SECURITY NOW!

Transcript of Episode #703

Out in the Wild

Description: This week we discuss a number of ongoing out-in-the-wild attacks, along with a bunch of other news. We have another early-warned Drupal vulnerability that has immediately come under attack in the wild, and a 19-year-old flaw in an obscure decompress for the "ACE" archive format, which until a few days ago WinRAR was supporting to its detriment. Microsoft reveals an abuse of HTTP/2 protocol which is DoSing its IIS servers. Mozilla faces a dilemma about a wannabe Certificate Authority, and they also send a worried letter to Australia. Microsoft's Edge browser is revealed to be secretly whitelisting 58 web domains which are allowed to bypass its "Click-to-Run" permission for FLASH. ICANN renews its plea for the Internet to adopt DNSSEC, NVIDIA releases a handful of critical driver updates for Windows, and Apple increases the intelligence of its Intelligent Tracking Prevention.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-703.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-703-lq.mp3

SHOW TEASE: It's time for Security Now!. Steve Gibson is here. The show "Out in the Wild" talks about a bunch of new vulnerabilities out in the wild, including a terrible Drupal flaw. He'll also talk about how Apple's doing a great job with privacy in Safari, and a shocking file hidden away inside your Windows 10 that doesn't do what you think it should do. It's coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 703, recorded Tuesday, February 26th, 2019: Out in the Wild.

It's time for Security Now!, yes, indeed, that show you look forward to all week long, with Security Now!'s head honcho, the man in charge, Steven Gibson. Hello, Steve.

Steve Gibson: Yo, Leo. Great to be with you once again. Episode 703 and counting down to 999. No, just kidding. So there was no huge single piece of news that happened in the last week, but there are a number of ongoing attacks in the wild. So I titled this week's podcast "Out in the Wild" to just sort of say, well, yes, unfortunately, this is what's happening right now. And for some interesting reasons. It's always now a consequence of a disclosure of a vulnerability that is then jumped on immediately by the bad guys.

So we're going to discuss, as I mentioned, a number of these ongoing, out in the wild attacks. And we have a bunch of other news. We've got another early warning Drupal vulnerability that - it's one of the sources of these problems that immediately came under attack in the wild. Believe it or not, a 19-year-old flaw in an obscure decompression DLL for the ACE archive format. And Leo, I don't think I've ever used ACE.

Leo: I remember ACE, though. That's pretty old.

Steve: You do? Yeah, yeah, yeah. And until a few days ago, WinRAR - which is my standard archiving tool. I'll talk about why a little bit later. But it was supporting this ACE archive format to its detriment. It no longer does. But we definitely need, if we've got a lot of people here who are also WinRAR users, as I am, they need to update. And we'll talk about that.

We also have Microsoft's revelation of an abuse of the HTTP/2 protocol which is being used to DDoS, well, actually not DDoS, just DoS. A single connection can bring down an IIS server. Mozilla is facing a dilemma about a wannabe certificate authority that it doesn't really know what to do about. And they separately sent a worried letter to Australia, sort of complaining about, well, not even sort of, clearly complaining about the recently enacted legislation about privacy. Microsoft's Edge browser was revealed to be secretly whitelisting 58 web domains which were allowed to bypass its click-to-run permission for Flash.

ICANN has renewed its plea for the Internet to adopt DNSSEC. And NVIDIA released a handful of critical driver updates, mostly for Windows, although Linux, because there's some overlap in drivers, gets updates too, although the problems for Linux are not nearly as bad as they are for Windows. And Apple in the beta of 12.2 of iOS, and also the Safari for the next macOS, is increasing the intelligence of its Intelligent Tracking Prevention. So lots of stuff to talk about.

And earlier this week, actually it wasn't early this week, it was late last week, but previously, I had run across - a friend sent me a link to a tweet that just blew my mind on Twitter. So it's apropos because we're going to be talking about cookie things and Intelligent Tracking Prevention with Apple. And this was a WebKit guy commenting about something that they have seen which just I couldn't believe. So a fun Picture of the Week, and then lots of great news.

Leo: Nice. Shall we do the Picture of the Week?

Steve: Well, sure. This was a tweet from one of the Apple engineers on the WebKit team. And I just had to shake my head. We'll be talking about this at the end of the podcast. But he tweeted - this is Don Marti.

Leo: Oh, yeah. I know Don Marti, yeah.

Steve: "We've investigated reports of news site subscribers getting spuriously logged out, and found that trackers were adding so many cookies that the news site's legitimate login cookie got pushed out." That's just...

Leo: Wow. I didn't even know that could happen.

Steve: ...so wrong. Yeah. There is, I mean, but it's Ks. I mean, it's thousands of characters, but there is a limit on the size of the cookie header. And if you have enough trackers, I mean, and I'm seeing the number 70. We'll be talking about there was an instance of 70, seven zero, seven times 10, 70 tracking cookies that just - they were hogging all of the available cookie space with lots of data, which is ridiculous also. A cookie doesn't have to be big, it just has to be unique. So it could be, you know, 12 characters. But no. We're just going to not - we don't care. We're a tracker. We don't care. And as a consequence, people were like, wait a minute. I thought - why do I have to log in again? Turns out too many trackers just pushed the legitimate cookie off the end. Oh, it's just - it's so wrong.

Leo: Wow.

Steve: Anyway, Apple has improved the intelligence of their Intelligent Tracking Protection, and we'll be talking about that at the end of the podcast.

In the meantime, once again, Drupal has trouble. I'll start with the takeaway, like the most important thing I could say about this to any of our listeners who are responsible for Drupal sites is you want to be absolutely sure that you get their email. This time they released the news of a forthcoming critical release only the day before. So the announcement was on February 19th, which was, what was that, last Tuesday. Yeah, it was a week ago. Highly critical, 20/25, critical release. They said there will be a security release of 8.5.x and 8.6.x on February 20th, so the next day, between 1:00 p.m. and 5:00 p.m., they said America/New York, and then they gave the UTC, 18:00 to 22:00. To see this in your local time zone, refer to the Drupal Core Calendar. The risk on this is currently rated at 20/25, highly critical.

So, okay. So the point is that, because of the fact that this is a PHP-based content management system, it is ridiculously easy for bad guys to reverse engineer any change that they make. But I'm kind of getting ahead of myself. So first of all, not all configurations are affected. That was their statement. And they said: "Reserve time on February 20th" - meaning the next day from the date of this release - "during the release window to determine whether your sites are affected and in need of an immediate update. Mitigation information will be included in the advisory." And unfortunately the mitigation information was wrong.

And then they said: "Contributed module security updates may also be required. If you're running Drupal 7, no update is required, but you may need to update contributed modules if you are using an affected module." They said: "We are unable to provide the list of those modules at this time." Oh, and they wanted to make it clear, because they figured this would generate a lot of concern, they said: "Neither the security team nor any other party is able to release any more information about this vulnerability until the announcement is made."

So in other words, don't bug us. We're making everybody wait because, again, unfortunately the way the industry is developing is that we are now seeing attacks pouncing on updates and turning them into, well, attackers pouncing on updates and then turning them into viable attacks with increased speed because they recognize there is some window of opportunity. And the window, depending upon the environment, it could be open for years, as we have seen, or it could be hours, minutes, days.

Oh, I was confused by my own notes because I said hopefully if you're running Drupal 8.5x or 8.6x, this is old news to you, meaning to our listeners, because you are on the mailing list, you did get the notice of this, and you made some time on Wednesday of last week, which was the 20th, in order to respond. Two days after Drupal's disclosure, two different sites published proof-of-concept code. And the day after that, using one of those proof of concepts as its foundation, attacks against Drupal sites began.

Yesterday, Monday, Imperva Security summed up their observations for the preceding two days where they watched attacks against their customers whose networks they were monitoring as part of their service. They said: "Latest Drupal Remote Code Execution" - and I should mention, yes, it is a remote code execution flaw - "Used by Cryptocurrency Miners and Other Attackers." Imperva wrote: "Another remote code execution vulnerability has been revealed in Drupal," they wrote, "the popular open source web content management system. One exploit, still working at the time of this writing, has been used in dozens of unsuccessful attacks," they write, against their customers, with an unknown number of attacks, some doubtless successful, against other websites.

So there are two vulnerabilities, and we've talked about this before. It was some data in the queries that were not properly sanitized, which allow bad guys to essentially inject commands through the open and exposed Drupal interface. Unfortunately, what Drupal said was that - and these are REST modules, which use HTTP protocol. So HTTP, we often talk about GET and POST, and there's also PATCH and PUT and DELETE and a few other less used verbs. Well, Drupal said that, if you disabled PUT, PATCH, and POST, then you were fine. It turns out that the GET request was able to perform this so that Drupal's mitigation advice was incorrect in this case.

So what the bad guys do is, all they have to do, and in fact from one of the reverse engineering sites we know this is what they actually did, was they "diffed," as the jargon is. They looked at the difference between the pre- and post-patched source, immediately spotted what it was that was changed, and said, oh, and then went about designing an exploit. Which is far easier because, unfortunately, we have, as I said, a PHP source base release, rather than, for example, in the case of a Microsoft patch, it's some DLL, and it's necessary to go in and have a much higher level of reverse engineering skills in order to figure out what it was that got changed.

The compiler that produces the DLL may also tend to rearrange things if it's set up to do so or chooses to or if it's got optimization turned on. So it can help to obscure what's going on. It's just much less easy to do that when you're patching source code in PHP, which is ASCII. It's just text. And so you can run some software that finds the differences between the various files and easily design an exploit around that. So that's what these guys did. They diffed 8.6.9 and 8.6.10, and then designed an exploit.

So Drupal has built a very nice system, but it does suffer from having, first of all, a large exposed attack surface, I mean, it sort of has to by design. That's what it's for. And it's written in PHP, which means that changes to it are easy to spot. So anyway, as I said, the takeaway for our listeners is, if you are responsible for Drupal, be very, very sure that you are subscribed to their email list, and that you take these notices to heart. And they're now, I mean, I remember the last advisory from them, what was it, maybe three or four months ago, Leo? They gave the industry several weeks of notice. They said, you know, at some point in the future, and they told us when. But there was a larger window. I don't know...

Leo: This is a zero-day. That's the problem; right? It's already been in the wild.

Steve: Well, it actually took three days from their...

Leo: Oh, it's a three-day.

Steve: Yeah. So I'm a little curious to know why they rushed this out. Maybe they were not sure that they had containment of the nature of the problem and so they thought, okay, we just can't sit on this for too long. And so what that says is that this is not something where you want to be lazy about even checking your inbox if you might be getting incoming mail from Drupal, if you have the ability to flag email for your emergency attention by the email's source. I would put Drupal on that list of mail you want to see because right now they're only giving, I mean, if this is an example, they're giving people 24 hours to get ready. And it's not like, oh, yeah, we'll get around to updating this soon. I mean, you really need to do it quickly. So there's the first of our out-in-the-wild problems.

Leo: And we'll be updating. We are Drupal, and I love Drupal, but we have professional maintainers who are keeping an eye on it. Any open source project's going to have this problem, PHP or not. If you can read the source of the patch and compare it to the previous code, you'll be able to see it immediately. The only reason Microsoft doesn't is because they're not open. It's a binary blob; right?

Steve: Right, right. Well, in fact it's a little different because the deliverable is PHP. I mean, it would be...

Leo: Yeah, but you could do - I think you do obfuscated PHP. I'm not sure, but I think, just like with JavaScript, you know, you can obfuscate it. But it doesn't matter because the source, it's open, and so the source is going to be somewhere, I think. So, yeah, I think any patch in open source has this problem.

Steve: So WinRAR is my favorite archiving tool. With a little bit of tweaking, it's able to produce significantly smaller compressed archives than, for example, ZIP or many other tools. Yesterday, meaning Monday, what, the 25th, security researchers at 360 Threat Intelligence Center detected an in-the-wild malspamming email campaign which is distributing a malicious RAR archive file that exploits what was only a few days old, a newly discovered - you know, the press is a little confused about this. This is one of these deals where you can't simplify it or it becomes inaccurate because technically it isn't a WinRAR vulnerability, although that's how it's being painted, because WinRAR supports a large array of archive formats.

This ACE format was developed 20 years ago. It hit its peak popularity in 1999 to 2001. It was developed by a guy named Marcel Lemke and later bought by e-merge. It always remained proprietary. And back then there was some advertising or ad-supported form of it, sort of like, you know, and I remember I got a copy of WinZip once that was also adware. And I thought, okay, these guys have gone further than I want to go with them. WinRAR has been well maintained over the years, but it supports many different formats: CAB, ARJ, LZH, TAR, GZ, TAR.GZ, BZ2, TAR.BZ2, UUE, I remember that one. Even Java Archives, JAR files. And it also can open up ISO images, CD images, and allow you to see inside them, as well as 7-Zip and other formats. And it supports this ACE. The support for ACE came from a DLL.

And WinRAR is not alone in being an archiver that supports ACE. There are some others that do. WinRAR, however, is the most popular, and it's got about 500 million users. So a half a billion users of WinRAR have this installed on their machines. So it represents a rather large attack surface for bad guys. Unfortunately, WinRAR, or I guess maybe cleverly, WinRAR looks at the content of the file rather than the file extension. So you can be a victim if you open a file that says .RAR, thinking it's safe, when in fact it's a .ACE format. WinRAR will look at it, say oh, that's actually a .ACE format. And what happens is it invokes this separate DLL.

The point I was making was that, because this was a proprietary compression archive format, nobody had access to the source for this. And in some of the coverage of this, the WinRAR developers said, yeah, I don't have source for this DLL, so we're just going to remove it. So that's the upshot of this is that the capability of decompressing ACE archives as of a few days ago has been removed from WinRAR.

So the takeaway for our listeners is, if you, like me, are a user of WinRAR, go update. I'm not sure whether WinRAR has an update check. Mine didn't give me any notification. I went to check, found that there was a new version, and that the website no longer mentions ACE among the archive formats that it supports. It boasts of everything else, but not that one.

So what happened was Check Point were the ones who found a problem. And we've also been talking a lot recently about these path traversal problems. Well, here was another one. They found an absolute path traversal bug in this unace.dll. That is, nothing could produce ACE archives, but if you encountered one somewhere - and I don't know. I never have. But Leo, you at least remembered them. I didn't even remember that there was a .ACE.

Leo: I'm an old-timer. Oh, wait a minute, so are you. Never mind.

Steve: Yeah. Even older than you. So it turns out that all they could do was open them. Nobody was able to produce them, and nobody cared because, you know, who cares if you can create an ACE archive. Nobody else can open them. So anyway, somehow they found this path traversal bug which allowed someone to design an arbitrary code execution exploit which would plant an executable file wherever in the user's system they wanted - for example, in their startup folder, such that next time they logged in, this file that had been maliciously unarchived thanks to WinRAR's support of the ACE format would get executed, and then you could be in trouble.

So we are now at - WinRAR is at 5.70 beta 2. The news coverage, which was fresh, it was only a day old, mentioned beta 1. So it looks like there was another change that was made for some reason since then. I now have 5.70. I went to look beforehand and afterhand. And sure enough, beforehand there was an unace.dll. That's gone after updating to the new version of WinRAR. So what's happening is there is a spam campaign that is, you know, it's like click here to, what, get the news of the sweepstakes that you just won or something. And what it's actually doing is, thanks to this bug, it's able to plant an executable file somewhere that will get run next time the user logs in. You don't want to get bitten by that. So if you're a WinRAR user, be sure to update. And of course you should never be clicking links that are offering you sweepstakes winning notifications either way.

Gal Goldshtein of F5 Networks found and reported a vulnerability which affects Microsoft's recent implementation of the HTTP/2 web protocol. Our listeners know that back in the beginning of the podcast we were at /1.0, and then we went to 1.1, where we've been for quite a while. Now the world is at 2, all IIS servers running Windows Server 2016, Windows Server versions 1709 and 1803, as well as all Windows 10 versions 1607, 1703, 1709, and 1803. And you know, Leo, as I was putting this together I was thinking, isn't it nice how Microsoft has executed on their "We're only going to have one version of Windows now."

Leo: Well, it's the same version. Just updates.

Steve: What a disaster. That's right.

Leo: Can't just sit still. Nobody would want that; right?

Steve: Well, but it's necessary now to enumerate this trail of debris that you've left behind. So anyway, I'm obviously a little biased about this.

Leo: They've got to polish their turd, my friend. Somebody's got to.

Steve: That's right. It turns out that by setting up a maliciously crafted HTTP/2 connection, it's possible for any remote attacker with a single connection to bring one of those many IIS servers to its knees, producing what they called, kind of euphemistically, "IIS resource exhaustion." And it's the CPU that gets exhausted. It's able to pin the processor at 100% so that it's just essentially stuck in an infinite loop, and nothing else happens. The server locks up. No other connections are answered. And it maintains that state until the connection times out. The default connection timeout is 120 seconds, two minutes. However, as soon as the server comes back to life, the bad guy simply initiates another connection, similarly, and brings it right back down again. So it's possible, with no bandwidth whatsoever, to pull an IIS server down.

Microsoft wrote: "The HTTP/2 specification allows clients to specify any number of settings frames with any number of settings parameters. In some situations, excessive settings can cause services to become unstable and may result in a temporary CPU usage spike until the connection timeout is reached" - and I don't mean to sound so gleeful about this - "until the connection timeout is reached and the connection is closed."

The default IIS connection timeout is two minutes. This makes overlapping connections - oh, this is me speaking now because I did a little bit of research. I thought, well, how long is a connection, because I thought it was a couple minutes. The default IIS connection, as I mentioned before, two minutes, allowing a chain of connections in order to keep the machine offline.

In their advisory, Microsoft states that there are no known mitigations or workarounds for the vulnerability. So essentially it says, you know, the spec says any number of settings frames, with any number of settings. So probably someone puts in something like two billion, a large positive 32-bit number or something, and it just, you know, it's something that Microsoft's system was not prepared to deal with, and so it collapses it.

They have updates. There are 14 editions of Windows which are enumerated in their update list, which are covered by a total of four patch editions among them. So if you are - first of all, technically this affects Windows 10 because Windows 10 has IIS in it. But no home user typically has a publicly exposed web server. We're all behind NAT routers and happy. But certainly Server 2016, and if you are using Windows Server 1709 or 1803, then you do have a publicly exposed presence. If you're curious why your server's been going offline recently, maybe you already know. Anyway, there is a patch for this. So you want to choose the proper one of the four patches and install that. And then that should get you back up and going again.

So who do we trust to be in our certificate root store? We've talked about this...

Leo: Hong Kong Post Office.

Steve: We've talked about this quite a lot. A company called - and here's the first problem - DarkMatter, based in the UAE, the United Arab Emirates...

Leo: Oh, no, no, no, no.

Steve: I know, I know. Kind of like, okay, if you want to be trusted, you're going to say, hey, let's call ourselves DarkMatter. It's like, okay.

Leo: Well, it gets worse.

Steve: It does. They're petitioning Mozilla to include their certificate authority root cert in Firefox. The problem is, DarkMatter has been known to sell surveillance and hacking services to oppressive regimes throughout the Middle East. And last month a report by Reuters further described DarkMatter's involvement in helping the Saudi government spy on dissidents. Reuters said - their coverage was titled "Project Raven: Inside the UAE's secret hacking team of American mercenaries. Ex-NSA operatives reveal how they helped spy on targets for the Arab monarchy - dissidents, rival leaders, and journalists." Okay, and I won't go into any more detail. For anyone who's interested, it's a fascinating report that Reuters put together, and I have a link in the show notes.

So as we know, I'm a Firefox user, and I don't want any CA root cert from a company who chose to name itself DarkMatter anywhere near my machine. No thank you. The only possible reason to be carrying such a certificate is if I were going to be visiting a website whose TLS certificate was purchased from DarkMatter. So that's a chance I'm happy to take. You know? I don't want that cert in my machine.

And it's true that certificate mis-signing is increasingly difficult to pull off in today's world with the degree of welcome certificate issuance oversight that we now have in our industry. But the benefit, at least to me, and to probably our listeners who are not maybe in the UAE and might visit a website that got their cert from DarkMatter - and, by the way, that's apparently zero right now so it's like, okay, why start? - the benefit seems marginal at best, compared to the risk.

So Mozilla, not surprisingly, is under pressure by the Electronic Frontier Foundation, also Amnesty International and the Intercept, to decline DarkMatter's request. But DarkMatter, which does have the ability to issue certificates because it is trusted by another well-placed certificate authority, QuoVadis, claims that it has never abused its TLS certificate issuance power to do anything bad. So there's no technical basis...

Leo: Of course not.

Steve: Oh, yeah. Oh, not us.

Leo: No, we never did anything like that.

Steve: So there's no technical basis for Mozilla treating it with less trust than other CAs that have applied for this same privilege in the past. And DarkMatter wishes, apparently, I mean, technically they can issue certs, but their certs are trusted only because QuoVadis has signed their cert, so they're an intermediate. They wish to move from sort of a second-class CA up to first-class status. And why not? Of course they want that.

So concerns are further heightened because Mozilla's list of trusted root certificates is also used by some Linux distros. That is, that's where the Linux distros get their root store. So there are fears that, once approved and added into Mozilla's certificate store, DarkMatter would then be able to issue TLS certificates to intercept Internet traffic without triggering any errors or warnings on Linux systems, which are often deployed in datacenters and cloud service providers.

So the point is it kind of - it sort of perniciously creeps out from just being browser authentication to cloud service providers authentication. So Mozilla has a dilemma. But there are many who are not the least bit ambivalent. The EFF's Cooper Quintin said - oh, Mozilla opened a Google Group discussion to sort of air this publicly and get opinions. And, oh, did they. The EFF's Cooper Quintin said: "Given DarkMatter's business interest in intercepting TLS communications, adding them to the trusted root list seems like a very bad idea." He wrote: "I would go so far as revoking their intermediate certificate, as well, based on these revelations."

Quintin expanded on his fears in a post on the EFF blog, reminding Mozilla that it went through a similar issue 20 years ago, back in 2009, with CN-NIC. That's the Chinese government's official certificate authority. Back in 2009, Mozilla approved CN-NIC as a trusted root CA in Firefox. Then, six years later, that CA was caught mis-issuing certificates for Google domains back in 2015, which we covered at the time, which allowed threat actors to intercept traffic meant for Google sites - which got CN-NIC banned from most certificate stores.

And the outcry against this CA addition is overwhelming in its support for not doing this. Mozilla publicly posted that: "Mozilla's Root Store Policy grants us the discretion to take actions based on the risk to people who use our products. Despite the lack of direct evidence of mis-issuance by DarkMatter, this may be a time" - you think? - "when we should use our discretion to act in the interest of individuals who rely on our root store." And to that I say, "Amen."

So, you know, if they feel like for some reason they absolutely have to do this, then I would vote for an option, maybe call it "Allow Sketchy CAs" and have it off by default, or have it produce a warning. And then, if you're absolutely sure that you want to do this, well, we do know about this root cert which is disabled by default. And if you're sure you want to turn it on, then okay. Anyway, it sounds like saner heads are prevailing here, and DarkMatter is going to have to go do something else. I mean, or keep signing with - and maybe the reason they're not happy developing a business around being an intermediate is that they do feel themselves endangered. It's like, well, sorry, but it's about trust. And if they're in the business of intercepting TLS communications, then you don't get to be in the root store.

Meanwhile, Mozilla worries that its employees could be subject to Australia's legislation. And I guess this was a close reading of this legislation that we've talked about before. Sophos, their headline, caught my attention and initially puzzled me because their headline was "Mozilla fears encryption law could turn its employees into insider threats."

So last Friday, on the 22nd of February, Mozilla wrote to the Committee Secretary of the Australian Parliamentary Joint Committee on Intelligence and Security. And they said: "Regarding Comments for Parliamentary Joint Committee on Intelligence and Security (PJCIS) Review of the Telecommunication and Other Legislation Amendment."

And so Mozilla said: "Thank you for the opportunity to provide comment as part of your review of Telecommunication and Other Legislation Amendment (TOLA)." Mozilla said: "This legislation grants sweeping and dangerous new powers to Australian law enforcement and intelligence agencies and, thanks to the foreign assistance provisions, extends these powers to foreign authorities, as well. In doing so, this legislation raises grave concerns for the security of Internet users and infrastructure in Australia and abroad, and fails to place appropriate limits on government surveillance. Given the serious threats to security and privacy posed by this Act, we welcome the Committee's review of this legislation and urge you to move swiftly to ameliorate its harms."

They said: "Mozilla's mission is to ensure the Internet is a global public resource, open and accessible to all. Our flagship product is Firefox, which is an openly developed and open source web browser used by hundreds of millions of people worldwide." I'm one. They said: "The Firefox code base is also used for the Tor browser, which allows anonymous browsing. In addition to protecting the security of our products, Mozilla has influenced core security protocols used in the Internet and backed the adoption of HTTPS, which encrypts website connections to enable more private and secure browsing. In addition, we have advocated to judges and policymakers in many countries on the importance of transparent and robust government processes to handle security vulnerabilities and surveillance requests.

"As we noted in our submission to this Committee when this legislation was initially under consideration: 'Any measure that allows a government to dictate the design of Internet systems represents a significant risk to the security, stability, and trust of those systems. Mozilla believes that TCNs or any similar device would significantly weaken the security of the Internet.'" They said: "We do not believe that this law should have been passed in the first place; and we believe the best possible path is to repeal this legislation in its entirety and begin afresh with a proper, public consultation." And then they go on to say that they know that's unlikely to happen, blah blah blah. They give eight points that they think really need attention. I'll just focus on the first because that was what created the headline that caught my attention.

They said: "1. Clarify that Australian authorities cannot target an employee of a Designated Communications Provider." And they said: "Due to ambiguous language in TOLA" - the T-O-L-A that I talked about before, that's this telecommunications and other stuff - "one could interpret the law to allow Australian authorities to target employees of a Designated Communications Provider rather than serving an order on the DCP itself through its General Counsel or otherwise designated official for process. It's easy to imagine how Australian authorities could abuse their powers and the penalties of this law to coerce an employee of a DCP [Designated Communications Provider] to compromise security of the systems and products they develop or maintain."

And of course we've talked about this back in the Snowden days, the idea that individuals within a corporation could have been subjugated one way or another. And Mozilla's concerned that this is essentially codifying this in legislation that could allow this to happen. They said: "In order to ensure due process, appropriate diligence, and full compliance where appropriate with orders issued under this legislation, we strongly believe that Australian authorities should only serve an order on the Designated Communications Provider (DCP) itself. Serving an order on an individual employee rather than a DCP would fail to allow a DCP to avail itself fully of the protections afforded under this legislation in regards to consultations, assessments, and legal challenges.

"Further, this potential would force DCPs to treat Australia-based employees as potential insider threats, introducing another vector for compromise that could undermine trust in critical products, incentivizing companies to move critical roles to other localities. Parliament recognized the wisdom of this limitation in regards to Contracted Service Providers, but not DCPs."

So anyway, Mozilla is point out that, hopefully as an omission that could be corrected, where there was a carve-out for contracted service providers, the legislation needs the same thing for Australian-located employees of non-Australian communications providers like Mozilla because they don't want their Australia-based employees to be subject to the legislation as written. So anyway, they finished in bold italics, saying: "We recommend the Committee: Add a clarification in the Section [and it happens to be] 317B definition of Designated Communications Provider to specify that this term 'does not include a person who performs such services in their capacity as an employee, agent, or vendor of the provider.'"

So anyway, this legislation is still obviously in some flux, and they're entertaining comments, and the industry is responding. So I salute Mozilla and thank them for adding this clarification, which we need. Or at least they feel we need.

Okay. This is a real headshaker, Leo. It turns out that Ivan Fratric - who we've talked about before, he's at Google's Project Zero - he was curious about a file located in the Windows\system32 directory whose name was edgehtmlpluginpolicy.bin. So the file was not named in any way to obscure its intent: edgehtmlpluginpolicy.bin. But what was obscured was its contents. For reasons that really could only have been to hide what it was doing, the file contains SHA-256 hashes of domains which are allowed to bypass Edge's click2play Flash blocker.

Leo: Hmm.

Steve: Yeah. And you should scroll down in the show notes because I have a list of them. And some of them are a little distressing. So what we now know as a consequence of Google's Project Zero work is that Microsoft's Edge web browser comes with a deliberately obscured whitelist, specifically, which, now, it's been whittled down as a consequence of Project Zero's disclosure. Now it only allows Facebook to circumvent the request for user consent with its...

Leo: But what's StupidVideos going to do? Or ontvtime.ru? What are they supposed to do now?

Steve: I know. Gee, they're going to have to ask for permission before they run Flash content in your Edge browser. Oh, boohoo. Yeah. I liked dilidili.wang. Oh, my goodness.

Leo: I don't even want to know what that is.

Steve: Okay. So back on November 26th, Ivan posted: "In Microsoft Windows there is a file, C:\Windows\system32\edgehtmlpluginpolicy.bin, that contains the default whitelist of domains that can bypass Flash click2play and load Flash content." And, I mean, we all know what, I mean, like, how bad this is. Like Flash can't die soon enough, and for some reason we're postponing its death, what is it, till next year, 2020? Or is it 2022? I think it might be 2022. It's like die now. Anyway, "to load Flash content without getting user confirmation in Microsoft Edge."

So today's updated version of the previously secret Edge whitelist only allows Facebook to bypass Flash's click2play policy. So www.facebook.com and apps.facebook.com are the two Facebook domains still allowed to run Flash content without getting users' okay. So technically, I mean, it was called a "bug report," but it certainly wasn't a bug. Ivan also highlighted the security implications of even having a Flash autorun whitelist bundled with a web browser.

Leo: I don't know why anybody would be worried if dilidili.wang could get through, or totaljerkface.com.

Steve: Oh, goodness.

Leo: It's very odd. I mean, I can understand, you know what, games.aarp.org or Facebook. I can understand that because they want to make it easy, and those are safe. And Vudu, maybe. But honestly.

Steve: I know.

Leo: Some of these it doesn't make...

Steve: NSEIndia.com? It's like, okay.

Leo: India.com is probably the government; right? Maybe not. I don't know who owns it.

Steve: Wasu.cn, so there's a Chinese site. I'd like to know.

Leo: Microsoft's in there.

Steve: What's WGT.com?

Leo: I don't know. See, I think a lot of these are kind of, okay, we know these are safe.

Steve: Deezer. Deezer. At least it wasn't Geezer.

Leo: It's a music service, actually.

Steve: Okay, it wasn't Geezer.

Leo: Deezer's a music service. I think some of these it's like, well, you know, we can whitelist these because they're established. But totaljerkface.com or dilidili.wang? I mean, what are those?

Steve: Yeah. And, I mean, I guess my argument is, if you're using Edge, and one of the selling features, not that anyone bought it, but one of the stuffed-down-your-throat features was that it had click2play protection for Flash content. And users are like, I mean, security-conscious users are like, yes, I want that. And then there's a - oh, and notice, Leo, this - okay. So the other thing is that they hid this. This wasn't ever documented. It wasn't stated. And in order for Ivan to determine the list, he had to reverse, he had to crack these SHA-256 hashes of these domain names in order to figure out what it was they were.

Anyway, but it's kind of worse because he says, as I was saying, he highlighted the security implications of having a Flash autorun whitelist bundled with a web browser, especially given the number of Flash security patches issued by Adobe nearly every month. "A whitelist," he wrote and explained, "is insecure for a number of reasons."

As we mentioned last week, by far the most common and most prevalent problem on today's web are cross-site scripting vulnerabilities. With this sort of domain name-based whitelist, a cross-site scripting vulnerability on any of the whitelisted sites would allow a bypass by a malicious site of the click2play policy. And moreover, there are currently publicly known and unpatched instances of cross-site scripting vulnerabilities on at least some of the whitelisted domains. So bad guys who knew of a Flash vulnerability and who knew of this Edge whitelist could employ a cross-site scripting vulnerability to essentially attack a user who was visiting a site not on the whitelist.

So anyway, as I mentioned above, the big issue reported by Ivan was partially addressed by Microsoft just two weeks ago. In this month's Patch Tuesday they trimmed the whitelist down to just Facebook. Oh, and they were also - I also forgot to mention that it also wasn't limited to HTTPS. So many of these domains did not support HTTPS at all, so you could only go to them without encryption, which allowed man-in-the-middle attacks to easily bypass this click2play policy that Edge was enforcing. I mean, it was a mess. So he reversed the hashes, found the 58 domains which were whitelisted. The good news is, as of a couple weeks ago, they are no longer all present. So dilidili.wang, which may or may not have supported HTTPS connections, won't be able to do anyone...

Leo: Nice job, Microsoft. Nice job.

Steve: Nice job. That's right. Super secure. So anyway, Tuesday...

Leo: What's N/A mean? Does that mean we don't know what it is, or it doesn't mean anything, or...

Steve: Maybe he wasn't able to reverse that one. Yeah, maybe that's like a mystery. And how about OK.ru? It's like, okay.

Leo: Yeah, that's not okay. That is not okay.

Steve: So Ivan tweeted, he said: "The default Flash whitelist in Edge," and he has a link, he said, "really surprised me. So many sites for which I'm completely baffled as to why they're there. Like a site of a hairdresser in Spain." And then he has a link. "I wonder how the list was formed, and if Microsoft Security knew about it." Anyway, so, yeah. Good news is, thanks to last Tuesday's update, that list has been whittled down to just Facebook.

ICANN, of course, I love ICANN, I-C-A-N-N, the Internet Corporation for Assigned Names and Numbers, last Friday put out a press release kind of begging for full DNSSEC deployment. The press release was titled "ICANN Calls for Full DNSSEC Deployment, Promotes Community Collaboration to Protect the Internet." And I'll share what they said. It's not very long. They said February 22nd, 2019, Los Angeles: "The Internet Corporation for Assigned Names and Numbers believes that there is an ongoing and significant risk to key parts of the Domain Name System infrastructure. In the context of increasing reports" - and we've covered this, like, maybe about a month ago - "of malicious activity targeting the DNS infrastructure, ICANN is calling for full deployment of the Domain Name System Security Extensions (DNSSEC) across all unsecured domain names." And I confess I'm guilty. "The organization also reaffirms its commitment to engage in collaborative efforts to ensure the security, stability, and resiliency of the Internet's global identifier systems.

"As one of many entities engaged in the decentralized management of the Internet, ICANN is specifically responsible for coordinating the top level of the DNS to ensure its stable and secure operation and universal resolvability. On the 15th of February, in response to reports of attacks against key parts of the DNS infrastructure, ICANN offered a checklist of recommended security precautions for members of the domain name industry, registries, registrars, resellers, and related others to proactively take to protect their systems, their customers' systems, and information reachable via the DNS.

"Public reports indicate that there is a pattern of multifaceted attacks utilizing different methodologies. Some of the attacks target the DNS, in which unauthorized changes to the delegation structure of domain names are made, replacing the addresses of intended servers with addresses of machines controlled by the attackers." In fact, we did go into this in detail about two weeks ago.

"This particular type of attack, which targets DNS, only works when DNSSEC is not in use. DNSSEC," they write, "is a technology developed to protect against such changes by digitally signing data to assure its validity. Although DNSSEC cannot solve all forms of attack against the DNS, when it is used, unauthorized modification to DNS information can be detected, and users are blocked from being misdirected."

Anyway, they go on. But we all get the gist of this. I've talked about DNSSEC for years, and about the benefit that we would have as an industry if and when we finally produce signed DNS. What it would give us is an incredibly powerful distributed ability to look up information which we still to this day, despite the fact that DNS has been around for a long time, don't have. So I just - I thought it was good that they're sort of reminding everyone and saying, look, here's a checklist of things to do. Please make this happen. And it reminds me that I need to get GRC's DNS zones signed also. There's just no excuse for not. I don't remember now what support Hover has for that. I'll have to take a look at that and see what they do.

Leo: You know, we register at Hover, but we use DNSimple for our DNS, and they seem to be very complete for kind of industrial-grade DNS. My guess is that Hover doesn't. Maybe we haven't covered this, but at some time we ought to. The DNS Flag Day from ISC, are you familiar with that? So February 1st, I guess some sites who weren't supporting DNS extensions would stop working on the Internet. So was it the ISC created DNS Flag Day to let people know.

And but then they did another thing, which was they flagged - they said, look, this site will continue to work. It's DNSflagday.net. This site will continue to work, but it's not fully compliant. And one of the things they said is, if sites aren't fully compliant, won't be able to implement DNSSEC in the long run, I think, was the idea. I'm looking at this again.

Steve: It must just be like ancient BIND DNS servers that do not support...

Leo: No, no, no, a lot of them don't. So it's extended DNS, which wasn't part of the protocol.

Steve: Right, EDNS.

Leo: EDNS.

Steve: Right.

Leo: And as you know, it's one of the reasons that you can get DNS amplification attacks, because you can use those extended fields to flood.

Steve: Yeah.

Leo: And Hover does not support the full extended DNS stuff. And I asked them because Greg Ferro was all over them for it. So I talked to their guy, and he said, yeah, we've looked at it, of course. We certainly could implement it. We decided not to for a number of technical reasons. One is that it would break DNS for some of our customers. Then you always want to err on the side of being functional.

Steve: Yeah.

Leo: And he said most people don't, you know, if you're doing, like TWiT, if you're doing big pulse, lots of traffic DNS, you wouldn't use Hover as your DNS provider anyway. You'd use, as we do, DNSimple. You'd use a bigger provider. It'd be like using your ISP for Amazon.com. You just wouldn't do that. So he said we're basically a retail DNS provider. We provide it as a service for our clients who just want simple websites and whatever. But we don't really think it's something - we're reluctant to implement it right away. And I think that, and I may be wrong on this, and I'm looking through the DNS Flag Day site, but I think that the EDNS is a prerequisite for Secure DNS, but I may be wrong on that. Maybe you know more.

Steve: It probably is. And I think, as I recall now, I mean, I'm still running - I'm running, I think it's FreeBSD version 4. And BIND.

Leo: You're using BIND; right?

Steve: I'm using BIND, like version 2.

Leo: Yeah, no, no. So you're not. So you're not, yeah.

Steve: So I have a newer Unix machine which I just haven't brought online yet. And that one would be using - I don't think I was using even BIND. There's a much nicer DNS. Shoot, I can't remember the name of it now. But I've got it all set up and configured. I just haven't deployed it.

Leo: Well, let me run - you want me to run GRC.com through this tester?

Steve: No.

Leo: You think it might not be compliant? No, you know what? All okay. Your domain is perfectly ready. You do not need to worry about DNS Flag Day. Your DNS administrator is doing a good job. Send them a sincere thanks. Are you your - I'm my own DNS administrator.

Steve: I am. Yeah, in fact, so what I do is I run the master DNS, and then I've got a pair of Level 3 servers that are my public-facing...

Leo: That's probably what they're hitting; right?

Steve: Yeah. Oh, yeah, exactly, yes. So they're quite happy with me.

Leo: They say BIND pre 9.14 does not. Not resolver does. PowerDNS Recursor 4.2 and earlier and Unbound 1.9 and earlier. Anyway, Hover was very aware of it. And they said, you know, given what we do in our business, we're really a registrar, not a DNS provider. We don't want to break people's DNS.

Steve: Yeah. One of the problems is - now I remember what's going on is that I have to do some rejiggering because GRC has pseudo DNS servers for its domain. Maybe that would - I have to see whether that would affect us or not. Because, for example, when the DNS Benchmark checks to see if there's a new version of it - and I did update DNS Benchmark after many years to add 1.1.1.1 and 9.9.9.9 - I actually do it with a DNS query because it's super easy and lightweight. And so the DNS Benchmark issues a DNS query which receives a response which can change. And so the whole deal with signing your zone is you can't have your DNS changing because it's, you know, you're basically - you have to have the zone signed as it is and then have the signature verifiable. So I may have to do some changing. But I wonder if maybe I could sign it without the variable. Anyway, I have to look into it because it is, you know, I should practice what I'm preaching, and I'm certainly loving the idea that we'll get DNSSEC widely deployed someday. That would be great.

NVIDIA has released some updates to their drivers, which if you are especially a Windows user, it's worth doing. I don't know if NVIDIA drivers will flag their own need to update. It turns out I am using NVIDIA-based display adapters. And I went looking, and my drivers were not affected, which is to say, mine were dated 2016, and there's nothing newer available. But NVIDIA has released, I think it's patches for their drivers covering eight flaws, which can lead to code execution, although not remote code execution because obviously it's a video driver. I mean, there's no public exposure. But code execution, escalation of privileges, denial of service, or information disclosure on both Windows and Linux machines, far more critical for Windows than for Linux, but still worth doing.

The exploitation of these problems which have been patched does require local system access. They're not remotely exploitable. But bad guys that have some means of running code on a machine could take advantage of these problems, if not patched, to execute code and elevate their privileges to allow them, you know, install a rootkit and so forth. The CVSS rating system is a 10-point scale, with five of the eight vulnerabilities for the Windows drivers receiving an 8.8 out of 10. So they're bad. NVIDIA's own documentation, they said, for example, there's a CVE-2019-5665 from their page, and I've got a link to the downloads in the show notes.

They wrote: "NVIDIA Windows GPU Display driver contains a vulnerability in the 3D vision component in which the stereo service software, when opening a file, does not check for hard links. This behavior may lead to code execution, denial of service, or escalation of privileges."

Another example: "The GPU Driver contains a vulnerability in the kernel mode layer create context command in which the product uses untrusted input when calculating or using an array index. But the product does not validate or incorrectly validates the index to ensure the index references a valid position within the array, which may lead to denial of service or escalation of privileges." And there are even some, oh, here's one: "Contains a vulnerability in the kernel mode layer handler in which the application dereferences a pointer that it expects to be valid, but is NULL, which may lead to code execution, denial of service, or escalation of privileges."

Anyway, as I said, I'm not sure that NVIDIA drivers proactively notify their users if they are in need of update. But our listeners, who are certainly security aware, if you are NVIDIA-based, I would go checking. It's just www.nvidia.com/download/index.aspx. Check to see if you have updates available to your drivers because it sounds like you don't want to have something there that bad guys are going to know about quickly and be able to reverse engineer and then leverage to take advantage of the vulnerability.

Let's see if there's anything else that I had to say about that. Oh, yeah. Mine have not been updated since 2016. Nothing new was available. But I think our listeners should check.

Leo: Steve Gibson continues.

Steve: So Apple has increased the intelligence of their already intelligent, now it's more intelligent, Leo...

Leo: Super intelligent.

Steve: It's super intelligent. Well, I wouldn't go that far. But I am impressed by how much they're really focusing on this. God, how many hours have we spent on this podcast talking about tracking. This is the Intelligent Tracking Prevention, ITP, which now goes to - they're calling it version 2.1. The beta release of iOS 12.2 and Safari 12.1 on macOS, which will be High Sierra and Mojave, they include this updated version of Intelligent Tracking Protection. One of these things is so obscure, I thought, okay, wait. And I had to read this several times to get a hold of it.

So their stated goal is to further reduce trackers' ability to establish user identities across sites. Of course that's what tracking is; right? To know that somebody who was at Facebook, then went to Amazon, and then went to, you know, did something at Google, and then went over to dilidili.wang or, you know, so forth.

Leo: Don't track me, man.

Steve: So previous versions of this Intelligent Tracking Protection allowed domains that were classified with tracking capabilities to store what Apple called "partitioned cookies." And those were cookies keyed off of the top site. I'll explain partitioning in a second. Then they said: "As of ITP 2.1, which is the more intelligent tracking protection, partition cookies are no longer supported, and third parties classified with cross-site tracking capabilities now have to use the storage access API as opposed to the standard cookie API." So here's how Apple explains what they call "verified partition cache."

They said: "WebKit implemented partitioned caches more than five years ago." They said: "A partitioned cache means cache entries for third-party resources are double-keyed to their origin and the first-party eTLD+1" - top level domain and then down one, so like Apple.com. So dotcom is the top level; Apple.com is one step down. They said: "This prohibits cross-site trackers from using the cache to track users." Okay, now let me stop for a second.

So the problem that we have is that the good news is Apple has been alone in the industry, and I've saluted them always, for by default, and here again this is the tyranny of the default, by default blocking third-party cookies. When I did that cookie forensics work years ago, I was compiling cookie - the visitors to GRC, compiling their browsers' cookie handling by browser. And all of the browsers were like up near 100% except Apple alone, with a very, very low percentage in third-party cookies enabled because it was off by default. Meaning that by default almost all of the Apple visitors or Safari users who visited GRC had third-party cookies disabled. So very nice.

However, by getting up to various tricks with JavaScript, as we've talked about before, it's possible to circumvent the third-party protection by essentially establishing a first-party presence using JavaScript and iframes and other tricks. So what partitioning does is very clever. It's something that Apple alone has been doing which tags the cookie that's being issued in a first-party context to the parent first-party context. That is, like, the actual domain on the URL of the page where a first-party cookie is being issued where it's different from the page's primary domain. That's what they mentioned when they say it's double-keyed to the origin of the first party and to the domain of the page.

So essentially it allows a cookie to be set by somebody, for example, an advertisement which is being hosted and then getting around third-party protection, yet it no longer - so it allows it to be set, but it's only going to be retrieved when that domain is being visited from the original domain where it was issued, thus preventing tracking.

Anyway, they say, continuing with Apple's description: "Even so, our research has shown that trackers, in order to keep their practices alive under ITP" - that's Apple's previous Intelligent Tracking Protection - "have resorted to partitioned cache abuse. Therefore, we have developed the verified partitioned cache." Which is what makes the new one more intelligent.

They said: "When a partitioned cache entry is created for a domain that's classified by the Intelligent Tracking Protection as having cross-site tracking capabilities, the entry gets flagged for verification. After seven days, if there's a cache hit for such a flagged entry, WebKit will act as if it has never seen this resource and load it again. The new response is then compared to the cached response." So basically WebKit is faking it out, saying, huh, don't know about you, when in fact it does. But that's only in the case if it's greater than seven days. It'll compare the cached response to the new response.

"If they match in the ways we care about," Apple wrote, "for privacy reasons, the verification flag is cleared, and the cache entry is from that point considered legitimate. However, if the new response does not match the cache entry, the old cache entry is discarded, and the new one is created," meaning the tracking is blocked because it won't be returning the one that it had from before. "The new one is created with the verification flag set, and the verification process starts all over." And if everybody's just completely gone cross-eyed listening to this description, I don't blame you. As I said, I had to read it several times in order to say, what what what what? And so this is the extent to which Apple has engineered tracking protection. So again, I tip my hat to them. I'm glad they're on our side and that they're doing this.

The upshot of this is that third-party cookies are blocked by default. Also another thing that they're doing is that session cookies for domains not visited for 30 days are deleted. So you can still have the "I want to remain logged in." But if you don't go back to a domain where you said you wanted to be logged in, like Google or Apple or Facebook or something you're doing enough, if you don't go back within 30 days, it will prune that. And this is Apple trying to get control of this massive cookie abuse that we talked about with our Picture of the Week where there are just so many ridiculously large cookies now that it's hard for them to be used for the actual valid purpose for which they were intended.

And the other subtle thing is cookies can be set in two ways. The normal way of setting a cookie is to receive it as a set cookie header in a response from the web server. That's how it's traditionally been done. And you want those to both be marked as secure, meaning the cookie can only be read over HTTPS, and you want it to be also flagged as HTTP only, meaning only with the HTTP protocol, not scripting. The reason that's important is that you don't want your scripting on the page to be able to see cookies which may have sensitive content in them, that is, to allow session hijacking if the cookie's value were readable by script. But that's one way of setting the cookie.

The other way is there is a document.cookie property that the Document Object Model that we talked about last week, the DOM, maintains. So script is able to set cookies. Now this more Intelligent Tracking Protection is separating cookies received from a web server, which are first-class cookies, from those that are set by script on the page. Cookies that are set by browser script are removed after seven days. So again, that's one of the ways, for example, that tracking is being done. And Apple, in seriously looking at where all these cookies are coming from, is working to pare this down. And so cookies set by script only get to last a week, and then they are removed.

So again, hats off to Apple for continuing to come up with ways, basically, they're struggling not to break anything that we want to have work while deliberately breaking all of the clever workarounds that are the extent to which trackers are going. And think about that. Think about the extent to which tracking is being, like, just held onto. It suggests that it is really valuable for advertising services and others to maintain a grip on users and profile where they go, the fact that they're willing to go to such degrees.

So again, Apple is alone in the industry, I mean, even in disabling something as simple as disabling third-party cookies by default. There was a point at which IE, I don't remember which IE it was, a while ago Microsoft said, yeah, we're going to turn off third-party cookies. There was such an uproar from the industry that they ended up never shipping a browser that had third-party cookies disabled, even though Apple always has. Safari has that.

Leo: Says something about market share more than anything else, frankly.

Steve: Yeah, probably does, yeah. And Leo, that's our show.

Leo: Wait a minute. You've got to plug SpinRite, the world's finest hard drive recovery and maintenance utility.

Steve: Well, yes.

Leo: Please buy it, folks.

Steve: I did want to thank our listeners. I saw, after I mentioned that the support that this podcast's listeners have provided me lo these five years while I have been working on SQRL, I saw an effect in SpinRite's sales from that.

Leo: Good. Thank you.

Steve: So I have been conscious of that ever since mentioning that. I really appreciate the help. It, like, keeps the ship afloat. And I actually did have a person, Van Zeck, posted in the SQRL Forum. His subject was "Wow!" exclamation point. He said: "The magic is real," he said, "and in my hands." And the good news is it will soon be in everyone's hands.

Leo: That's cool.

Steve: He said: "Steve, I just used Jeff's" - that's Jeff Arthur, the guy who's done the iOS client. He said: "I just used Jeff's iOS client to log into the SQRL Forums for the first time. As I posted in Jeff's area, fantastic!" exclamation point. He said: "I have read about and watched others use the magic, but it was a real rush to have the magic right in my own hands. Thanks for persevering with SQRL and showing how it is possible to potentially eliminate the biggest hassle and risk in Internet life - passwords." He said: "I have been lurking around SQRL for all five years, and it is truly exciting to see things come to fruition. Van."

Leo: Nice.

Steve: So Van, thank you for sharing your reaction. And again, to our listeners for your patience, those who are waiting for the next version of SpinRite, I can't wait to get back to it. And we're getting close. Where I am now is Rasmus is actually adding some additional features to SQRL's support for XenForo, which is our forum software. I'm building, finishing up on the static content for new SQRL users who will go and wonder what's going on and how they get started and so forth. And that'll be separate from all of the dynamic forum content. As soon as that's in place, then I start working on catching the specification stuff up to speed. But basically we're there. So it won't be long.

Leo: Nice. Yeah, when you actually - I think that's going to be the key to getting it out there is people using it. Because as soon as you see it, you go, oh. It's a little hard to understand, maybe. But when you use it, it's like, oh.

Steve: Oh, Leo.

Leo: Wait a minute, what happened?

Steve: That's just it. I mean, it is - and I think, you know, I realize, I mean, I respect your skepticism and the skepticism of others who are like, okay, what chance does some guy have of getting some random protocol established? Well, first of all, it's more than a guy now. There's a community. But the experience is, first of all, it is so secure, and the experience is so friction-free, that when you use it, you're like, wait, what just happened? And, like, could this possibly be secure? And then it's like, okay, why don't we have this everywhere?

Leo: Just to be clear, my skepticism is not the technology. I completely believe the technology and understand that it works.

Steve: Right.

Leo: And I would be thrilled to no end if the world would adopt it. I just - I worry that it's kind of an uphill battle, given you've got FIDO and Google and Microsoft and everybody in the world; you know.

Steve: It is. I recognize that. But how could I not...

Leo: You're got to do it; right?

Steve: ...offer it if it's possible?

Leo: Yeah.

Steve: So...

Leo: No, absolutely. I mean, if Tesla had said, eh, this Edison guy's got it all, I'm just going to stop, he's got it under control...

Steve: Oh, my god, we'd have DC, and nothing would work. We would. Edison was DC. Tesla was AC.

Leo: Yeah, yeah.

Steve: And AC is so much better. I mean, you could have transformers. You can't transform DC.

Leo: You see? You see? And they thought he was crazy. In fact, he was basically persecuted to death.

Steve: Yeah.

Leo: So I'm not - I don't think you're going to get persecuted to death for SQRL.

Steve: I'm not worried.

Leo: But if it happens, I'll be the first to send you a card.

Steve: Gee, thanks.

Leo: No, I love you, Steve, and I'm so glad you're here every Tuesday, 1:30 Pacific, 4:30 Eastern, 21:30 UTC to do Security Now!. You can watch live at TWiT.tv/live. You can watch or listen live. You can join the chatroom if you're doing that at irc.twit.tv. Of course Steve has on-demand audio that you can get anytime at your convenience. That's at his website, GRC.com. And he also has transcripts there, which is really nice. Not computer-generated, real human written transcripts.

Steve: Yeah, boy, she got snowed in, too. Elaine's been going through some...

Leo: She's in the High Desert? Where is she?

Steve: Yeah, wherever she is, I mean, I remember snow on a cactus. So that would be a clue. Whew.

Leo: Wow.

Steve: Yeah.

Leo: It sounds like the High Desert. Wow. If you're at GRC.com, yes, pick up a copy of SpinRite, the world's best hard drive recovery and maintenance utility. And then, once you've paid your fare, you can visit all the rest of the site - well, you can do it anyway - free. There's all sorts of good stuff there, including all the deets on SQRL.

We have audio and video at our website, TWiT.tv/sn. Hey, if you're over there, you might take the survey. I want to make sure that a lot of Security Now! listeners respond to the survey because we want to get your unique point of view, get the nerdy point of view. Yes, I'm looking at you, Dave Redekop. Go to TWiT.to/survey19. It's our yearly attempt to get to know you a little bit better. But don't worry. No salesman will call. We're not tracking you. We're not interested in any personally individually identifiable data. It's the aggregate that makes the difference. Thank you, though. If you don't have to do it, but if you do, thank you in advance at TWiT.to/survey19.

And I think we can wrap this sucker up for the day. Don't forget to subscribe to the show. You'll get it every Tuesday when it's available. Thank you, Steve.

Steve: See you in March, my friend.

Leo: What? You taking a one-week vacation? What?

Steve: No, next podcast.

Leo: I know, yeah.

Steve: Yeah, it's like, whoa, March already? Wow. Yeah, a one-week vacation, exactly. Seven days.

Leo: See you in seven days. Thanks, Steve.

Steve: Bye.

Leo: Take care. Bye-bye.