Next Gen Ad Privacy

Description: This week we check in on Texas, and on the Kazakhstan government's attempt to be their own CA. How did that work out for them? We note a troubling increase in attacks on the open source software supply chain. Google's announced plans to add data breach notification to Chrome. We look at a surprising Apple iOS v12.4 regression (whoops!) and at another Microsoft RDP component in need of updating. I update our listeners on the state of SQRL (another of its documents is completed) and on SQRL presentations past and future. I share some news from my ongoing file sync journey. We conclude by looking at some very interesting and promising moves as browser-based advertising matures from the ad hoc mess it has always been into a privacy-respecting Internet citizen.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-729.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-729-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. Of course we've got security flaws - what would Security Now! be without security flaws? - in Webmin. Another RubyGems flaw. We'll talk about the ransomware attacks in Texas. Steve has an interesting theory about what's really going on. And then Google's proposal for ad privacy. Why Google thinks cookies should stick around. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 729, recorded Tuesday, August 27th, 2019: Next Gen Ad Privacy.

It's time for Security Now!, the show where we cover your safety online, your privacy, very much privacy online and all of that jazz with this guy right here, Security Now!'s own Steve Gibson. Hello, Steve.

Steve Gibson: Hey, Leo. Great to be with you once again.

Leo: Yes, for a thrilling [crosstalk].

Steve: We're going to be recording in two weeks on Monday, rather than on Tuesday, for any avid stream followers, because of the big anticipated - I guess it's not anticipated anymore, it's real, isn't it, the Apple event on Tuesday.

Leo: Yes. Well, no. Still a guess.

Steve: Ah.

Leo: Still a guess. So Apple has not yet sent out invitations; to my knowledge hasn't confirmed it. But I'm hoping, we're expecting September 10th. And if that's the case - now, if it's not the case, we won't move.

Steve: Yeah, good, okay.

Leo: Yeah.

Steve: And we'll know by then, so...

Leo: Well, we should know, well, I hope we'll know by then. We should know by Thursday is what Lory Gil said. She thought we'd get the invitations Thursday.

Steve: So Episode 729 for August 27th, the last episode of August, our second podcast of our 15th year. And this is interesting. There have been, well, we often talk about the whole tension, I mean, the increasing tension that exists between online advertising and the propensity it has for tracking, and kind of the icky feeling that the idea of being tracked across the web gives people. I've had, in like watching talking heads shows, from time to time they'll talk about how creepy it feels for them to be texting about something, and then they go to the Internet, and they start seeing ads for that. And it's like, okay, wait. I'm not sure that actually happens that way. But we know that tracking is an issue.

Apple has been making a series of announcements over the last couple years about standing up more for their users' privacy. We know that Firefox has. Well, of course the biggie in the industry is Google with Chrome. And also Google, who purchased - what was the big ad company, used to be independent and is now part of Google...

Leo: Yeah, I want to say Overture, but that was Microsoft. I can't remember what they bought, yeah.

Steve: Yeah, I can't think of them. Anyway, it was the biggie. So what finally has happened is that Google has recognized that, unless they do something proactive, they're in danger of having their own revenue impacted by other browsers proactively taking measures. I mean, and to Google's credit, we've been talking about how Google is trying to react against fingerprinting. Chrome and the fingerprinters have been going back and forth.

Well, it turns out that it's looking like what we're on the cusp of is the adoption of, and the development of, and then adoption of some technologies to essentially legitimize profiling without tracking, the idea being that, for example, where we've only had cookies, and cookies would allow you to get tracked across the Internet, there really hasn't been an alternative to tracking individuals and building profiles around individuals, which is what privacy advocates are so upset about.

But when you think about it, I mean, if the advertiser knew you anonymously as a member of a cohort of similar people such that they could get what they want, which is targeting, without knowing what we don't want, which is us, that is, at the individual granularity level, then everybody could win. And to the degree that it's true that ads run the 'Net, and in fact Google has some instrumentation where they're claiming to have seen a 50% drop in advertising return for anonymous ads versus targeted ads, then that's significant.

Anyway, this week's podcast is titled "Next Gen Ad Privacy." And we will wrap up the podcast by talking about in more detail some of what Google is saying, what Apple has said. It turns out that even Cloudflare has a technology which they developed along with a plugin with a summer intern a couple years ago to minimize the number of times you have to do the CAPTCHAs on Cloudflare to say that you're a human and not a bot. And so I think a really interesting podcast. And this feels like the right thing. Advertising and tracking had just been ad hoc. What really, and our listeners know, those who have been listening for 14 years, that I've always been annoyed because this wasn't what cookies were meant for. And it was just the idea of the abuse of a mechanism from a technology standpoint always got under my skin.

And so the idea of saying, okay, look, advertising is here, it's real, it's not going anywhere, so let's embrace it, if we can do so in a way that respects our privacy. And we know that cryptographic technology, as I've often said, can do anything we want it to. Turns out it can solve these problems, too. So we're going to talk about, first, we're going to check in on Texas and find out what happens when you mess with Texas.

Leo: Uh-oh.

Steve: Also on the Kazakhstan government's attempt to be their own certificate authority, how did that work out for them. We'll note the troubling increase in attacks on the open source software supply chain that's getting worse. And also on Google's announced plans to add data breach notification to their own browser. We know that Firefox had done that previously. Also we look at a surprising Apple iOS version 10.4 regression which was rapidly fixed in 10.4.1, and at another Microsoft RDP component which is in need of updating for, oddly enough, for Android users.

Then I'm going to update our listeners on the state of SQRL. Another of its four documents, that is, the third of four is completed. And also on SQRL presentations past and future. I also have some news on my ongoing file sync journey. And then we'll wrap up by talking about what I think is the soon to come legitimization of advertising and its need to understand something about who we are without telling them who we are. And so I think another great podcast for our listeners.

Leo: Yeah, it's a really interesting topic. And I've also read some comments from people that say what Google is doing benefits Google. But, you know...

Steve: Well, of course it does.

Leo: Yeah. They're not getting rid of tracking. They just have other ways of doing it besides cookies.

Steve: But it's different than that, though. People who said that did not understand what Google is doing. And so we'll talk about that.

Leo: All right.

Steve: So our Picture of the Week, Leo...

Leo: This reminds me of another one we've done.

Steve: We've done several like this, yes. The other one I loved because you could see the tire tracks that were rolling around on either side of it. And we've done several of the locks, the combination locks that had the combinations printed on the window. I was looking at the latch on this, and I was thinking, I hope that's not a lock on that gate so that, you know, presumably it keeps it from swinging through, so it only opens in one direction, rather than them imagining that they can lock that gate. But really, I mean, could anybody explain this? I mean, I just, like, I don't understand.

Leo: It's to keep baby strollers out or something. I mean, it's got to be, I mean, there's a gate across a path, but of course there's no wall attached to the gate, so you can just walk right around it.

Steve: Yeah, I mean, or bicycles?

Leo: Yeah, I think it's kind of...

Steve: Not a car. It wouldn't block a car. It wouldn't block anything.

Leo: Not really. Anything that wanted to get in.

Steve: Because, I mean, it's narrower than the space on either side, so it's not like it's going to block something of a certain width.

Leo: Maybe they plan to build a wall.

Steve: Maybe it's so you can get arrested if you're on the wrong side; and they can say, look, there's a gate.

Leo: Yes, we told you.

Steve: You didn't go through the gate, did you. And it makes it very clear. It's painted yellow.

Leo: It's very funny.

Steve: So, like, see the gate. I just, really, I just...

Leo: It also clearly was added after the path was built because there's a patch of concrete where the gate is.

Steve: Yeah. So they thought, you know what this path needs?

Leo: A nonfunctional gate.

Steve: We forgot. Someone, like, we forgot the gate. Out in the middle of nowhere, out in the open. We're going to put a gate here just because. We have an extra one, and so we have to put it somewhere. I don't know.

So Texas ransomware update. We of course talked about the number, 23 townships in Texas. Now, with a little bit of reading between the lines I'm developing a theory of what happened, and it sort of puts - it changes the profile of what we thought. Everybody got off on, like, the coordinated attack against 20 or 23 cities at once, and on my god, this ups the ante and changes the territory. It's looking like maybe that didn't happen.

NPR carried a story of this, and I've been checking back because one of the weird things is there was so little information that was coming out. We've been talking about all of these attacks. Certainly the ones in Florida were interesting. We got lots of interesting details. This one was odd because, like, due to its lack of specificity. It was like, yeah, this happened. But no one could get any information about it. And again, this also follows my theory as to kind of why. But theirs was the only reporting, NPR's, which caught the mayor of one of the Texas towns stating that the attackers were demanding $2.5 million.

Leo: Wow, that's a lot more than the Florida attacks.

Steve: That is. And of course, as we know, it took out - it did something to - the number dropped also, I guess it was just overstated, from 23 to 22. So that's now the number that I'm seeing reported is 22 municipalities affected. So they're wanting $2.5 million to provide the decryption keys for these machines that were successfully attacked and encrypted.

And what we know is that, as of last Wednesday, only two cities have come forward to say that their computer systems were affected. Officials in Borger in the Texas Panhandle said the attack has affected city businesses and financial operations. Birth and death certificates are not available online, and the city cannot accept utility payments from any of its 13,250 residents. City officials said that "Responders have not yet established a timeframe for when full normal operations will be restored." And then Keene, Texas, a small city with 6,100 residents outside of Forth Worth, was also hit, officials announced. That city's government is also, and this was the first clue I had, also unable to process utility payments.

And to me that suggested that these two small cities, also notable because of their small size, might be sharing outsourced payment systems to handle their utilities. And then that appeared to be confirmed by Keene's Mayor Gary Heinrich, who told NPR that the

hackers broke into the IT system used by the city and managed by an outsourced company. And so that, which is what he said, also supports many of the other municipalities targeted.

So I think that's what happened. I don't think this was a big, mass, ooh, the game has changed coordinated attack because we also heard from the three big Texas cities, like Dallas and Fort Worth itself, yeah, there were, like, three that were asked. And they said, as far as they knew, nothing in their city had been attacked. So what the model appears to be, there are small towns in Texas that aren't big enough to do their own thing, so they outsourced these services. And I'll bet it was that servicing company.

Leo: You bet.

Steve: Yes. And that explains the simultaneity of this, how it was suddenly, like, oh my god, everybody on Sunday morning woke up and, well, everybody, 23 townships that this happened to. So anyway, I think that's what's happening. We still haven't heard about the payments. It's looking like it won't be up to the cities. The cities, these 22 cities are just waiting for their systems to come back. They're not theirs to fix. They're this outsourced company.

And that also explains why there's sort of been this weird silence. It's because stuff that they were contracting for, the services they were contracting for, stopped working. And they're like, hmm. Well, we hope they're going to come back soon. And it's also - it's looking like they're not going to be paying any ransom. It's going to be this outsourced company, or hopefully that company's insured, blah blah blah. So anyway, that's the update on Texas.

We talked about this Kazakhstan certificate, where the reporting was murky and sort of self-contradictory. Some of the country's users apparently received themselves mixed messages. Some citizens of Kazakhstan were told that they would need to install the government's private certificate if they wished to retain Internet access. And of course we all know what that means. We've worried ourselves in the U.S. that at some point we may find our ISPs telling us that we need to install certificates if we want to continue receiving service. And I dread the day that that happens.

Although, boy, after the pushback that we're about to talk about, maybe that's looking less good or less feasible, less likely. Some users reported that indeed they had been cut off from the Internet until they had installed that certificate, whereas others said they had not installed anything and that their service was just fine. Then we heard that only the area surrounding the capital city would be affected. But then others reported that they were in the capital, and everything was fine. So it seemed really weird. We also heard that Kazakh ISPs were forcing their customers to install this government-issued root cert on their devices to regain access to Internet services.

And we found the certificate. So there definitely was one. In a prior episode of this podcast we showed the certificate and put it side by - we showed the fake Facebook certificate and the real Facebook certificate side by side and how it just said something like Internet security service or something on the fake one, whereas the real one said Facebook, Inc. And it was a clone except for the signature, the cryptographic signature, which cannot be spoofed. And of course Facebook wasn't happy that somebody had a clone, a working fake of their certificate, which was only trusted because the person who had it had installed the Kazakhstani root certificate.

Anyway, so what we do know is that, amid all this confusion, Google, Mozilla, and Apple didn't find any of this confusion humorous, and they took immediate action. Mozilla's blog

posting about this last Wednesday contains some new information, so I'm going to share it. And Google blogged, and Apple blogged. Everybody was like, you know, these made - the major browser vendors did not find any of this amusing.

Mozilla said: "In July, a Firefox user informed Mozilla of a security issue impacting Firefox users in Kazakhstan. They stated that Internet Service Providers in Kazakhstan had begun telling their customers that they must install a government-issued root certificate on their devices. What the ISPs didn't tell their customers was that the certificate" - this is Mozilla talking - "was that the certificate was being used to intercept network communications. Other users and researchers confirmed these claims and listed three dozen popular social media and communications sites that were affected."

And I'll break here just to note that this perhaps suggests that the interception was selective, which would make sense and would also explain the inconsistency of the connectivity reports. So my point being that it may well have been that ISPs were filtering specific domains like Facebook, but not the Internet as a whole. So that would explain why people who were trying to go to Facebook or Instagram or whatever, three dozen popular social media and communications sites, were having a problem; whereas others who weren't doing that said, no, everything's working just fine. Yeah, until you try to go to one of the proscribed sites.

So anyway, Mozilla continues: "The security and privacy of HTTPS-encrypted communications in Firefox and other browsers relies on trusted Certificate Authorities to issue website certificates only to someone that controls the domain or website." Which of course is what you have to prove in order to qualify for a certificate is you prove domain ownership. "For example," they say, "you and I cannot obtain a trusted certificate for www.facebook.com because Mozilla has strict policies for all CAs trusted by Firefox which only allow an authorized person to get a certificate for that domain.

"However, when a user in Kazakhstan installs the root certificate provided by their ISP, they are choosing to trust a CA that doesn't have to follow any rules and can issue a certificate for any website to anyone. This enables the interception and decryption of network communications between Firefox and the website, sometimes referred to as a Monster-in-the-Middle (MITM) attack. We believe," they write, "this act undermines the security of our users and the web, and it directly contradicts Principle 4 of the Mozilla Manifesto that states: 'Individuals' security and privacy on the Internet are fundamental and must not be treated as optional.'" And note that privacy is in there, too. So tracking.

They said: "To protect our users, Firefox, together with Chrome" - and now we know also with Safari - "will block the use of the Kazakhstan root CA certificate. This means it will not be trusted by Firefox, even if the user has installed it. We believe this is the appropriate response because users in Kazakhstan are not being given a meaningful choice over whether to install the certificate, and because this attack undermines the integrity of a critical network security mechanism," the whole PK, public key certificate system.

They said: "When attempting to access a website that responds with this certificate, Firefox users will see an error message stating that the certificate should not be trusted. We encourage users in Kazakhstan affected by this change to research the use of virtual private network software" - in other words, they're saying go around your ISP - "or the Tor Browser to access the Web. We also strongly encourage anyone who followed the steps to install the Kazakhstan government root certificate to remove it from your devices and to immediately change your passwords, using a strong, unique password for each of your online accounts." And so that concludes Mozilla's posting.

And that bit about changing passwords is a very good point, too, when you think about it, because any login username and password credentials which were used on those sites

would have been compromised by the ISP's ability to intercept everything. So even the login with username and password, that's all broken if they're able to intercept TLS. In any event, Google and Apple both followed suit with Mozilla, with Google adding the errant cert's fingerprint to Chrome's CRLSET. So the CRLSET is good for something - if not blocking revoked certificates, at least explicitly blocking certificates that you absolutely don't want to accept. Also, Google has added that to the Chromium source code base so that blocking of that root will eventually filter out to all Chromium-based browsers. And presumably that includes Edge, as well.

So this has been an interesting little dance that we've just watched happen, and one has to imagine that other countries, and hopefully ISPs in the U.S., are recognizing that they don't have the ability. Yes, technically you could ask users to do this. But the whole point is to intercept communications by browsers. The only recourse would be to go the next step, and we talked about this at the time, which would be for Kazakhstan to take an open source browser and then make it their own, essentially. And nobody wants that to happen because it would be very unlikely that they would keep up with current security fixes on an ongoing basis. And, boy, forcing everyone in the country to use the Kazakhstan browser if they want access to 30 particular social and communications websites seems like a heavy lift. So again, we're at this juncture where governments are saying it's not okay for them not be able to monitor the communications taking place within their borders, and this is creating a lot of tension. We don't have a resolution to that yet.

So the headline on The Hacker News site was "Hackers Planted Backdoor in Webmin, Popular Utility for Linux/Unix Servers." And that did happen. "With over three million downloads per year, Webmin," they wrote, "is one of the world's most popular open source web-based systems for managing Unix-based systems, primarily servers running Linux, FreeBSD, or OpenBSD." And I have to say I was unfamiliar with it. I'm always happy at a command line. I'm very comfortable at the command line of a FreeBSD or Linux server.

But so I took a look at the screenshot. It's really pretty. I mean, I found myself thinking, huh, that looks like a nice way to admin a remote server. On the other hand, if you thought that, and you had downloaded this particularly vulnerable version of Webmin, you might change your mind. Although, I have to say, nothing of those sorts of services is ever exposed to the public. I can't bring myself to do that. So I'm a big believer in filtering traffic so that anything that you can't really vouch for is just not available to the open public. And so if I were to ever use that, I would make sure that that web portal was only available, over VPN, to systems that I have secured through other means. But it offers a very neat-looking UI for managing users and groups, databases, the BIND DNS server, Apache, Postfix, Sendmail, Qmail, your backups and firewalls, monitoring alerts and logs and so forth. So it looks like a nice piece of work.

Anyway, what happened was, it was with some concern that, following the surprise and irresponsible public disclosure of a critical zero-day vulnerability in this Webmin on August 10th is when this irresponsible public disclosure occurred during a presentation at DEF CON, that the project's maintainers revealed last Tuesday that the flaw, which was until then believed to be a bug, was not actually the result of a coding mistake made by the programmers. Instead, it was the result of a secretly planted by an unknown - it was the result of software, malware, that was secretly planted by an unknown hacker who had successfully managed to inject a backdoor at some point into the project's build infrastructure which had persisted through a number of releases. It appeared first in v.1.882 and was there through 1.921, which meant that it had been present and hidden for over a year.

So this all began when a Turkish researcher presented a zero-day remote code execution vulnerability in Webmin at DEF CON on August 10th. He'd given the Webmin project no

advance notice of either his discovery or his plan to disclose his finding. Joe Cooper, who's one of the Webmin project's developers, said: "We received no notification of it, which is unusual and unethical on the part of the researcher who discovered it. But in such cases," he said, "there's nothing we can do but fix it ASAP." And they did.

Besides revealing the flaw to the public, the Turkish researcher also released a Metasploit module for the vulnerability to automate its exploitation using the Metasploit framework. The vulnerability, which is now tracked as 2019-15107, was introduced in a security feature that was designed to let a Webmin admin force a password expiration policy for other user accounts. Which, and as we've talked about many times on this podcast, I've never understood the logic behind forcing people to change their password. It just creates a mess, certainly in their little password notepads, where they're recording all their passwords.

According to this Turkish researcher, the security flaw resides in the password reset page. And get this. It allows a remote, unauthenticated attacker to execute arbitrary commands with root privilege on vulnerable servers just by adding a pipe command - that's the vertical bar - into the old password field through POST requests. So in the first place, I mean, nothing about that feels like a mistake; right? So the backdoor is, if you send a POST query to the server, which is what the browser would normally submit, if you were using a web admin, and you were going to be updating your password, you would fill in your old password and then your new password maybe even twice.

Well, that would be submitted with a POST. So what was discovered was that, if you put the pipe vertical bar into the old password field, you could then pipe in commands through that POST submission, that POST query, that would be executed with root privilege. So this is like, okay, that's not going to happen by mistake. That wasn't a coding error. It has all the feeling of something that was deliberately snuck in.

So in a blog post published yesterday, Cooper said that the team is still investigating how and when the backdoor was introduced. But he confirmed that the official Webmin downloads were replaced by the backdoored packages only on the project's SourceForge repository, though that actually is the primary download point, and not on Webmin's GitHub repositories. SourceForge was like the official distribution for this particular package. Joe also stressed that the affected password expiration feature is disabled by default for Webmin accounts, which means that most versions will not be vulnerable in their default configuration, and that the flaw would only affect Webmin accounts, or Webmin admins who had manually enabled this feature.

He said: "To exploit the malicious code, your Webmin installation must have," and then there's the nested menu tree: "Webmin > Webmin Configuration > Authentication > Password expiration policy set to prompt users with expired passwords to enter a new one." He said: "This option is not set by default; but, if it is set, it allows remote code execution." Which is not a good thing.

However, another security researcher on Twitter later revealed that Webmin v1.890 is affected in its default configuration, as hackers appear to have modified the source code to enable the password expiration feature by default for all Webmin users. So the plot thickens. An unknown attacker made a subtle change to a Webmin script called password_change.cgi. This change gave attackers the ability to send a command through a special URL that an infected Webmin server would then execute with root privileges.

Anyway, it turns out in version 1.890, which had more than 421,000 downloads between June of 2018 and last weekend, when all of this got fixed, the backdoor was turned on by default. In 1.90 through 1.92, which collectively had more than 942,000 downloads, the backdoor was active only when changed and deliberately enabled.

So a Shodan search showed that Webmin has nearly a quarter of a million Internet-exposed instances, 218,000 Internet-exposed instances. So it is popular. So a quarter million exposed Webmin portals, of which more than 13,000 are running the default vulnerable Webmin version 1.890. And of course since this one is a bit old, this suggests also that those Webmin instances are not being kept current since a whole bunch of versions since then don't have it on by default. So the threat intelligence firm Bad Packets tweeted that there are several actors now actively exploiting the Webmin vulnerability, and that that exploitation began the day after its zero-day disclosure at DEF CON.

So, I mean, this demonstrates a number of things. First of all, these systems are not updating themselves. It really, really should be, especially for something like this. It also demonstrates the danger, I mean, that in fact it's hard to argue that this was not unethical for this disclosure to be made without notifying the Webmin admins because it did take them time to respond, and attacks began the day after, the next day after this disclosure. And of course this was all aided by the fact that the guy created a Metasploit exploit, which took all the mystery out of how to do this, although I just explained it. I mean, it's trivial to exploit. And unfortunately, Shodan finds these for people so they don't have to even scan for them. 13,000 instances now of as a consequence of this web portal server administration tool having this backdoor installed which was present for quite a while and is still now present. They're all subject to this remote vulnerability.

The Webmin developers promptly removed the malicious backdoor, of course, and released clean versions. But again, it's not looking like these systems that are going to be attacked are going to get any benefit from that because they're already using one, which is very much older. So needless to say, if any of our listeners are Webmin administrators, that is, you are using Webmin on any of your Linux/FreeBSD/OpenBSD systems, absolutely, I mean, I guess check them. Of course, this is a big problem; right? If something gets in, if you happen to be using the vulnerable one, or you have turned that option on such that you're vulnerable, how do you ever trust that server again? So that's the first of several stories about the problems we're seeing with attacks on the open source supply chain.

Leo: So.

Steve: So RubyGems is once again in trouble.

Leo: Uh-oh.

Steve: I know.

Leo: This just happened.

Steve: Yup, just happened. A second backdoor came to light on Monday in 11 libraries available in the RubyGems repository. According to an analysis by the developer, Jan Dintel, the backdoor allowed attackers to use preset credentials. So they had, like, they built in, right, username and password, preset credentials to remotely execute commands of their choice on infected servers. The malware included a variety of other capabilities including code that uploaded environment variables which are often the source for static credentials used to access databases, service providers, and other sensitive resources. The exfiltrated material was sent to a server. And I had to look

up .ua. That's the Ukraine. So located in Ukraine. RubyGems officials also found the malicious code included a miner for cryptocurrencies because, you know, why not while you're at it. Download counts showed that the backdoor libraries had been downloaded nearly 3,600 times. Rest-client versions 1.6.10...

Leo: Wait, wait, what was that?

Steve: That was email coming in - through 1.6.13, accounting for more than 1,200 of those downloads, they were backdoored by someone who compromised an old developer account protected by a previously cracked password. So it's unclear how the other RubyGems libraries were infected. That was one of 11 that were. So standing back from that a bit, both the Webmin and RubyGems libraries turn out to - because we've touched on these over time. They're only the latest in a series of supply chain attacks which have impacted the software of the open source community.

And most people don't think twice about installing software or updates from the official site of a known developer. I know I do it all the time. But as software and websites have continued to tighten down their security and become more difficult to exploit, which is what we're seeing over time, the black hats have increasingly turned to exploiting this trust that we have in well-meaning developers who are doing their development in plain sight to poison the code at its source. This increase in focus first came, I think, to light last October, when two unrelated supply chain attacks against two open source projects were discovered. The first was the Vespa control panel, the VespaCP control panel interface; and the other was the "colourama" package that was slipped into the Python repository. And I remember we talked about both at the time.

Then a month after that, malicious code designed to steal funds from bitcoin wallets found its way into event-stream, which is a code library with two million downloads that's used by Fortune 500 companies and small startups. Officials from NPM, the Node JS Package Manager, said that they were hosting backdoored software, saying that the malicious code was designed to target people using a bitcoin wallet which was developed by Copay, which is one of the companies that incorporated event-stream into its app. So it would make sense that, if you could compromise this event-stream code and get that into this NPM package, then that would achieve your ends from a malware standpoint.

Then, finally, a few months ago, in March, researchers found, as we talked about it at the time, that that first RubyGems library called bootstrap-sass was also backdoored. And early last month we talked about the RubyGems library called strong_password being backdoored. And then, like the attack discovered this week which infected the 11 RubyGems projects, that bootstrap-sass and the strong_password backdoors used a browser cookie function to give attackers the ability to execute code on infected servers. That strong_password backdoor also interacted with a server in Ukraine, smiley.zzz.com.ua. And of course the previous Ukraine server was mironanoru.zzz.com.ua, which obviously bears a strong resemblance to smiley.zzz.com.ua, making it feel like this is all a single actor.

Anyway, all of this raises the specter of what hasn't yet been discovered; right? I mean, so there's an incredible number of open source packages that represent a tremendously beneficial resource to the computing industry. All we can talk about are the problems that have been found. We don't know what may be present and not found.

So HD Moore, whom we've spoken of many times, he's a network security expert, codes on open source projects. He's the hacker who originated and developed the Metasploit framework. He said: "The recent discoveries make it clear that these issues are becoming more frequent and that the security ecosystem around package publication and

management isn't improving fast enough." In other words, that security needs to get tightened up. He said: "The scary part is that each of these instances likely resulted in even more developer accounts being compromised through captured passwords, authorization tokens, API keys, and SSH keys."

He said: "The attackers likely have sufficient credentials at hand to do this again, and repeatedly, until all credentials are reset, and appropriate multifactor authentication and signing is put in place." So he said the impact of open source supply chain infections is often hard to gauge also because backdoored applications might be included as an upstream dependency by another package. So it's not just direct downloads, but it's indirect downloads as a consequence of package dependency managers pulling things down that are needed by other packages.

Anyway, finally, he said: "The way that dependency management tools push for the latest packages by default makes a successful attack in the case of a backdoored dependency even more likely." So anyway, it's something that we need to keep in mind, which is that we're not sure who's watching. You know, we've talked about, for example, a concern that maybe there was subtle influence by the NSA on the entropy being produced by pseudorandom number generators, like maybe there was a way for the NSA to have skewed the output in a way that benefited U.S. law enforcement at the expense of the absolute entropy that was presumed to be generated by these things. I mean, so there's that. That happened in plain sight.

But it does look like the security of these open source projects is something that really needs to be paid attention to because this to some degree represents some low-hanging fruit. While other sources of compromise are getting more secure, the bad guys are going to look for the easiest thing to infect. And if these resources are not kept secure, then they could be subject to attack.

Chrome is going to be adding data breach notification. We talked about Mozilla doing this. The genesis of Chrome's decision was Mozilla's partnering with Troy Hunt's HaveIBeenPwned service to create its Firefox, they call it the Monitor, data breach notification service, which checked for the presence of users' credentials, their passwords or emails, to see whether they are among those leaked in past data breaches. I'm sure, although that's not what the Firefox Monitor Data Breach Notification Service does, I'm sure I've had the experience, I've talked about it, in fact, on the podcast, of going to a site and receiving one of those little dropdown notifications from Firefox saying this site has suffered a breach in the past.

So it turns out that's where Chrome is going to be going. What Google first created was what they called their "password checkup," which was a Chrome extension to perform a function similar to Firefox's Monitor Data Breach Notification. In this case Google has been maintaining a list of credentials, much as Troy has been. Apparently they've got four billion leaked and disclosed credentials, which the password checkup Chrome extension was able to check against. But using the analytics which was built into this password checkup extension, Google was able to get a sense for what people were doing. They conducted a study using the analytics which showed that 1.5% of all logins have been compromised in data breaches, 1.5%. The study also showed that 26% of users, so just over one out of every four, who were shown a data breach notification, changed their password as a result. So it wasn't ignored. It was like, oh, crap, I'm going to change my password.

Since this study demonstrated to Google that providing notifications of compromised login credentials was beneficial to users, Google is now building the support directly into Chrome. So once this feature is in place, Chrome will begin alerting users when they are logging into sites with credentials that have been exposed by breaches, much as that notice I'm sure I received from Firefox was. So not that your credentials specifically have

been exposed, or not that there's been a collision of your password and email with, for example, the HaveIBeenPwned database that Troy maintains, but just this website has been exposed to a breach.

And I remember when I talked about this now with Firefox. I liked that because that's not something that any website has control over, if Firefox and Chrome both note that for future visitors. And that's a blemish that no website wants to carry. And I think I remember noting that maybe it ought to expire after some number of years. It's not something that ought to be, like, forever and ever. But still it would provide additional incentive for a website to work for that not to be true of them.

So there's also a Check Your Passwords button. In the reporting of this and in Google's own disclosure, they showed some screenshots, and there was a Check Your Passwords button. But nobody that I could find was sure what it was that that would do. Maybe that will - it might be that the dropdown says this website has suffered a breach in the past, and then you can have it explicitly check your password to see whether there's any collision that Google is able to find with those that have been disclosed in the past.

I wanted to briefly mention that I was surprised that so much was made of basically a code regression mistake that Apple made in their recently released v10.4...

Leo: 12.4.

Steve: Oh, right, I wrote that [crosstalk] 12.

Leo: Yeah, I saw you had 10; and I thought, it's 12.

Steve: Yeah, yeah, yeah, it is 12.4 because of course we're all waiting for 13.

Leo: Right.

Steve: Lucky 13 in a couple weeks.

Leo: Right.

Steve: So, yes, 12.4, thank you. What they did was they had a bit of a code regression. There was a jailbreak which had been found in an earlier version of iOS and fixed. And a hacker noted that it had reappeared. Oh, it was fixed in iOS 12.3, and then reappeared, to everyone's surprise, in 12.4. Apple quickly said whoops and fixed it in 12.4.1, which we all now are using, those of us who've updated our iOS devices. That got fixed. But it was interesting, I mean, it was just like, all of the tech press was breathless over the idea that a jailbreak was now available. Of course, you know, I guess it's because they're so rare and would be sought-after. If they were kept secret, and if nobody told Apple, it would be handy for people who like the flexibility of jailbreaking their phones. I mean, once upon a time it was, like, common. And now it's like, no, that's not something Apple wants to allow to have happen.

And lastly, Microsoft's RDP client for Android needed and received an update. We've recently been discussing the various concerns over vulnerabilities in the Remote Desktop

Protocol, and then subsequently in the Remote Desktop Services, which is an enterprise-grade service where users can actually connect to their desktop remotely, you know, multiple people hooking to a licensed server, that's licensed for that. Whereas those of us who have standard Windows, we're able to connect one session to Windows, and it logs off the desktop session if you happen to be logged on at the same time. And then, of course, the most notable of the RDP Protocol was BlueKeep that everyone keeps expecting to go into a worm, but hasn't so far. I've seen other commentary, I mean, my theory is it's not difficult enough to be worth doing. Others have said it is not easy enough for a worm to take over. So anyway, we'll see.

But recall that some time ago, predating these, there was also a means for Remote Desktop Servers to sort of reverse compromise remote desktop clients. The trouble was the recurring problem with vulnerabilities in interpretation. And in this case RDP clients, which are displaying the desktop, they are inherently interpreting and trusting the data they receive from the RDP servers, in order to draw the desktop. They are doing lots of interpretation. We talked about it at the time, that there was a way that a client could be taken over if it connected to a malicious server. And at the time I said, you know, that seems unlike - it's not like a big worry because people are not running around connecting their remote desktop clients to random Windows servers. Normally you're just connecting to your own machine at home or to an enterprise system. But still, it should get fixed.

Anyway, it turns out that Microsoft has an RDP client for Android, and that there were also similar exploits possible against it. So Microsoft has patched that. So although the threat is minimal, I just wanted to put it on our listeners' radar that it was worth, if you were a person who used an Android device for RDPing to some Windows system, again, if you're RDPing to your own system, it's unlikely that your own computer is going to attack your Android client. But I'm not sure what the update channel is for Android software from Microsoft. So if that's something that you use, why not update?

I did mention, and I'll say more, of SQRL. I just finished the third of the four SQRL documents. That was SQRL's cryptography documentation. And I am now on, well, I will shortly be on SQRL's what I call "On the Wire," which is basically, it's the fourth of the four documents, and this is all of the - basically everything that's left, since it's the last document, which is the way you format the keys and the signing and coding and so forth into ASCII for transmission on the wire and what the various bits mean in the protocol and so forth. So I will be starting that.

The reason I haven't quite is that after Thursday's Orange County OWASP SQRL presentation, I decided to spend a little bit of time on my PowerPoint presentation. First I should say that Thursday's OWASP presentation was truly terrific. I think it was 117-some people had signed up, which was between two and three times the group's normal size based on previous numbers that I saw in Meetup. And I asked for a show of hands of how many people there were Security Now! podcast listeners. And I expected, you know, maybe half? Everybody's hand went up. Which was surprising. And then I said, okay, wait a minute. Who here is not a Security Now! podcast listener? And maybe two or three hands were short of sheepishly raised. And then I said to them, I said, well, maybe there's a clue there.

Anyway, it was really a fun evening. And I got to meet a lot of the podcast listeners for the first time. But my concept, and I think I mentioned it earlier, was I wanted to see if I could just have all the diagrams on the screen and randomly select them, rather than following a bullet-pointed presentation. And I found that that didn't really work, either. But I got a sort of an updated inspiration for the SQRL presentations. And so I'm spending a little time working on that. And then as soon as that's finished, I'm back to the fourth of the SQRL documents, which will be "SQRL on the Wire."

Also the two upcoming presentations which are taking me out of the country, in response to a tweet from OWASP in, and I can't pronounce this the way they do, so I'll just say Gothenburg, which is in Sweden. So OWASP...

Leo: You'll learn it before you get back. You'll know it.

Steve: Probably will. Yeah, GBG said...

Leo: I'll buy you some Slivovitz if you can do it right. It helps. The more Slivovitz you have, the more you sound right.

Steve: So they tweeted two days ago: "Yihaa," I guess, "Yihaa, what an amazing response we got for the event with Steve Gibson @SGgrc!!" They wrote: "200 tickets, and we're fully booked."

Leo: Oh, great.

Steve: And they actually said 218 in a private email, so he says: "As usual, you can sign up on the waiting list through Eventbrite." And then after that I shot a note to Mick Ryan, who's the organizer of the Dublin, Ireland event. And he wrote back: "It's full, Steve! 300 people."

Leo: What? What?

Steve: He said: "The room fits 200 people seated, with lots of standing room also."

Leo: You've got fans all over the world, my friend.

Steve: I guess so.

Leo: That's awesome. That's awesome.

Steve: So anyway, but so I just thought I would give a hint to our listeners. Unfortunately, maybe they're all going to be our listeners. They're 50% overbooked. So it might be a good idea to show up early if you want to get a seat. And I was mentioning this to Lorrie last night, and she said, "Well, yeah, but they could sit in the aisles and on the floor and so forth." So it's like, yeah, well, okay, that's true. But the full - this presentation ends up being like one of our podcasts. It's about two hours of just nonstop information transfer. So you really would rather not stand for that, if you didn't have to. So I imagine that people will be sitting. And, boy, I mean, the time just flies by. So we're going to have a lot of fun a month from now in Dublin and then wherever that is in Sweden, Gothenburg, something. And again, GRC.com/calendar will inform you of those events.

Leo: So great. Wow. I'm impressed.

Steve: Yeah, yeah. And I'm still on the trek with the file sync journey, which you and I, Leo, will be doing a podcast devoted to that here in a couple weeks. I'm looking around at different things. I'm keeping up with Twitter and submissions. I'm beginning to get a sense for what I like. So we will do a complete rundown of options and sort of clever tricks that I've come up with on the podcast a few weeks from now. So it'll be good.

Let's take our last break and then talk about the ratification of ads on the Internet. I mean, they're going to stop being, I think I see this, it's going to make sense, stop being sort of an ad hoc, shoehorned-in thing that no one likes and become probably well-behaving citizens, thanks to some work that Google is doing.

Leo: You're such an optimist. But we need it. It's the ecosystem. We want great web content. It's got to be paid for somehow.

Steve: Yup. Exactly.

Leo: There's got to be a way to do it. All right. Ad privacy with Mr. Gibson.

Steve: So Apple, who of course doesn't depend upon revenue from web ads, has been moving forward on the anti-tracking front for some time.

Leo: Same thing with Firefox; right? I mean, it's not their business; right? So they don't need to worry about it.

Steve: Yes, yes, yes. Exactly. So on June 5th of 2017, so a little over two years ago, we got from Apple Intelligent Tracking Prevention. And we talked about it at the time. We liked it. The idea was that they were going to be expiring cookies after some length of time so they didn't persist forever. So it was sort of a nice compromise. You would sort of, like, have the right to be forgotten if there wasn't something keeping the cookie persistent.

Then on March 14th of 2018, we got the update, Intelligent Tracking Prevention 1.1. Then on June 4, a few months after that, 2.0. Then February 21st of this year, 2019, we went to Intelligent Tracking Prevention 2.1. And then in April of this year, ITP 2.2. And then, most recently, toward the end of May, on May 22nd, Apple posted Privacy Preserving Ad Click Attribution for the Web, which is Apple's attempt to essentially solve this problem. My point is I'm not going to go into this in great detail because I don't think it makes sense until we get some consensus, and we're still pre-consensus. But Apple has a proposed solution - Google refers to it in their own brainstorming for a solution - of preserving ad click attribution while also preserving privacy.

So my point is that there is a, you know, we've seen Apple continuing to refine tracking prevention. Apple has put their stake down as "We're the security and privacy protecting company," and we're seeing them move on this. Certainly, as you said, Leo, Google's revenue is from advertising. We know that, and Google tells us. That's why all this stuff can be free is advertising revenue. So what appears to be happening is that we're approaching an inflection point for the industry, the fact that it's sort of a convergence.

As I said at the beginning, advertising and tracking has just been ad hoc to this point. What annoyed me was that it was abusing a system that was never intended for third-party tracking. Cookies were meant to create a stateful relationship with the site you were visiting. But then things like the Like buttons and where there is a little presence of Facebook Like all over the place, well, your browser, that's a link to Facebook which your browser retrieves. And in doing so, it sends Facebook the cookie you have with them in that query. So they know where you are.

And of course the same thing is true for ads. Apparently, I mean, it's not just theoretical that the ad companies desperately want to track us because when Chrome does something like blocking fingerprinting by - like we were just talking about the Incognito mode. Well, in this case that wasn't advertising tracking, it was The New York Times and the Washington Post saying, wait a minute, we're not going to show you anything if you're in Incognito mode because we want to be able to meter our free sample of our paid content, and you need to drop Incognito mode. So Chrome said, if users are incognito, then they need to be incognito.

Last Thursday over on the Google side, Justin Schuh, I guess that's how you pronounce it, S-C-H-U-H? Schuh, maybe? He's the director of Chrome Engineering. He posted two short overviews about what Google is calling their Privacy Sandbox initiative. And they're both short. I want to share them, and then we'll dig more into the details.

The first blog posting was titled "Building a More Private Web." He said - and yes, I understand we have to consider the source. He said: "Privacy is paramount to us, in everything we do."

Leo: [Buzzer sound]

Steve: Okay. Right. He said: "So today" - yeah, we who read all of your email on Gmail.

Leo: That's a bad lead. That's really a bad lead.

Steve: Yeah.

Leo: Okay, go ahead.

Steve: Okay. Again, consider the source.

Leo: I don't believe anything he says after this. Go ahead.

Steve: Have your buzzer ready, Leo. He said: "So today we are announcing a new initiative to develop a set of open standards to fundamentally enhance privacy on the web. We're calling this a Privacy Sandbox." He says: "Technology that publishers and advertisers use to make advertising even more relevant to people is now being used far beyond its original design intent, to a point where some data practices don't match up to user expectations of privacy."

Leo: Okay. That's true.

Steve: That's true. "Recently, some other browsers have attempted to address this problem," which is like saying we're not happy with what Apple is doing because they're really saying they're going to...

Leo: And Firefox, yeah. Yeah, Firefox is really doing a great job on this.

Steve: Yup. He says: "But without an agreed-upon set of standards, attempts to improve user privacy are having unintended consequences. First, large-scale blocking of cookies undermine people's privacy by encouraging opaque techniques such as fingerprinting. With fingerprinting..."

Leo: Okay. Put an asterisk there because I disagree with that, but go ahead.

Steve: Okay. "With fingerprinting, developers have found ways to use tiny bits of information that vary between users." And of course we know, and maybe this is the nature of your asterisk, is that fingerprinting is not unique to an individual. You know, when you go to Panopticlick, for example, to look at it, it shows you how many bits of entropy they found, and how many other users had the same fingerprint as you. Meaning, you know, you're not a unique little flower. You're one in a cohort.

Leo: My asterisk is really, yes, that's true, that's a bad thing. But that doesn't mean, well, we've got to keep cookies. They're both bad things. But go ahead.

Steve: Yes, yes, thank you.

Leo: That doesn't justify cookies, let's put it that way.

Steve: Correct. So they say little "...bits of information that vary between users such as what device they have or what fonts they have installed to generate a unique identifier which can then be used to match a user across websites. Unlike cookies, users cannot clear their fingerprint, and therefore cannot control how their information is collected. We think this subverts user choice and is wrong." Okay.

Leo: And I should point out that both WebKit (Apple) and Firefox have done some things to block fingerprinting.

Steve: To thwart fingerprinting, yes.

Leo: Fingerprinting is thwartable.

Steve: Yes. He says: "Second, blocking cookies without another way to deliver relevant ads significantly reduces publishers' primary means of funding..."

Leo: On that I'll agree, yeah.

Steve: Yes. And I was really affected, Leo, by - I was listening before our podcast once sometime ago to your MacBreak Weekly, where you were able to talk to, I think it was Rene, about the degree to...

Leo: Oh, it's killing iMore, yeah.

Steve: Yes, the degree to which they're dependent upon advertising for revenue. And it was the more effective blocking of cookies and what - or, well, blocking of ads actually.

Leo: And yet...

Steve: It was the ad blockers is what it was.

Leo: And yet I would point out, we can, without any tracking at all, make a good living here at TWiT. Advertisers continue to buy ads. NBC, CBS, and ABC continue to sell ads without any tracking. So there's a lot of media that works quite well without tracking. The problem is advertisers want it. They wish they could get it.

Steve: Right, exactly.

Leo: Despite the evidence it doesn't help, by the way. There's very little evidence that tracking does much good.

Steve: Well, now, and that would be worth digging into because what he's about to say is the first number I've seen. And I've always wondered whether this was just completely bogus or not. And we still don't know. But he said: "Many publishers have been able to continue to invest in freely accessible content because they can be confident that their advertising will fund their costs. If this funding is cut, we are concerned" - yeah, "we," right, Google - "are concerned that we will see much less accessible content for everyone. Recent studies" - and here it is. "Recent studies have shown that when advertising is made less relevant by removing cookies, funding for publishers falls by 52% on average."

Leo: Okay. That is a study from Google.

Steve: Uh-huh.

Leo: Based on analysis of a random selection fraction of traffic on each of the 500 largest Google ad manager publishers over three months. That's not study after study. It's one paragraph in a Google blog post from Google ads.

Steve: Okay.

Leo: So that's also a little disingenuous. There isn't study after study. This is a Google kind of not such a great study.

Steve: Right. So when they say "recent studies" have shown, it would be a study we conducted showed...

Leo: Yeah.

Steve: Okay.

Leo: Which is a randomly selected fraction of traffic on each of the 500 largest Google ad manager publishers.

Steve: Right.

Leo: And that's where the 52% comes from is that study.

Steve: Okay. So he said: "We are doing something different. We want to find a solution that both really protects user privacy and also helps content remain freely accessible on the web." He said: "At I/O, we announced a plan to improve the classification of cookies, giving clarity and visibility to cookie settings, as well as plans to more aggressively block fingerprinting. We are making progress on this." But where we're headed is different. He says: "And today we are providing more details on our plans to restrict fingerprinting. Collectively, we believe all these changes will improve transparency, choice, and control.

"But we can go further. Starting with today's announcements, we will work with the web community to develop new standards that advance privacy, while continuing to support free access to content. Over the last couple of weeks we've started sharing our preliminary ideas for a Privacy Sandbox, a secure environment for personalization that also protects user privacy. Some ideas include new approaches to ensure that ads continue to be relevant for users." And I've never found them to be relevant for me, but what the heck.

Leo: That's to me the real issue; right?

Steve: Yeah, I mean, I just...

Leo: More relevant ads? No.

Steve: No. Okay. "But user data shared with websites and advertisers would be minimized by anonymously aggregating user information, and keeping much more user information on-device only." That is on-hyphen-device only.

Leo: Sounds good, yeah.

Steve: It does. It sounds good. "Our goal is to create a set of standards that is more consistent with users' expectations of privacy. We're following the web standards process and seeking industry feedback on our initial ideas for the Privacy Sandbox. While Chrome can take action quickly in some areas - for instance, restrictions on fingerprinting - developing web standards..."

Leo: Wait a minute. I thought you couldn't do that.

Steve: Right.

Leo: It's kind of contradicting what he said earlier. Oh, no cookies, but you get more fingerprinting. Oh, you can block that. Oh, okay.

Steve: Well, what he said was, if you restrict cookies, then you're driving advertisers to find some other means, and he used fingerprinting as an example of what they would resort to if you blocked cookies. He said: "Developing web standards is a complex process, and we know from experience that ecosystem changes of this scope take time. They require significant thought, debate, and input from many stakeholders." And of course we want everybody involved because there are a lot of privacy advocates, you know, the EFF, they'll jump into this and get their two cents' worth in. So that's good.

He says: "To move things forward as quickly as possible, we have documented the specific problems we are trying to solve together, and we are sharing a series of explainers with the web community. We've also summarized these ideas today on the Chromium blog." Which is where we're going to go next. "We look forward to getting feedback on this approach from the web platform community, including other browsers, publishers, and their advertising partners. Thank you in advance," he finishes, "for your help and input on this process. We believe that we must solve these problems" - and I want, I mean, this is all sounding good, depending upon what it is - "together to ensure that the incredible benefits of an open, accessible web continue into the next generation of the Internet."

So now we switch to his posting on the Chromium blog, which gets into more nitty-gritty, titled "Potential Uses for the Privacy Sandbox." And he says: "Today on The Keyword" - which is what I just read - "we outlined our vision for an initiative aimed at evolving the web with architecture that advances privacy while continuing to support a free and open ecosystem. In order to work toward that vision, we have begun publishing a series of explainers that are intended to be shared and iterated on across the community. Below, we've summarized each of these early proposals, which we are collectively referring to as the Privacy Sandbox."

So first is user information. He says: "First, let's identify how user information is currently used in the advertising ecosystem so that we can explore the development of the Privacy Sandbox's privacy-preserving APIs." And here I'm kind of glad that they've got an advertiser, that they own an advertiser, because they also own a browser. And it would be nice, I mean, so they understand what the advertiser's view is.

Leo: No coincidence, of course.

Steve: I know, of course. So then ad selection. For ad selection they said: "One of the most challenging questions is what your browser could do to allow a publisher to pick relevant content or show a relevant ad to you" - and I still say good luck with that - "while sharing as little information about your browsing history as possible. We're exploring how to deliver ads to large groups of similar people without letting individually identifying data ever leave your browser, building on the differential privacy techniques we've been using in Chrome for nearly five years to collect anonymous telemetry information." There are some cool crypto hashing things that can be done.

Anyway, he said: "New technologies like Federated Learning show that it's possible for your browser to avoid revealing that you are a member of a group that likes, for example, Beyonc and sweater vests until it can be sure that group contains thousands of other such people." In which case...

Leo: Let me interrupt before you move on.

Steve: Yeah, good.

Leo: Because some people have talked about this. And the issue is not merely that I don't want you to know who I am, but that you don't want to assign me to a group. Let's say that you've collected over time information that leads me to think that I am bipolar, and that I'm in a manic phase. And, oh, good, I've got 50,000 other people bipolar in a manic phase. I know they're going to be particularly susceptible to a certain kind of advertising. Admittedly, you don't know who is, what their name is, but you still can target them in a vulnerable situation. You could still use it for a targeted ad in political situations. Are you a gun owner?

I mean, in other words, this is a little disingenuous in the sense that, yes, you're protecting an individual's personal information, but it's not stopping the kind of targeting that could be and has been used in the past, thanks to Facebook and Cambridge Analytica, in a very detrimental way. So I just - I want to point out it doesn't eliminate all hazards. There's still this hazard.

Steve: I agree. And in fact, to broaden that a little bit, what you're saying is that there is a fundamental tension - I mean, fundamental, not about technology - between what advertisers want and what users want.

Leo: In a nutshell. That's exactly it. That's exactly it. And you can come up - because we're technologists, we love technological solutions. The idea of differential privacy is fascinating. But you nailed it. That's exactly what the tension is. Advertisers want this information.

Steve: And that's not going to go - yes. And there was, many, many years ago, what was the show with Paul Reiser? He was married to a woman...

Leo: I know what you mean, yeah. "How I Met Your Mother" or something. Anyway, yeah, yeah.

Steve: Yeah. It was just - it was about their relationship. And he was there, and TiVo was sort of a new thing. And he was watching sports.

Leo: "Mad About You." "Mad About You."

Steve: "Mad About You," yes. "Mad About You." He was watching football or something. And his wife walked in and says, "What are you doing?" And he says, "I'm watching sports." And she says, "Why?" And he said, "TiVo thinks I'm gay."

Leo: There was a whole article about that, actually. That really happened to somebody, yes. Too many thumbs up to the wrong shows.

Steve: And so TiVo was profiling and doing recommended shows, and it had decided - anyway, I got a kick out of that.

Leo: The Netflix recommendations, the Amazon recommendations, the TiVo, it just doesn't work. But even - but what if it did? That's almost even more terrifying.

Steve: Well, and exactly to our point is that you brought up the point, which I think is accurate, I mean, it obviously is, that advertisers want something we don't want them to have.

Leo: Right.

Steve: And isn't it our right to deny them that?

Leo: And I would say of course they want it. They ask us for this all the time. Can we put ad trackers in, you know. And there's a huge movement afoot in the podcast industry to add ad tech to podcasting, which I'm virulently against because you see what ad tech did to blogs and online periodicals. I don't want that to happen to podcasting. But there is a strong movement because it's demand from advertisers.

Steve: And so you ran into this when you were in Florida two weeks ago?

Leo: Absolutely. In fact, there's new ratings techniques. There's all sorts of ways that - get ready. Somebody came up to me, said, "Oh, there's this new technology. We put a little subsonic sound in the podcast, and then we can match it." It's bad. It's really bad. And I don't - I know our audience doesn't want that. And I don't want to participate in it. And, oh, by the way, our ads work without that. Of course advertisers are going to want everything they can get. But we don't have to give it to them.

Steve: Yeah. And it's not like ads become zero effective. We saw, I mean, if we believe - well, first of all, we don't know if we can believe the number that Google said, that is, it cuts the revenue in half.

Leo: No, that's made up. And there's a lot of evidence that it doesn't, that in fact there's academic studies that show...

Steve: I've always wondered.

Leo: ...a small, like, well, I'm looking at an article from Princeton's Center for Information Technology Policy about people. I don't if you've read this: "Deconstructing Google's Excuses on Tracking Protection."

Steve: No, I've not read it.

Leo: Highly recommend it. And one of the things they say is one of the academic studies that's out there says that tracking helps by about 4%, which is, I think, small enough that we can, without putting people out of business, say, yeah, let's not do that.

Steve: Yeah.

Leo: New York Times International Edition recently switched from tracking-based behavioral ads to contextual and geographic ads, and it did not experience any decrease in advertising revenue.

Steve: That's interesting. So not who you are at all, just what your IP reveals about where you are.

Leo: Yeah. We don't do that, by the way. But there are, and I think we probably will end up doing that, there's something called "direct insertion ads." Our ads are live, and they're built into the show. But most podcasters are moving to ads that are inserted, and it has to do with, if you download it from an IP address, we know where you are geographically. So somebody - see, right now when you buy an ad, you buy an ad for the entire TWiT audience. But if somebody wanted to buy an ad for the West Coast or England, they could do that with direct insertion. Only people in that geographic region would get that ad.

Steve: Well, and there's clearly some means for that happening in a TV stream because I sometimes get, like, local car dealership ads on CNN, which is...

Leo: Right. So CNN sells national ads - same thing with my radio, my national radio show - which are heard throughout CNN. And there's a gap. In that gap, local affiliates can put in their local ads. So your cable company is doing that. They're putting in local ads in the national feed. And that happens across all networks. So that's geographic. That's a good example. Because if you're a car dealer in Irvine, you don't want to have Leo see your ad. It's a waste of money. So there are things that - this is a Wall Street Journal article from August 27th, 2019. Why, that's today. "Behavioral ad targeting not paying off for publishers, study suggests. Are creepy

ads necessary to support the free web? A new academic study suggests they're not." So look at The Wall Street Journal today.

Steve: So "creepy," I mean, that reflects...

Leo: Creepy's subjective, I understand, yeah.

Steve: Well, but that's the way I've heard...

Leo: Feels that way.

Steve: ...people describe it. If they're, like, talking to somebody - in fact, I think it was Joe Scarborough who thought that, like, Siri was listening to him because he was talking about something, then when he went on the web he saw ads for it. It's like, uh, no.

Leo: Joe's a little...

Steve: I don't think so, yeah.

Leo: So this is a study from University of Minnesota, University of California Irvine, Carnegie Mellon. It suggests publishers only get about 4% more revenue for an ad impression that has a cookie enabled than for one that doesn't. The study tracked millions of ad transactions at a large U.S. media company over the course of a week. That's still a small sample. However...

Steve: Yeah, but still, it's not coming up with a lot of uncertainty. I mean, 4%...

Leo: It's pretty low.

Steve: If they're showing 4% even on a small sample, it's not actually going to be 50%.

Leo: But this confirms your hypothesis. A 2009 study showed that advertisers are willing to pay 2.68 times, 268% more, for a behaviorally targeted ad than one that wasn't. It's worth, they think - they're wrong - 268% more. That's why this is happening.

Steve: Interesting. Yeah.

Leo: Yeah. Isn't that interesting? This is today in The Wall Street Journal. I think this is - you are exactly right. This is a conversation and a battle that has just - that is important and is just beginning because nobody wants to kill the web.

Steve: No.

Leo: We want these people who - ad supported media is a good model. I like it. That's what I've lived on for 43 years. So we don't want to kill it. But we want to do it in a respectful way to our listeners or readers.

Steve: Yeah. Well, it'll be interesting to see how this evolves. It sounds as though we may get more controls. And moving this from an ad hoc - although it also sounds like advertisers may choose to go around this if we are given too much control.

Leo: Well, that's true, too, yeah.

Steve: Because of course why wouldn't they? So they have a series of explainers. I've got the links in the show notes for anyone who is interested. They have one called "conversion measurement." And they describe that as "click-through conversion measurement event-level API." It is an explainer for a potential new web platform feature which allows for measuring and reporting ad click conversions. Then there's something called the "trust token API." This document is an explainer for a potential future web platform API that allows propagating trust across sites. And that uses what's called a Privacy Pass protocol that actually is something that the Cloudflare guys came up with a couple years ago.

Leo: I trust them.

Steve: I do, too. So, I mean, I do feel like there are some interesting concepts here. They have something called the "privacy budget," and they said "combating fingerprinting with a privacy budget." And they explained it as the current state of affairs is browsers have been making changes to how cookies are treated. "Blunt approaches" - and here's that Google again, not happy. "Blunt approaches to cookie blocking have been tried. And in response, we have seen some user tracking efforts move underground, employing harder-to-detect methods that subvert cookie controls. These methods, known as 'fingerprinting,' rely on various techniques to examine what makes a given user's browser unique. Because fingerprinting is neither transparent nor under the user's control, it results in tracking that doesn't respect user choice."

So their "end state to aim for" is: "Fundamentally, we want to limit how much information about individual users is exposed to sites so that in total it is insufficient to identify and track users across the web, except for possibly as part of a large, heterogeneous group, or heterogeneous groups." And of course, Leo, we keep coming back to your point, which is people don't even want to be grouped. You know? People just want to be anonymous.

Leo: Yeah. Anonymous is anonymous, not like in a big group.

Steve: Right. And in fact they address - this next one they call "FLoC," Federated Learning of Cohorts. They said: "This is an explainer for a new way that browsers could enable interest-based advertising on the web, in which the companies who today observe the browsing behavior of individuals instead observe the behavior of a cohort or FloC of

similar people." They said: "The choice of what ads to show on a web page may typically be based on three broad categories of information: One, the site or page, irrespective of who is reading it." And that's a good point. Put this ad on web pages with motorcycles, for example.

Leo: Yeah. That's what people who buy TWiT are doing; right?

Steve: Yes, exactly. And that makes absolute sense. I mean, and why not? "Two, general information about the interests of the person who is going to see the ad, for example, show this ad to classical music lovers." Okay. And, "Three, specific previous actions the person has taken, for example, offer a discount on some shoes that you left in a shopping cart. This document," they said, "addresses category two, ads targeting based on someone's general interests."

They said: "In today's web, people's interests are typically inferred based on observing what sites or pages they have visited, which relies on tracking techniques like third-party cookies or less transparent mechanisms like device fingerprinting. It would be better for privacy if interest-based advertising could be accomplished without needing to collect a particular individual's browsing history. We plan to explore ways in which a browser can group together people with similar browsing habits, so that ad tech companies can observe the habits of large groups instead of the activity of individuals. Ad targeting could then be partly based on what group the person falls into." And I don't, you know, how do you do that? I mean, there are so many different, like, talk about a Venn diagram that would give you heartburn.

Leo: Well, that sounds like tracking still.

Steve: It does. "Browsers would need a way to form clusters that are" - browsers, that is, web browsers - "would need a way to form clusters that are both useful and private, useful by collecting people with similar enough interests and producing labels suitable for machine learning, and private by forming large clusters that don't reveal information that's too personal when the clusters are created, or when they're used."

So in other words, what that would sort of - the way that would happen would be that websites would have like a huge bitmap of things they could turn on that describe what they are. The browsers would collect those bitmaps as you went from browser to browser and accrue them. And then the browser would be able to present, would then be able to re-aggregate its user based on the aggregation of bitmaps that it had received from websites over time. Which, if I'm right, that's a lot of work. So who's going to do all that? It's not clear.

Anyway, they finish, saying: "A FLoCK," that is, what was the FLoC, was Federated Learning of Cohorts, so they then add a "K" to the end. The FLoC Key, or F-L-O-C-K, the flock, "is a short name that is shared by a large number, thousands of people, derived by the browser from its users' browsing history." So again, there's got to be a way for browsers to compare notes in order to do this. "The browser updates the flock over time as its user traverses the web. The value is made available to websites via a Client Hint." Okay, so actually this wouldn't have to be the individual sites. The browser's publisher, like Mozilla, could know about the sites you visit, and then it could be doing the aggregating in order to create these flocks that then its users are part of. So that would be a way of doing it.

And then, finally, they said: "A Privacy Model for the Web." "Sharding" is what they call it, sharding web identity. "The identity model for the web has been the implicit result of two interacting browser capabilities: per-domain state, especially cookies, which let one top level domain plus one" - that is, like GRC.com, TWiT.tv - "maintain a consistent notion of a visitor's identity. This identity extends across top-level sites due to third-party cookies, storage within frames, et cetera." And then the second thing is "in-browser passing of information among the parties co-occurring on a web page," in other words, by sharing state, the DOM, JavaScript, HTTP redirects and so forth, like where an ad is on a site. That site is able to get information shared with the third parties that it's sharing content with.

They said: "This combination has led to widely shared cross-site identities, and so to an ability to perform web-wide tracking of a person's browsing activity. Global static identifiers like device fingerprinting or like personally identifiable information provided by or covertly taken from the person browsing also offer an independent path to global identity. Limitations on cookies, fingerprinting, and other browser state all aim to reduce this ability to create or access a global identity."

So they say: "On the one hand, global identity gives rise to the capacity to weave together a record of much of a person's browsing history, a core privacy concern with today's web. Browsers are well-positioned to take action on this issue" - and of course this is the genesis of all this is Google is seeing this is happening, as you noted, Leo, with Apple and Firefox - "by imposing limits on the underlying capabilities exposed to developers. On the other hand, global identity also plays a substantial role in today's web advertising ecosystem. Browsers that impose limitations on these technical capabilities can directly affect publishers' economic viability and encourage workarounds, if they haven't provided for the legitimate needs of the ecosystem. This document describes a way the web could potentially work that would not require cross-site tracking, but would still let publishers support themselves with effective advertising."

They conclude: "We need a dialogue within the web platform community including browsers and the many stakeholders that thrive in and contribute to the ecosystem so that we can clearly describe a new identity end state that works for everyone. For browsers, this gives us a framework for evaluating proposed changes, standards, and APIs; what boundaries we need to enforce; and where we should innovate to improve the web. For developers, a clear long-term vision provides stability and predictability of the changes along the way.

"Any discussion of a new identity model must answer two specific questions. First, across what range of web activity does the browser let websites treat a person as having a single identity? And, two, in what ways can information move across identity boundaries without compromising that separation? This document offers one possible answer to these questions. The goal is a balanced way forward, dramatically improving web privacy" - or I guess we would say dramatically altering the nature of web privacy because we've seen that advertisers really don't want true privacy - "while allowing enough information flow to carefully support key needs of publishers and advertisers."

So anyway, I'm interested to see how this develops. And it does look like there is a non-tracking solution, but it's still going to be profiling because that's what advertisers want. And it's certainly - so this sort of changes our complaint. Our complaint has been "I'm being tracked, and I'm being profiled based on where I'm going." Google is saying, okay, we can arrange not to track you if you will still let us profile you. And so then we need to say, okay, was it tracking we objected to, or was it profiling?

Leo: Yeah.

Steve: And I think it's probably all the above.

Leo: It is such a challenge. I don't know if there's a good answer, and certainly Google's voice is welcome in this. I just, you know...

Steve: And the idea that they're opening this discussion, and it's going to be open standards, and maybe we'll get something to replace cookies and give users more control.

Leo: Is it cynical to say...

Steve: But, you know, I would just turn it off, Leo. I never - I don't click on an ad. I don't think I have in my entire life. I'm sorry.

Leo: Well, I have bought way too many...

Steve: Is it cynical to...

Leo: I've bought way too many Instagram things to think that ads don't work, Instagram things I don't need. I'm wearing my Instagram...

Steve: You were saying is it...

Leo: Is it cynical of me to think that Google's intent here is to enter the standards process and come out with nothing in a few years? That's what happened with Do Not Track, by the way. Six years in the making.

Steve: I know. I know.

Leo: And often...

Steve: And when you said to me, Steve, you're such an optimist, I was thinking, oh, yeah, I was all for DNT.

Leo: DNT. I mean, I don't want to be cynical, either. I want to be open to the idea that this is an important discussion that we have to have. And there are competing interests, and not all of them are malign. But I also - you could say cynically, well, Google's doing - this is a common strategy, which is to kind of embrace - the bear hug strategy. Embrace it. It makes you look good. And just make sure nothing happens.

Steve: Yes, I think Microsoft had...

Leo: Yeah, it's a bear hug. So I just don't know. I don't know. We will watch with interest. And, you know, the funny thing is this is a conversation that happens on Wednesdays on TWiG almost every week. You know, I mean, it's a really tough conversation, and I don't know if there's any answer to it, yeah.

Steve: Yeah.

Leo: I have been lately using Firefox because it does seem to have the best privacy protections.

Steve: I heard you say that on - I think it was on last Sunday's show, yes. I'm glad.

Leo: And I had to get used to it because it does look a little different.

Steve: Yeah.

Leo: But I turn it on, you know, even DNS over HTTPS is in there. And so I think as a single browser goes, this is a pretty - you have to go through the settings. But there's a lot you can do including turning off all the trackers, cryptominers, fingerprinters. You can say I don't want Twitter or Facebook or Google to know I'm here. And it does change the way your experience, you know, how you experience the Internet, but not necessarily for the worst. I can understand why people want to do that, I really can.

Steve: Yeah. Well, and I've mentioned, our listeners know that I'm a fan of uBlock Origin.

Leo: Which is not going to work anymore on Chrome, by the way, because of Manifest V3. They're really limiting what uBlock can do. Has Gorhill said anything about that?

Steve: Yes. He's not happy about it.

Leo: Right.

Steve: Yes, he has talked about it. And it's unfortunate because, boy, I look, sometimes I'll be on someone else's machine and just look at all this crap jumping up and down, and it's like, whoa. Whoa.

Leo: Not to mention the security, the weight of all that stuff, the megabytes you're wasting, the security of running these third-party - you talk about this all the time. JavaScript from god knows where.

Steve: Yup. And actually that's a place, it turns out.

Leo: God Knows Where. It's where all the bad scripts...

Steve: You do not want to go there, that's right.

Leo: Steve Gibson, we know where he comes from, GRC.com, the Gibson Research Corporation. That's his home on the web. Great place to go just to hang out, browse around. We have a Colonel, Colonel Newton in the studio, a long-time Army chaplain who loves your sleep formula. He uses it, and he says, "Boy, I travel a lot, do a lot of jet lag. This sleep formula really works." So that's great.

Steve: Cool.

Leo: That's there. That's free. SQRL. So many good things. While you're there, do us a favor. Support Steve, his bread and butter, the one things he sells, SpinRite, the world's best hard drive maintenance and recovery utility. You can pick yourself up a copy there. You can also pick up the 64Kb audio version of the show, the 16Kb audio version of the show, and it's the only place you can get the transcripts. Elaine Farris writes those, and they come out a few days after the show comes out. GRC.com. Leave him feedback at GRC.com/feedback, or his Twitter. He takes DMs from anybody, crazy man. His Twitter handle is @SGgrc.

We have audio and video of the show, if you want to watch Steve's moustache grey over the years. We should do a montage of you and me, just 14 years in.

Steve: How we have evolved.

Leo: How we've changed over the years. All the episodes are at TWiT.tv/sn, starting with TWiT.tv/sn1, going all the way through TWiT.tv/sn729, this episode, audio or video. Best thing to do, though, you don't need to do any of that, is subscribe. That way you'll get it automatically on your phone or on your laptop, on your device. Or you can ask your Echo or your Google Assistant. Just say, "Hey, Goog, play Security Now! podcast," and it'll just play the most recent version for you, which is the most easy way to do it.

I just looked at some stats. About 5% of the audience listens on a device. But that's pretty good. I think that's a number that's growing fast.

Steve: As opposed to...

Leo: Your phone or, you know, the way...

Steve: Oh. Oh, you mean, like...

Leo: A voice assistant. Actually, that's a bigger number than I would have thought.

Steve: And we should tell people about Boston; right?

Leo: Yeah, I'm so excited. This is going to be amazing. We are going to be the guests of LogMeIn and LastPass. We are going to do an - I'm calling it an "authentication summit." It'll be a live - now, tickets are going to be limited. I think it's only a hundred people. And I don't yet know how those are going to be allocated. But they will be available. It's for charity, so I think they'll probably be selling tickets, again, as a benefit. I'll have more details of that. And you know the date. I don't remember the date. It's in early November.

Steve: Yeah.

Leo: Or October.

Steve: October 3rd. October 3rd.

Leo: Thank you.

Steve: Is a Thursday.

Leo: October 3rd, Thursday. I will be moderating the panel. Steve will be there. We'll definitely talk about SQRL because that's what it's all about. Also the CISO for LogMeIn will be there, and the world-famous creator/inventor of the firewall, Bill Cheswick. He's also recently been writing a lot about why passwords don't work and has lots of thoughts. Really a smart fellow. Lots of thoughts about authentication, passwords, and what we do going forward. So I think it's going to be a fascinating summit. We will record it and offer it later for download. So if you can't get to Boston October 3rd, don't worry. You'll be able to see and hear it all. But more details to come on how to get tickets and all that other stuff.

Steve: Cool.

Leo: Thank you, Steve. Always a pleasure. Have a great day.

Steve: Okay, buddy. Bye.