC). This leads to a lower-quality experience on Yahoo by blocking certain editorial content, including embedded tweets and YouTube videos, and third-party ads that are relevant to your interests.

To enhance your Yahoo experience, allow us to share and sell your personal information. This includes technical identifiers, like your IP address and cookie IDs, but does not include things like personal emails or contact information.

This won't affect your GPC settings for other websites and you can always change this preference in Privacy controls.

The pop-up overlay notice's headline was "Review your Global Privacy Control preferences." And the notice read:

You're using Global Privacy Control (GPC). This leads to a lower quality experience on Yahoo by blocking certain editorial content, including embedded tweets, YouTube videos and third-party ads that are relevant to your interests.

Huh. That's interesting. And it gets better...

To enhance your Yahoo experience, allow us to share and sell your personal information.

Right. Because it's very clearly in my best interest to have my Yahoo experience enhanced by allowing them to share and sell my personal information, no doubt about it. And there's more:

This includes technical identifiers like your IP address and cookie Ids, but does not include things like personal eMails or contact information.

And then it concludes with:

This won't affect your GPC settings for other websites and you can always change this preference in "Privacy Controls."

And then I was presented with two options: **Allow** or **Don't Allow**. You can probably guess which one I chose.

There are several bits of good news here. One is that someone made them do this. The other is that the only reason I received this notice was that I took my own advice back on May 3rd of last year, in 2022, as a result of our podcast #869 which was titled *"Global Privacy Control"*. I flipped a switch that's built into Firefox and then promptly forgot about it. But that switch remained flipped. And I should note that perhaps it's no surprise that the switch is missing from Chrome. However in addition to Firefox, which incorporates it natively, it's present in both the Brave and the DuckDuckGo privacy browsers; and it can be added to Chrome with the use of a third-party extension.

So we have many things to talk about, here. First of all, to clear up one question, I was visiting TechCrunch and was informed by that pop-up that Yahoo wanted me to drop my pants. Wikipedia explains this, writing:

> *In 2010, AOL acquired [TechCrunch] for approximately $25 million. Following the 2015 acquisition of AOL and Yahoo by Verizon, the site was owned by Verizon Media from 2015 through 2021. In 2021 Verizon sold its media assets, including AOL, Yahoo, and TechCrunch, to the private equity firm Apollo Global Management, and Apollo integrated them into a new entity called* ***Yahoo! Inc.****.*

The next thing that caught my eye in that pop-up was their term *"technical identifiers"* which the notice was hoping I would be willing to allow them to share and sell. Seems we should know what those are. That term in the pop-up was also a link which took me to a Yahoo! Page titled: *"Collection, Use, and Linking of Technical Identifiers"*. They write:

> *Yahoo uses different technical identifiers to make its consumer services available on most platforms, browsers, and devices. Yahoo also uses these technical identifiers to provide our digital advertising services on our properties and for our business partners.*
>
> *As detailed in our Cookie Policy, these technical identifiers include:*
>
> - *Browser cookie identifiers (sometimes referred to as "cookie IDs") and browser local storage identifiers*
> - *Mobile device identifiers, such as the Android advertising ID or the Apple Identifier for Advertising (IDFA)*
> - *Platform or operating system-based identifiers, such as those offered on smart or connected TVs or media streaming devices*
> - *Partner-supplied technical identifiers*
> - *Encrypted or one-way cryptographic hashes of personal information such as email addresses, phone numbers, account identifiers, derivatives, or escalated versions of these identifiers*
> - *Household-based identifiers*
> - *IP addresses*
> - *Probabilistic (non-unique) identifiers*
> - *Identifiers generated from the combination of various device, browser, or operating system attributes, such as the operating system or browser version*

> - *"Cohort", audience, or group identifiers, such as "sports enthusiasts"*
>
> *The storage, generation, and collection methods of these identifiers may also vary, depending on the context. For instance, some browsers and devices offer limited technical identifier support and/or limited cookie support, so non-cookie-based identifiers may be used in these cases. Examples of these devices include:*
>
> - *Smart or connected TVs, over-the-top (OTT) streaming devices (such as a Roku device), and similar interactive media players*
> - *Digital-out-of-home (DOOH) billboards and similar media devices*
> - *Browsers enabled with intelligent tracking prevention (ITP), privacy sandbox, or similar cookie-blocking technology*
> - *Certain apps, mobile devices, or installed software, where permitted and applicable*
> - *Certain internet-of-things (IoT) devices*
>
> *The collection methods for technical identifiers and associated data depend on the context, as described here. When using the Internet in a browser (for example, Chrome), our consumer services and digital advertising services may use standard cookies, Javascript code, libraries, and/or dynamic HTML tags, web beacons, and similar technologies. In mobile apps, our consumer services and digital advertising services may use mobile software development kits (SDKs), local or remote application programming interfaces (APIs), and similar client or server-side code. In other cases, we may exchange data and files (such as log files) with our partners in "offline" contexts using secure server-to-server transfer methods, APIs, cloud services, mutual agents or technology service providers, or other industry-standard methods.*
>
> *Technical identifiers may be used to identify a user across multiple devices, often referred to as "cross-device-linking" or "cross-device identifier resolution". As a result, technical identifiers that are presumed to belong to a particular user, device, or household can be linked to one another, and the associated technical identifier may be used to reference data, personalize advertisements, or tailor experiences. This process may be implemented and used by us or in coordination with our advertising partners as part of our digital advertising services.*

Okay. So, in short, "Technical Identifiers" amounts to pretty much anything and everything they can possibly get their hands on to track me and associate me with any members of my family and presumably coworkers through instances of shared IP addresses – tracking us across any and all of our devices using every trick and technique that's available to them.

Thank god I said no. But also thank goodness I was asked, and had the **opportunity** to say no. ***Not everyone is given the option – you have to ask for it!*** And we have California's state legislature and Attorney General, as well as those in Colorado and Connecticut to thank for this. I'll explain all that in a moment. But this is certainly not my first visit to TechCrunch recently. I've been popping over to TechCrunch from time to time with Firefox, following links to news to share with this audience, and this is the first time I've ever seen that pop-up. So this is new behavior. I received that notice because ever since we first talked about this in May of 2022, my Firefox browser has been broadcasting the standardized GPC – Global Privacy Control – signal to indicate that I do not wish to have my ***"online experiences enhanced"*** at the cost of my, my family's and my company's privacy.

Before I move on, I need to note that while digging deeper into what was up with TechCrunch, I followed TechCrunch's *"Your Privacy Choices"* link which you can find at the bottom of their pages. The first interesting thing is that the right hand side of the page has specific pages for California, Colorado and Connecticut as I mentioned before. But there's also one for Virginia.

The California page I was taken to has two switch settings. Due to my previous reply to the pop-up, the first one was turned off and set to "Don't Allow". That corresponded to "Allow the Sale and Sharing of My Personal Information." But there was a second switch and it was still turned on and set to "Don't Limit". And that one corresponds to "Limit the Use of My Sensitive Personal Information". The page explains:

> *In connection with providing our services, we may use sensitive personal information such as precise location data and email content data. Among other purposes, we use such data to help understand your interests so we can show you more relevant ads and content.* ***To opt out of this use and limit our use of such information to only those purposes permitted by California law, select "Limit".*** *This may make the content and ads that you see less relevant to you.*

After doing some additional research I figured out what's going on for anyone in California and why Yahoo! is showing two switches: Under the California Consumer Privacy Act (CCPA), California consumers have the right to opt-out of the sale and the sharing and the use of their personal information. Those three things: Sale, Sharing & Use. But the Global Privacy Control as it is presently defined – and it's quite unlikely to ever have its strength broadened – ONLY applies to those first two of the three: personal information sales and sharing. The GPC does not also cover the **use** of personal information. But, California law does. So if Californian's want to prohibit the **use** of their personal information (beyond its sale and sharing which can be done globally with the GPC) that will still need to be done on a site-by-site basis.

Here's how the GPC places itself, its need and the role it's filling. Its specification explains:

> *Building websites today often requires relying on services provided by businesses other than the one which a person chooses to interact with. This result is a natural consequence of the increasing complexity of Web technology and of the division of labor between different services. While this architecture can be used in the service of better Web experiences, it can also be abused to violate privacy. While data can be shared with service providers for limited operational purposes, it can also be shared with third parties or used for behavioral targeting in ways that many users find objectionable.*
>
> *Several legal frameworks exist — and more are on the way — within which people have the right to request that their privacy be protected, including requests that their data not be sold or shared beyond the business with which they intend to interact. Requiring that people manually express their rights for each and every site they visit is, however, impractical.*
>
> ***[The spec then quotes the California Attorney General]***
>
> > *"Given the ease and frequency by which personal information is collected and sold when a consumer visits a website, consumers should have a similarly easy ability to request to opt-out globally. This regulation offers consumers a global choice to opt-out of the sale of*

> *personal information, as opposed to going website by website to make individual requests with each business each time they use a new browser or a new device.*
>
> *This specification addresses the issue by providing a way to signal, through an HTTP header or the DOM, a person's assertion of their applicable rights to prevent the sale of their data, the sharing of their data with third parties, and the use of their data for cross-site targeted advertising. This signal is equivalent, for example, to the "global privacy control" in the CCPA regulations.*

What's also annoying, now that I've woken up to this, is that I should have never received that pop-up in the first place. My browser's GPC setting is not the default. And it's not even available from Chrome without an add-on. So if a browser is broadcasting it, it's because this is what its owner means and wants. Which means that the pop-up I received was TechCrunch's *"are you really sure this is what you want here? – Would you consider changing your mind, pretty please with a cherry on top?"*

I'll also note that all four of the states that have enacted GPC-specific legislation have differing definitions and language in their laws. So each of those four pages, where TechCrunch's parent company Yahoo! is juggling legislation, is different from the others. This means that we now have state- by-state privacy laws, and that Yahoo! is desperately clinging to the leverage of every bit of personal information – its sales, sharing and internal use – that they can, on a state-by-state basis.

Okay, so now let's step back a bit to get some perspective on the whole Global Privacy Control issue. I found a great write up at a site called **Firewalls Don't Stop Dragons:** https://firewallsdontstopdragons.com/how-to-enable-global-privacy-control/

They wrote:

> *You are tracked mercilessly today when you surf the web, either on your computer or your smartphone. Websites use several different techniques to identify you and record as much data about you as they can. While marketers will claim that you have the power to opt out of most tracking, this is frankly impossible to do, practically speaking. There are simply too many trackers, many of which you'll never know about. There's a new(ish) initiative that aims to address this problem called Global Privacy Control, or GPC. GPC is a browser setting that lets you automatically tell every website you visit to stop collecting your data. Sounds good, right? But it may also sound familiar…*
>
> *Back in 2009, a group of researchers had a brilliant idea: why don't we give users a way to tell every website they visit that they don't want to be tracked? They came up with a simple, global **Do Not Track (DNT)** flag that users could set on their web browser once and forget it. Their browser would, in turn, tell every website you visited that you did not wish to be tracked.*
>
> *The obvious problem here is that websites (at that time) were under precisely zero obligation to comply. But there were also a couple interesting twists to the story. At one point, Microsoft took it upon themselves to automatically enable the DNT flag for Internet Explorer users. Advertisers were outraged because the flag was supposed to be an affirmative action taken by the user. They used this move as another reason to ignore the flag. And in an ironic twist, the very fact that your browser set this flag made you more trackable.*

> ***“Global Privacy Control: DNT 2.0”***
>
> *It turns out that DNT was a little ahead of its time. Without any legal reason to comply, it never caught on and was eventually abandoned. If it had only held out a little longer, it might have been relevant. The European Union’s General Data Protection Regulation (GDPR) was just coming online around the same time DNT was abandoned. However, the GDPR user consent verbiage didn’t seem to explicitly recognize DNT.*
>
> *Enter Global Privacy Control.*
>
> *From everything I can see, it’s really just “DNT 2.0”. However, this time there are legal requirements, at least in some regions, to actually require compliance. In particular, the California Consumer Privacy Act (CCPA) and subsequent California Privacy Rights Act (CPRA) have explicit language requiring sites to honor these automated requests not to be tracked. Similar laws have been passed in Nevada, Utah, Colorado, Virginia, and Connecticut – with others coming. GPC may yet succeed where DNT failed.*

I should pause here to note that just as the terms “Do Not Track” and “Global Privacy Control” sound like different things, they are indeed. So, as much as I like what this author has written, everyone who follows this podcast knows that I’m a stickler for detail. So when he says that GPC is really just DNT 2.0. That’s only true in as much as it’s a global beacon that browsers can be configured to send. That part of GPC is the same as DNT. But just to be clear, GPC is explicitly **not** about tracking. As I’ve been careful to say, it’s about prohibiting the sales and sharing of personal information. This author continues to make some good points about “How to Enable Global Privacy Control” he writes:

> *This is not a slam dunk. For one thing, there is no US federal law requiring companies to respect GPC. Also, the GDPR interpretation of GPC sadly seems a little weak. There are still too many regions that have no privacy regulations. And the various regulations that do exist need to be “harmonized” with one another on what GPC really means. For example, does the request apply only to further data collection or should it apply to data already collected? Does it apply to the user or just the device that sent the GPC flag?*
>
> *If you’re lucky enough to live in a region that has privacy laws, it’s a no-brainer – just enable it. But even if you don’t, there’s no reason you shouldn’t go ahead and register your desire not to be tracked.* [Which I’ll correct to say “not to have your personal information sold or shared. But otherwise he’s right.] *Then whenever and wherever this request is required to be honored, you’ll get the benefit.*
>
> *Thankfully, it’s pretty easy to do. And if you’re already using privacy tools, you may find that GPC has already been enabled. The test is simple: go to the Global Privacy Control website. If you see a green dot and “GPC signal detected” at the top, you’re good!*
>
> https://globalprivacycontrol.org/

Colorado’s Privacy Act (CPA) and Connecticut Data Privacy Act (CDPA) both recently went into effect on July 1, 2023. And like California’s CPRA, those states’ legislation require companies to honor the GPC.

Today, Firefox, Brave and the DuckDuckGo Privacy browser all support the GPC.

https://globalprivacycontrol.org/orgs

As for browser extensions for Chrome or other Chromium browsers that do not yet natively offer GPC, there’s “Abine” (The DeleteMe people), Disconnect, OptMeowt and Privacy Badger by the EFF. I have a link here at the end of the show notes to the GPC page which maintains a list of available extensions.

At the moment those in California, Colorado and Connecticut have the advantage of state laws which compel compliance with their resident’s GPC request. It doesn’t appear that websites serving Virginians, which does have similar privacy laws, are similarly bound to follow the GPC. But what we need now that the GPC exists and is gaining some ground will first be for additional states to step up and add their voices. Then we need the US Federal government to take this initiative national. At that point everyone will be on equal footing with the ability to opt-in globally. And then we can imagine the day when a federal law won’t require the presence of a GPC beacon. Well... we can dream, can’t we?