Revisiting Global Privacy Control

Description: 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.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-934.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-934-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. Whither Voyager 2? We've got an update on that story. You know Steve paid a lot of attention to Amit Joran's screed against Microsoft, on Microsoft's own social network, LinkedIn. What did Amit say, and why does Steve agree, coming up. Also Bruce Schneier on the recent attack on Azure by Chinese hackers. And then we'll talk about Global Privacy Control, how to turn it on, why you want to turn it on, and how Yahoo is responding. It's pretty amazing. All of that coming up and a whole lot more, this week on Security Now!. Stay tuned.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 934, recorded Tuesday, August 8th, 2023: Revisiting Global Privacy Control.

It's time for Security Now!, the show where we cover the week's security news, reassure you that there is no reason to set your hair on fire with this guy right here, whose hair apparently was set on fire some time ago, Mr. Steve Gibson. Hello, Steve.

Steve Gibson: Now, the question is, is it W-E-E-K security or W-E-A-K security.

Leo: No. Nothing weakly about the security here.

Steve: Nothing weak, weaklings, no.

Leo: Weaklings, no weaklings.

Steve: Okay. So today's topic, we are going to revisit something that we talked about a little over a year ago, in May of 2022. And I didn't plan this, but when I went to TechCrunch, something happened. And so I thought, whoa, what? Wait. So today's topic for Episode #934 for 8/8/23 is Revisiting Global Privacy Control. But before we get to that, we're going to update on what happened last week with Voyager 2. Something new happened in that, I mean, it's providing more news than JPL and NASA wishes it was producing.

Leo: Or that anything 12 billon miles away should, really.

Steve: Yeah, that's right, yeah. The good news is they haven't lost it, although it'd be easy, you know, to lose track of it. Also we're going to answer the question 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? Also, 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, as I said, a topic that 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 our personal information would hope. And of course we've got another great Picture of the Week. So I think a fun podcast for our listeners.

Leo: Always. Always fun with Steve. And informative. We can't forget that. Tell you another thing we need. We need a Picture of the Week, Mr. Gibson.

Steve: So, okay. Now, normally I'm able to explain on an audio podcast what the photo is that we're looking at.

Leo: You're not going to attempt that with this one, huh?

Steve: Well, this is a bit challenging. Okay. So here's the problem that an industrious person set out to solve. They had one of the newer style stereo plug connectors, like remember the old days, headphone jacks were large plugs. Technically they're called "quarter-inch phone plugs," or "phono plugs." But then newer headphones have the little smaller stereo connector, you know, like the kind that would plug into our smartphones when smartphones still had headphone connectors. So the problem is, so a person has a signal source with that kind of a connector on the end. Yet he wants to connect it to the old-school RCA-style mono plug that we used to have on the back of stereos, for example, and in this case it's the tape-in connectors. And so there's two of them, right, because each one is monaural.

Leo: All you have to do is you state the problem. How do I get a mini jack output into dual RCA, left and right RCA connectors? How do I solve that?

Steve: That was said much more simply, Leo. Yes. There you go. That's the...

Leo: And this is the wildest solution ever. So he's got an adapter, first thing he's got, he's got an adaptor on the thing; right?

Steve: Exactly. And remember that during that transition phase in headphones, and you still see it around, there is a - it is an adapter from the mini headphone connector to the old-school large one, if you want to plug it into something with a regular quarter-inch headphone socket. So he plugs the little guy into the adapter. Now it's a quarter-inch size. The problem is, you know, you've got a pair of RCA monaural connectors.

Leo: Now, you could go to the store and buy...

Steve: Yes. That would have been...

Leo: ...a connector.

Steve: I'm sure I have some...

Leo: I do, yes.

Steve: ...in my random adapter box, you know.

Leo: But no, this is an emergency.

Steve: Well, it was an emergency, or he had some extra thick, very heavy-gauge copper wire around.

Leo: Yeah. And praise to him for using heavy-gauge copper here.

Steve: Oh, yeah, yeah. You've got to do that. It's sort of like a coat hanger. I mean, you know, it took some effort to bend this copper, this solid copper wire. Again, now, here's where I've already, like, painted myself in the corner. I don't know how to describe this except to say imagine that you created like an eyelet with the copper so that it was an eyelet that then had a straight piece. And then you slid the eyelet down onto the quarter-inch plug. And now remember that the quarter-inch headphone plug has rings; right? It's got some insulator rings. And so it's actually three different conductive zones that you want to connect to.

Leo: You've got your tip, and you've got your ring; right?

Steve: Your ring and your tip, exactly.

Leo: Ring and tip, yeah.

Steve: And so he's got copper wrapped around the ring and the tip, which then goes into the RCA, the monaural RCA red and white for left and right plugs.

Leo: Right.

Steve: Then, to anchor the whole thing, near the base of this he wraps the copper around, and then does a big loop, a big U-shaped loop because he needs to connect to the ground side of these RCA plugs. Anyway, really you need to download the show notes. It's worth it if you download the show notes to look at the first page.

Leo: You'll never figure it out if you're just listening. It's crazy.

Steve: No. But beautiful. I mean, it's just...

Leo: Elegant.

Steve: Absolutely, yeah,

Leo: It's an elegant solution.

Steve: And maybe this was just until the Amazon delivery came.

Leo: Probably, yeah.

Steve: Couldn't wait to get the - had to have an adapter right now for whatever purpose. Anyway, hats off to the anonymous inventor of this approach. I don't think there's a market for this because, you know, Amazon will sell it to you for $3 and deliver it in the afternoon.

Okay. 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. Now, at its present distance of 12.3 billion miles, 2 degrees might as well be 90 degrees. I mean, it's missing the Earth by a long shot. So this meant that no more data could be received, nor could any corrective commands be sent to the probe. It wouldn't hear them. Now, the good news is that as long as all is going well, Voyager has a fail-safe system that was expected to perform an automatic reorientation this coming October. Actually it's on the 15th of October. But, you know, that's still three months away; right?

So NASA wrote: "Voyager 2 is programmed to reset its orientation multiple times each year to keep its antenna pointed at Earth. The next reset will occur on October 15th, which should enable communication to resume. The mission team expects Voyager 2 to remain on its planned trajectory" - let's hope that happens - "during the quiet period."

But then last week we received an update from NASA, on August 1st, 2023: "Using multiple antennas, NASA's Deep Space Network 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 at least 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" - that's the abbreviation for Deep Space Network - "a Deep Space Network 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, communication should resume."

That was on August 1st. That was when we did the podcast last week, on Tuesday. Then August 4th, three days later, "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'" - meaning I guess they cranked the power up to max, or maybe 11, and said, you know, point here.

Anyway, "More than 12.3 billion miles 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 the mission controllers to learn whether the command worked. At 12:29 a.m. EDT on August 4th, the spacecraft began returning science and telemetry data, indicating it is operating normally and that it remains on its expected trajectory."

So, yay. You can just imagine the breath-holding that was going on during those 37 hours. But really, the entire project is an incredible engineering accomplishment. You know, these guys should be so proud of what they have done.

Leo: And you can't, you know, you can't just fire up a telescope and look for it. It's beyond Pluto. I mean, it's not invisible, it's gone.

Steve: It's gone, Leo, it's gone. I mean, we fully expected it to simply dissolve when it left the Earth simulation, you know.

Leo: This is so cool.

Steve: But it's just incredible.

Leo: It's really neat. It's just the neatest thing.

Steve: You know. And when it didn't die after its first planetary encounter, they said, well, huh. Let's keep going. I mean, what the hell? You know?

Leo: This stuff is so over-engineered, I mean, look at Perseverance. Look at, I mean, it's amazing.

Steve: Exactly. The Rovers that just, like, you know, like they got covered in dust, and they wound down, and then the dust blew off, and it came back. Hello.

Leo: Hello. I'm back.

Steve: What did I miss?

Leo: What's up? What's up?

Steve: Wow.

Leo: Just very feel-good story, just great.

Steve: Okay. So everyone who listens to this podcast knows that I often become upset with Microsoft's behavior and with their performance.

Leo: Oh, you had a friend this time, didn't you. I know where you're going with this one.

Steve: 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 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 related to security.

Well, 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 - I'm getting a little excited, as you can see. For what this individual wrote to have any weight and bearing, you need to know something about the posting's author, as I mentioned at the top of the show, Amit Yoran.

Wikipedia informs us: "Amit Yoran is chairman and chief executive officer of Tenable, a position held since January 3rd, 2017. Previously, Yoran was president of computer and network security company RSA." We've heard of them. "Yoran joined RSA during his tenure as CEO of NetWitness Corp., which was acquired by RSA's parent company, EMC, in April of 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 U.S. Department of Defense Computer Emergency Response Team. He resigned from his position at US-CERT in October of '04.

"Earlier in his career, Yoran was a co-founder and CEO of Riptech, which was acquired by Symantec in August 2002. He 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 U.S. Department of Defense's Computer Emergency Response Team. He has a master's degree in computer science." In other words...

Leo: Very impressive guy. Very.

Steve: ... this guy has earned, yes, he has earned some street cred by being, like, 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."

Leo: Oh, boy.

Steve: So here's what Amit wrote and posted publicly on LinkedIn, a platform Microsoft purchased. He wrote:

"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, 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."

He writes: "Microsoft's lack of transparency applies to breaches, irresponsible security practices, and to vulnerabilities, all of which expose their customers to risks that they are deliberately kept in the dark about. In March 2023" - so, right, just this past March - "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" - and he wrote this last week - "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" - meaning end of next month - "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," he says, "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."

Leo: Wow.

Steve: "A culture of toxic obfuscation."

Leo: Toxic obfuscation. I love that.

Steve: Okay, now, 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 some of 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 also 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, CyberScoop, 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 30th, and confirmed it four days later. Tenable asked for an update on June 27th, 90 days later, and was told on July 6th that it was fixed, but Tenable says it was merely a partial fix." Okay, now, where have we heard that before? How many times on this podcast have we noted that someone at Microsoft, who was shown a serious vulnerability by a security researcher, and even given a fix for it, apparently didn't even 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, CyberScoop continues: "On July 21st, Microsoft told Tenable that it would take until September 28th for a complete fix. Tenable agreed to withhold technical details and proofs of concept until September 28th. 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're at risk and therefore can't make an informed decision about compensating controls and other risk mitigating actions.'"

And then we heard from Microsoft. Get this. They wrote: "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.'" Wow.

"Microsoft said Friday in a blog post that the issue has 'been fully addressed for all customers.'" So I guess Yoran got their attention with his blog post in LinkedIn. And, oh, what do you know, just two days later it's been completely fixed. It was going to take until September 28th, but shine a bright light on the problem, oh, look, it's all fixed. They said 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."

Leo: No one's using this. We don't have to worry.

Steve: Yeah, no, don't worry about it. So they say: "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 Department of Justice investigate whether Microsoft's actions in the incident broke the law." Okay, now, I'll just say 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." You know, right. Because Microsoft dismissed this as speculative because, after all, "a delicate balance is required between timeliness and quality, while ensuring maximized customer protection."

They said: "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" - and everybody agrees - "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." Okay. We can hope. Unlikely.

"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." Basically it was charging people for better logging of the activity on their cloud platforms; and they got a lot, they got into a lot of hot water for making money for, like, just offering more logging that, like, cost them nothing.

"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." To which I say, "Exactly."

And finally: "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." Well, yeah, because no one can make them do anything.

So, you know, 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. 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 here? 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 it today 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." Bruce wrote: "A bunch of networks, including U.S. government networks, have been hacked by the Chinese. The hackers," he said, "used forged authentication tokens to access user email, using a stolen Microsoft Azure account customer signing key. Congress wants answers. The phrase 'negligent security practices' is being tossed about, and with good reason," says Bruce. "Master signing keys are not supposed to be left around, waiting to be stolen."

He said: "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 (HSM), 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," says Bruce, "that the details are really bad."

And he says: "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." He says: "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."

Bruce finishes: "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 U.S., 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. Well, I've got a news flash for you.

I also thought of Bruce Schneier recently in another context because I love to quote one of his pithy observations, which is "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. And with that, keys, Leo. We're going to listen to you telling us about an advertiser.

Leo: I'll type out an ad. All right, Steve. Yeah, boy, when I read that screed I thought of you immediately.

Steve: Yeah.

Leo: You know, I think part of the problem with Microsoft is just that they have so many, it's such a big install base. I think anybody with a big install base like that would have similar problems. But anyway.

Steve: Well, yes. When I saw him note the percentage of zero-days that were theirs, it was like, well, yeah.

Leo: Of course. It should be 100%. It could be 100%; right?

Steve: Well, or if you were to break Microsoft up into individual organizations, each responsible for one of their different products, then, you know, the zero-days would be spread out the way they are among everyone else. So, you know, it is the fact that they're such a behemoth. On the other hand, that's a problem. I mean, I used to know people at Microsoft who were really nice people. Brad Silverberg was a great guy. And, you know, the really good guys, they're gone. You know, they took their money and left. And now we're just left with kind of an unaccountable monstrosity.

Leo: There's always Linux and Mac, if you ever want to change.

Steve: Hey, I love my unaccountable monstrosity.

Leo: I know you do. I know you do.

Steve: Well, yes. So, okay. As I was saying, Bruce Schneier reminds us always that attacks always get better, they never get worse.

Leo: In the sense that they always get better for the bad guy. They always get worse for the good guy.

Steve: Yes.

Leo: Let's be clear.

Steve: Yes, exactly that. Right, right. Okay, so in this case, this explains exactly what happened recently with the classic attack of listening to someone typing on a keyboard. Although significant controversy, understandably I think, 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, as we said, they always get better.

Last Thursday, on 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. 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.

They said: "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."

So this is a phenomenal level of recognition, 95%, for an outboard external microphone that's simply listening to keystrokes from a keyboard nearby. And to only lose 2% accuracy when significantly compressing the audio through Zoom is equally astonishing. Imagine being able to process the recorded sounds of someone typing after the fact through a compressed connection to be able 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 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 all 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 just begun to be explored. I expect we're going to be seeing some very significant discoveries in the future relative to security, once these relatively new capabilities become more widely available. And 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.

Leo: They've really made this their thing, haven't they.

Steve: Oh, my god. Once they add deep learning modeling to their many bags of tricks.

Leo: Oh, you're right. You're right.

Steve: Yeah. I mean, oh, boy.

Leo: Wow. Wow.

Steve: Yeah. I don't think we're going to have to wait long, either, because these guys tend to be on top of things.

Okay. So some feedback from our listeners. Rusty, tweeting as @rusty0101, he has another take on the "in the cloud or on the ground" discussion. He said: "Listening to this week's SN, with the discussion of running things in the cloud. I'd noted 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 counter argument," he said. "Additionally," he said, "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 the way."

Okay. So 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 everything will be. I think the reality is there's probably a place for both. And that's the point that I had intended to make.

I also sort of liked that whole notion that we're losing the inherent distributed nature of the Internet, which is one of the ways that it got so much of its strength and robustness. You know, we all feel that everybody aggregating around a Chromium-based web browser is not a good idea because we end up with a monoculture. So it certainly is the case that if anything like a really bad problem ever hit AWS, it would impact a huge portion of the users of the Internet, the providers of services. And so that's not the way it used to be. But we'll see what happens.

Alan C. Bonnici, he said: "Hi, Steve. I heard you speak about Authy and decided to give it a try." Actually, I was not talking about Authy.

Leo: Yeah, you're not in favor of Authy, yeah. I've mentioned Authy back in the day, it was from Twilio. But you and I both have come up with better solutions, I think, since.

Steve: Right. Anyway, so basically he's talking about some two-factor authenticator, you know, a TOTP-based approach. He said: "So I reset the two-factor authentication 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."

Okay. So Alan did follow 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, you know,

authenticators, 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 one that was just technically obsoleted, 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, and there there's no fudge factor, of course, because that keys the pseudo-random sequence.

But they must also agree upon the current time of day. In today's Internet-connected world, it's easy for devices to be within very close clock agreement because we're all able to get the time from the Internet. But it's also reasonable to make some allowances for them not being. And of course there is the whole man-in-the-middle, literally a person in the middle, reading the code from their phone and then turning around and transferring it by keying it in on their keyboard. And so that's going to introduce a little bit of a delay, and thus some need for fudging.

So that's all that was going on was that the authenticator would much rather be a little bit tolerant and accept, you know, as long as only one use of the code is possible, that is, as long as the code when used cannot be used immediately again, I don't see any reason for giving the user a bit of leeway. And also cutting down on user frustration of the system saying sorry, that code's invalid when you're quite sure that it just was.

Joe LaGreca, he said: "I'm finally ready to leave the Google Chrome browser. Which browser do you use or recommend?" And he said, "Firefox?" And of course Firefox is where both, Leo, you and I are, and I'm completely happy with the choice.

Leo: Me, too. I'm never changing.

Steve: Yeah.

Leo: Well, I shouldn't say that, but...

Steve: I know, it's hard to say never. So some time ago I did try using Bing because I was curious about tabs down the left-hand column of the browser. But frankly, I was stunned when I encountered some sites that it would not render. I thought, what? You know, who knows? Anyway, I had been using Chrome for a while, again, just to see how it compared to my longtime previous use of Firefox.

Today, having satisfied a little bit of my wanderlust, I'm back to Firefox. And, you know, many things about it are just exactly 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. And I use that when I'm working on the podcast in order to change locations and have all the tabs sent to a different location. So anyway, yes, I am 100% Firefox. And when we get to talking about today's topic, which we'll get to shortly, you'll learn another reason why it continues to be my choice.

Someone whose name is Seven in Twitter, he said: "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 extracurricular listening. I've since subscribed to Club TWiT" - thank you very much.

Leo: Yay.

Steve: "...and returned to attending weekly services."

Leo: Good.

Steve: He's back. He's got the religion again. He says: "I don't know if my question will be simple enough to address in a DM, but perhaps at least with a suggestion where to start. After receiving a notification that one of my accounts was compromised INCLUDING" - he has in all caps - "the password, I have come to fully realize that no passwords are safe. Period. I use 2FA wherever possible, but of course two-factor authentication support isn't consistent across all services. Is there a best way to simplify the process of not relying on passwords alone?" He says: "Is there a simple answer to the question WWSGD?"

Leo: What would Steve Gibson do?

Steve: Exactly.

Leo: I love it.

Steve: Took me a moment...

Leo: I need a button.

Steve: Took me a moment to parse WWSGD, but it's clear that that's what he meant. Okay. So I'll expand a bit on Seven's question 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 that spreads is that the appearance of new and better solutions doesn't automatically kill off the older and less secure solutions.

Despite the fact that two-factor authentication has been widely available now for, what, a decade or two, 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 to some degree. And now we have Passkeys which represents yet another step forward. But will Passkeys kill off two-factor authentication and passwords? No. Over time, more sites will be offering passkey support, maybe again as the perceived need at the higher end of the security spectrum manifests.

Eventually, support for two-factor authentication and Passkeys will be baked into servers, and servers will be taking more responsibility for authentication. After all, servers were doing HTTP without any security not that long ago. And it's only recently that sort of by universal agreement the whole industry just said, okay, we're just going to all decide we need to have it. And EFF and Let's Encrypt made that more practical by not having people paying for certificates constantly.

But we also know that even as two-factor support and Passkeys becomes more available, 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, after all, to create an account as a requirement just 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 turn out to be relevant. We'll see.

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 "I forgot my password" link appearing underneath every password prompt makes a strong case for not even bothering with remembering any passwords. Just bang on the keyboard for a while when you're creating an account, then click the "I forgot my password" link whenever you want to come back. And I actually think that people would probably 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?

Leo: Usually you have to jump through some hoops, though.

Steve: Well.

Leo: I mean, it's not just the password at that point, is it? I guess it depends on the site.

Steve: Yeah. It's your email. Wikipedia now has an entry on the topic "Security Theater," 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 rightfully, 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 always 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 named "total authentication bypass" - any entity who is able to obtain access, even transiently, to your email flow obtains unfettered access to your entire online life.

Leo: That's why you've got to secure your email. Let's face it.

Steve: Exactly.

Leo: That's really a...

Steve: Email security is really ultimately important.

Leo: There are a number of sites now that don't use passwords, that just use email. So they say, what's your email. You say it. They say, okay, we're going to send you a link. Click that, and you're going to log in. That's become more and more prevalent. I guess that's not more secure, but that's just in response to people not doing passwords right.

Steve: Right, yeah. I mean, I'm sure that their poor support people said they got just tired of receiving email, I forgot my password.

Leo: Yeah, right.

Steve: So they said, okay, screw it. I mean, after all, they were going to send a password recovery link anyway.

Leo: Right.

Steve: So short-circuit the whole process.

Leo: Right, and just assume everybody's forgotten their password.

Steve: And I have to say, Leo, when I was putting this together yesterday, I sort of had to pause and think, okay, what's wrong with just using email? Why don't we automate, somehow automate, like, the email link receipt and forwarding process somehow?

Leo: Yeah.

Steve: Like, you know? Because as I said, that's what this all devolves into anyway.

Leo: Well, we've always said the weakest link in your password recovery process is your security. Right? If there's a way to get a password, that's the weakest link. So if

it's email, it's email. So you might as well just say, okay, we're just going to use email, not bother with anything else. Especially if you're going to ignore two-factor.

Steve: Everything else gives you an oops, it doesn't work.

Leo: Email me, yeah.

Steve: Or I'm standing on my head, so send me an email link and then let me in. And you can always get in.

Leo: That is a good way, that is a good reason, I guess, to use these unique email generators; right? So that every site has a unique email. Because, well, mostly it's just - really it would...

Steve: It forwards in, it always forwards into your master email, yeah.

Leo: Yeah. But mostly what it says is keep your gosh darn email secure as you possibly can. Have two-factor on that, yeah.

Steve: Yup.

Leo: And I don't think Google, if you lose your two-factor, just says oh, fine, we'll email you something. I think you have to jump through more hoops than just that. I believe so. But I might be wrong. I haven't tried it.

Steve: Yeah.

Leo: Have an email provider that really lets you lock the sucker down.

Steve: Well, and actually this is why one of the features that I built into SQRL was after you got comfortable with SQRL, there was a checkbox in the config where you could say, "Please disable all other authentication."

Leo: No recovery, yeah. I don't want any...

Steve: Period.

Leo: Yeah, there are no other ways to recover.

Steve: Otherwise you don't actually get any additional security.

Leo: Yeah. I'm now looking, what happens if I lose my 2FA token on Fastmail? I'm just curious because that's where everything is. It uses your phone. You get a code sent to your phone.

Steve: Okay. That's...

Leo: Better than nothing; right?

Steve: Better than nothing. Although, you know, SMS is not great.

Leo: Right. That's another - so one more thing, another in the list of things to do is make sure that you have either PIN-protected or somehow protected your cell phone account so that somebody can't SIM jack you. And actually the FTC is moving forward on regulation on that because that's really important. I'm trusting Google Fi not to let people, you know, steal my number. I think Google's pretty good at that.

Steve: Trust is good.

Leo: Well, yeah, but trust is good if it's merited.

Steve: I know. That's kind of what I meant.

Leo: It's got to be earned first. All right. Now I have no slogan whatsoever for this ridiculously named, but I think very potent technology, Global Privacy Control.

Steve: So 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. It had their TechCrunch logo, which I've always thought was kind of cool, in the upper left-hand corner. And the headline read: "Review your Global Privacy Control preferences." And the notice said: "You're using Global Privacy Control (GPC). This leads to a lower-quality experience on Yahoo..."

Leo: What?

Steve: "...by blocking certain editorial content, including embedded tweets, YouTube videos, and third-party ads that are relevant to your interests." Huh. Okay. That was interesting. It gets better. "To enhance your Yahoo experience, allow us to share and sell your personal information."

Leo: This would be so much of a better experience that way.

Steve: Right, right, because, you know, it's clearly in my best interest to have my Yahoo experience enhanced.

Leo: Sure.

Steve: By allowing them to share and sell my personal information. No doubt about it.

Leo: Makes it so much better.

Steve: And there's more. They continue: "This includes technical identifiers like your IP address and cookie IDs, but does not include things like personal emails or contact information." And then this notice concludes with: "This won't affect your GPC settings for other websites, and you can always change this preference in Privacy Controls." And then at the bottom of this little pop-up I was presented with two options, Allow or Don't Allow. And you can probably guess which one I chose.

Okay. Now, there are several bits of good news here. One is that someone made them do this.

Leo: You made me.

Steve: That's right. They didn't want to do this. They didn't want to ask me this question. 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, 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 is 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. We'll get to all that later.

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

Leo: I like that. Might as well.

Steve: Wikipedia explains this by writing: "In 2010, AOL acquired TechCrunch for approximately $25 million."

Leo: Yeah, they've got to make that back somehow.

Steve: Uh-huh. "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."

Leo: Sigh.

Steve: So that's why I went to TechCrunch, and I was talking to Yahoo. 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. The term in that pop-up was also a link which took me to a Yahoo page titled "Collection, Use, and Linking of Technical Identifiers," where they write: "Yahoo uses" - meaning TechCrunch because that's a parent, global - "uses different technical identifiers to make its consumer services available on most platforms, browsers, and devices." And again, right, "to make its services available." Well, they work just fine without those, but no, we need those technical identifiers. Sorry, didn't mean to editorialize. They continue: "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.

"Now, of course, when you cryptographically hash something like your phone number, it becomes an identifier for you, not of you, but enabling cross-everything association of you. And of course email addresses, phone numbers, account identifiers and so forth. Household-based identifiers; IP addresses; probabilistic non-unique identifiers, right, like fingerprints. They can't be sure that's you, but eh, close. 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.'" Okay. That won't work for me, but you get the idea.

"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." Oh, no. "Certain apps, mobile devices, or installed software, where permitted and applicable; certain Internet-of-Things devices." In other words, every effing thing we could...

Leo: Everything you've got.

Steve: Yes, we paid. We have rooms of techies trying to find some way to track you, hook onto you, see where you went, what you're thinking, what you like, what you're doing. We want it all.

Leo: This is a great document because they're basically, I mean, they figure nobody's going to read this. But they're revealing everything. Right? They're just...

Steve: Yes.

Leo: I mean, it's really telling, yeah.

Steve: Yes. It is every frigging thing anybody we ever had was able to think of. So they say: "The collection methods for technical identifiers" - and remember, that's where they wanted me to click OK, yeah, let's go. "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" - that's rather generic, but okay - "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" - again, generic - "local or remote application programming interfaces" - same - "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 services providers" - maybe alien technology, who knows? - "and 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." Just click YES here to proceed. Oh, my lord.

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.

Leo: Good point.

Steve: 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," thank you very much, at the cost of my, my family's, and my company's privacy.

Okay. 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 subpages for California, Colorado, Connecticut...

Leo: Interesting.

Steve: Uh-huh, as I mentioned before. And there's also one for Virginia. The California page I was taken to has two switch settings. Oh. And I should also note that somewhere I did see something about IP geolocation, which is obviously how they decided that I'm in California.

Leo: Right.

Steve: Because I hadn't logged into them. So that's going on also. And I tried to go back and find it, but I couldn't. So anyway, 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 was set to "Don't Allow." That corresponded to "Allow the Sale and Sharing of My Personal Information." That's now off. 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." And the page explains - and Leo, if you go to TechCrunch and scroll to the bottom and click on Your Privacy Choices, you'll see this.

So under this "Limit the Use of My Sensitive Personal Information," the page explains: "In connection with providing our services, we may use sensitive personal" - listen to this - "sensitive personal information such as precise location data and email content data. Among other purposes, we use such data to help understand" - oh, and both of yours are on - "to help understand your interests so we can show you more relevant ads and content."

Leo: Look how long it's taking to turn it off. This is definitely wait for 30 seconds before refreshing page. Holy cow.

Steve: Actually, I think you'll find that the Firefox has died now. I don't think you can close that tab. I don't think there's anything you can do.

Leo: Oh, even worse. Oh, my gosh. I broke - you broke Firefox. Wow.

Steve: Yeah, they really don't want you to turn that off.

Leo: Oh, my goodness.

Steve: It's there. Uh-huh.

Leo: Don't flip that - how did you turn on your Global Privacy Control? Is that now a switch in Firefox, or is it automatic?

Steve: Yes.

Leo: Okay.

Steve: Yes. It's a switch in Firefox. And I presume that if you find it and turn it on, and then go back to TechCrunch, you'll get the pop-up experience that I got.

Leo: Yeah. Yeah.

Steve: So that would be interesting to see.

Leo: Do I have to do about:config? That's how we did it last time, when we talked about this last time.

Steve: I'm sure it's there, yes, about:config. And then maybe just in the search box maybe "GPC" or type "global privacy" and see if it finds it.

Leo: Yeah. Okay.

Steve: So they said: "To opt out of the use and limit our use of such information" - that is, my sensitive personal information - "to only those purposes" - get this - "only those purposes permitted by California law, select 'Limit.' This may make the content and ads that you see less relevant to you." Oh, boohoo. Okay. So after doing some additional research I figured out what's going on for anyone in California, and to some other degree the other states - "why Yahoo! is showing two switches. Under the California" - and you found it in Firefox?

Leo: I found it.

Steve: Yeah.

Leo: And now I want to enable it. And now it's enabled. Now let's go to TechCrunch, see what they say. Wow. You know, I go to TechCrunch every day. Nothing so far.

Steve: At least get the pop-up. For me it grayed out. So it'll see, it'll be interesting to see if at some point, since you do go every day, and presumably it sees you in California...

Leo: Yeah, it does, because I'm getting California Privacy Laws.

Steve: Ah, ah, good point, exactly. Right, right, right.

Leo: Yeah. Yeah. But every time I turn this off it crashes. How about if I limit the use of my sensitive personal information. Oh. Oh, this is so bad.

Steve: This is so bad, Leo.

Leo: The FTC needs to jump right in on this. This is terrible.

Steve: It's so bad.

Leo: Oh, my goodness.

Steve: Okay. So I figured out what's going on in for anyone in California and why Yahoo is showing two switches for us. Under the California Consumer Privacy Act (CCPA), California consumers have the right to opt-out of the sale and sharing and the use of their personal information. Okay. Those three things: sale, sharing, and 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 Californians want to prohibit the use of their personal information, beyond its sale and sharing, which can be done globally with the GPC setting, that will still need to be done on a site-by-site basis. So, and Leo, one thing you might try is restarting Firefox.

Leo: Oh, just close it all out, yeah, you're right. Start over.

Steve: Yeah, having put that on. Because now that your GPC is on, that upper switch in TechCrunch should really not be on. Because your browser is now shouting no, no, no.

Leo: Do not. No, no, no.

Steve: And California law requires that this be obeyed.

Leo: We can only hope that there are massive penalties for not because I restarted it, and I'm still getting a normal page. Let me try one more time to switch that switch. Geez, Louise.

Steve: Or see if the switch is already flipped for you.

Leo: Oh, yeah, good, okay. Yeah, let me check that.

Steve: Yeah, it ought to be turned off.

Leo: Oh, okay. I didn't get any pop-up, though, of saying anything.

Steve: No, no.

Leo: All right. Let's see. Is it already turned off? No.

Steve: Oh, boy.

Leo: No, they're just ignoring it. And of course - so I clicked a button. We'll be back later. Bye-bye.

Steve: Okay. So here's how the GPC, the topic of today, places itself, its need, and the role it's filling. The GPC's formal specification explains: "Building websites today often requires relying on services provided by businesses other than the one which a person chooses to interact with. The result is a natural consequence of the increasing complexity of web technology and the division of labor between different service providers. 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." And then the spec quotes the California Attorney General, saying: "Given the ease and frequency by which personal information is collected and sold when a customer visits a website, consumers should have a similarly easy ability to request to opt-out globally. This regulation [in California] 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."

So the spec says: "This specification addresses the issue by providing a way to signal, through an HTTP header or the DOM (Document Object Model), 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." And also, subsequent to the passing of the regulation, the Attorney General formally acknowledged or asserted that this GPC signal is within the scope of what California considers a global assertion that an individual wants this to hold.

Okay. So what's also annoying, though, 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, and for example yours hadn't been, Leo, if a browser is broadcasting it, it's because this is what its owner meant 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, differs 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 available - 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 firewallsdontstopdragons.com. And the guy gets most of this right, and I'll note where he made a couple mistakes.

But he writes: "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 newish 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 also may 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 they visited that they 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 now made you more trackable.

"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 bit 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," he wrote, "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."

Okay, now I'll pause here just to note that just as the terms "Do Not Track" and "Global Privacy Control" sound like different things, indeed they are. 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 inasmuch 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, however, about how to enable global privacy control.

He writes: "This is not a slam dunk. For one thing, there is no U.S. federal law requiring companies to respect GPC. Also, the GDPR interpretation of GPC sadly seems a little weak." Well, give them a little time. I bet they'll fix that. "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 set the GPC flag?

"If you're lucky enough," he writes, "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 remind everybody not to

have your personal information sold or shared. But otherwise he's right. He says: "Then, whenever and wherever this request is required to be honored, you'll get the benefit."

And he finishes: "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." And so the site is globalprivacycontrol.org. Just go https://globalprivacycontrol.org. And there is a little banner that you get at the very top. You should see a green dot to confirm that your browser is currently sending that out.

Leo: Oh, mine's not. Which is why I perhaps was getting...

Steve: Yes, good. So now google how to enable GPC for Firefox while I keep going.

Leo: I turned it on in the about:config, but I guess that wasn't enough. I'll have to...

Steve: There must be something else, yeah.

Leo: Yeah.

Steve: So Colorado's Privacy Act (CPA) and Connecticut's Data Privacy Act (CDPA) both recently went into effect...

Leo: Oh, there's two options, that's all. I have to do both. Okay.

Steve: Ah, good. They both went into effect on July 1st; right? So only a little over a month ago. And like California's CPRA, those states' legislation require companies to honor the GPC. But Virginia's apparently doesn't, so that's unfortunate. Virginia has some laws, but it's not - it didn't anticipate GPC. I see you've got two things there now.

Leo: There we go. There, now they're both enabled. Ah, okay.

Steve: And so first go to globalprivacycontrol.org.

Leo: And make sure, yeah.

Steve: And see if you get a green dot.

Leo: Mm-hmm. Okay.

Steve: And then back to TechCrunch.

Leo: Yeah.

Steve: So today, Firefox, Brave, and the DuckDuckGo privacy browser all support the GPC. As for browser extensions for Chrome and other Chromium browsers that do not yet natively offer - yup, you're green now.

Leo: Green light. All right.

Steve: There's Abine, which is from the DeleteMe people, I think a sponsor of the TWiT network. Disconnect, OptMeowt - where Meow, remember, is the cat's meow, so O-P-T-M-E-O-W-T, OptMeowt - and Privacy Badger, which - and that's a name I just hate, but that's from the EFF. Badger, really, Badger? Couldn't like, you know...

Leo: It's like a honey badger. It just doesn't give up.

Steve: How about Privacy Fairy or something, I don't know, but not badger. I have a link here at the end of the show notes to the GPC page which maintains a list of available extensions. So at the moment those in California, Colorado, and Connecticut, the three C's, have the advantage of state laws which compel compliance with their residents' GPC request.

Leo: Bravo.

Steve: It doesn't appear that websites serving Virginians, which does have similar privacy laws, are similarly bound to follow the GPC signal. But what we need now that the GPC exists and is gaining some traction will first be for additional states, which is easier to do, to step up and add their voices with their own statewide legislation. Then we need the U.S. federal government to take this initiative national. At that point everyone will be on equal footing with the ability to opt into this, and thus opt out of having your personal data sold and shared. And then we can imagine a day when a federal law won't require the presence of a GPC beacon. Well, we can dream; can't we? Leo, you and I may not see that day, but maybe our grandchildren. Well, I don't have any, but, you know.

Leo: Nor do I, yet.

Steve: You will.

Leo: Yeah, yeah. Very good stuff. And gosh, I hope you're right. That's all I can say.

Steve: If nothing else, it doesn't look like it's going to die. California, Colorado, and Connecticut are requiring it. As a consequence of my having turned it on, I got the pop-up with Yahoo saying are you sure? And, boy, am I glad I was asked. So it'll be interesting to hear back from our listeners as they experience the effects of having this turned on. I can't imagine why everybody would not turn this thing on.

Leo: Yeah, yeah. I'm going to go through all my browsers. Obviously, you can't in Chrome. But are there Chrome plugins that will let you enable it?

Steve: Yes. Abine has one, and Disconnect is one. And that horrible Privacy Badger thing, I mean horribly named, I'm sure it's a good thing.

Leo: They'll let you turn it on in Chrome because Google doesn't want this.

Steve: No.

Leo: But we're going to make them. And Yahoo, it's still spinning.

Steve: Oh, my lord.

Leo: Yeah. Who knows, you know. I turned off all of the tracking protection because, you know, I have to say Firefox has very good tracking protection, disabled all of that. Turned off UBlock Origin for it.

Steve: That was just what I was going to ask.

Leo: Yeah, no, turned it off.

Steve: What about starting it up, there's like a startup with no extensions.

Leo: I might have to do that.

Steve: Where it starts up clean.

Leo: And we also have quite a bit of stuff in our network, our company network.

Steve: Well, Leo, it happened to me. I couldn't do it either. I got that thing, I couldn't even close the page. I had to go to Task Manager and abort the process to get out.

Leo: Wow. I can at least close the page. You know, we have so many, and rightly so, perimeter protections on here and stuff. I just don't know if it's something we're doing or something they're doing. It doesn't look like they're compliant right now because when I go to TechCrunch.com and click the privacy choices, those switches are still on, and I cannot turn them off.

Steve: Yup. And your browser is broadcasting a, you know, eff off signal.

Leo: It's really frustrating. So frustrating. Okay. But Steve, see, if we listen to the show, right, we know. We are informed. And that is the first step into changing the world. That's because of this guy right here, Steve Gibson. GRC.com is his website. He has the podcast there.

Copyright (c) 2014 by Steve Gibson and Leo Laporte. SOME RIGHTS RESERVED

This work is licensed for the good of the Internet Community under the Creative Commons License v2.5. See the following Web page for details: http://creativecommons.org/licenses/by-nc-sa/2.5/