SECURITY NOW!

Transcript of Episode #638

Quad Nine

Description: This week we discuss Windows having a birthday, Net Neutrality about to succumb to big business despite a valiant battle, Intel's response to the horrifying JTAG over USB discovery, another surprising AWS public bucket discovery, Android phones caught sending position data when all permissions are denied, many websites found to be watching their visitors' actions, more Infineon ID card upset, the return of BlueBorne, a new arrival to our "Well, THAT didn't take long" department, speedy news for Firefox 57, some miscellany, listener feedback, and a look at the very appealing and speedy new "Quad 9" alternative DNS service.

High quality (64 kbps) mp3 audio file URL: http://media.GRC.com/sn/SN-638.mp3
Quarter size (16 kbps) mp3 audio file URL: http://media.GRC.com/sn/sn-638-lq.mp3

SHOW TEASE: It's time for Security Now!. Yay! Steve Gibson is here, our Turkey Day edition, and there are a few turkeys to talk about, including the Uber breach. We'll talk about a new DNS service that Steve's going to switch to right away. And of course our Picture of the Week, exactly why you have to be careful if you're an IT professional where you go to work. It's all coming up next on Security Now!.

Leo Laporte: This is Security Now! with Steve Gibson, Episode 638, recorded Tuesday, November 21st, 2017: Quad 9.

It's time for Security Now!, the show where we get together with Steve Gibson, our security officer - CSO, Chief Security Officer - to learn about what's going on on the worldwide - what the heck is it? Hello, Steve.

Steve Gibson: Yo, Leo.

Leo: Good to see you.

Steve: Great to be with you again for our 638th episode.

Leo: Wow.

Steve: Had we begun numbering from zero, of course, 638 would be the 639th episode. But we're unable to go back and change time.

Leo: Is that the first thing you'll do if you get a time machine?

Steve: Yeah, we really should have renumbered this. We should have. Negative one, that'd be a fun start.

Leo: Oh, yeah.

Steve: Especially since it was about the Honey Monkeys. So that one could have been the negative one episode. But no. So this is the week of a Windows birthday, which we're going to talk about. Not a big birthday, but it was yesterday, so I thought, oh, that's kind of fun. The RTM of Windows 1.0 happened on November 20th. That was yesterday since this is the 21st that we're recording this. So there's that.

And I heard you talking on MacBreak Weekly about Net Neutrality, so we just - we'll spend much less time on the topic than you guys all did. But I just wanted to bring up this disturbing announcement that will come to pass about a little over three weeks from now, on I think it was December - was it the 21st? No, I don't remember. Anyway - oh, the 14th, I think. Anyway, that's happening.

We've got Intel's response to the horrifying JTAG over USB discovery that we talked about last week. Another surprising AWS, that's Amazon Web Services, public bucket discovery. Android phones caught sending positioning data, even when all permissions to do so are denied. A surprising number of websites, sort of a disturbing number, found to be watching their visitors' actions, even without submitting any information to the website. More fallout from the Infineon ID card bad public key synthesis problem. The return of BlueBorne that we discussed in early September. A new arrival to our, "Well, that didn't take long" department. And unfortunately this is regarding the Amazon key service that we just talked about.

Leo: Oh, yeah. Oh, yeah.

Steve: That didn't take long. We've got speedy news for Firefox 57; a bit of miscellany; some listener feedback. And then, because this episode is titled "Quad 9," we're going to talk about a very appealing and very speedy privacy and security-enhancing DNS service known as Quad 9, and you can guess the IP because...

Leo: Quad 9.

Steve: It's in the name, yeah. So I think an interesting podcast for all of our listeners. Not surprisingly, because the industry keeps giving us good things to talk about.

Leo: There's plenty of juice in this lemon, yes, Steve.

Steve: Yeah, we're never running out.

Leo: Yeah. Okay, Steve. I've got a Picture of the Week right in front of me, right now.

Steve: Yes, we do. And I love it. It's someone sitting in front of the Human Resource guy's desk, and he is speaking. And the caption says - this is the HR guy: "We couldn't hire the cybersecurity candidate you sent us. He was saying too many scary things about our computers." It's like, yeah, right.

Leo: That's like - who was it that they're going after, they're threatening to go after the guy who revealed that they were revealing information about their - oh, DJI, the Phantom...

Steve: Yes, the Chinese Phantom manufacturer.

Leo: It's like, yeah, oh, yeah, go after that guy, the guy who told you what you're doing wrong. Good idea.

Steve: Yes. After he said, "Does penetrating your servers qualify under your bug bounty," and they said yeah.

Leo: Oh. They even gave him a go-ahead. Wow, wow.

Steve: Yes, exactly. And then they decided, oh, we're not happy with you anymore. So, yeah. I mean, it is a dicey relationship. Anyway, I love that we wanted to hire the cybersecurity guy, but he really scared us about our computers, so we said no, no.

Leo: Just put your head in the sand, yeah.

Steve: We want someone who just - just bring ice cream, please. We don't want anything scary. That's right.

Leo: No troubling news, no.

Steve: That's right. So anyway, yesterday was the 32nd birthday of Windows. The v1.0 of Microsoft Windows was - it went RTM, release to manufacturing, as we used to call it back then. Now it's kind of rolling disclosure. Back once upon a time when they actually said, okay, we're done, that was RTM. And that was 1985.

Leo: Wow.

Steve: So 32 years ago and a day, 32 years yesterday, Windows turned 32. Then, okay, so that was November 20th, '85. December 9th, '87, so a little more than two years later, two years and kind of half a month, was Windows 2.0. So that's December 9th, '87. Then May 22nd, 1990, that was the big one. That was Windows 3 where they finally, like, okay, they kind of - they were beginning to get this right. Remember? And everyone remembers Windows 3.11. That was like, everybody was happy. We had memory management. We were able to access more than 640MB. You could do more than one thing at a time. It was like, okay, this is kind of usable. I mean, our screens were still freezing, and it was crashing constantly. You had to always be saving all your work, otherwise it would all be lost. But those were the fun pioneering days.

Leo: Yeah, yeah.

Steve: Yeah.

Leo: Very much remember that. We've come a long way.

Steve: We really have, yes. It kind of works now, pretty much.

Leo: Amazingly.

Steve: After 32 years, yeah. So I did want to mention just briefly - you guys covered it extensively, and I'm sure you'll be talking about it with Jeff and Stacey tomorrow on This Week in Google. Today, this morning, FCC's Chairman Ajit Pai, who has made the reversal of the previous administration's Net Neutrality protections one of his top priorities, unveiled the FCC's plan to give Internet providers broad powers to determine what websites and online services their customers can see and use, and at what cost. So in other words, the end of federal government-enforced Net Neutrality.

So that's the announcement. The final decision, which will be put to a vote next month on Thursday, December 14th, just over three weeks from now, is expected to pass if the votes fall along political party lines, as they're expected to, since Ajit Pai's Republican Party holds three of the Commission's five seats, with two Democrats in the other two. So if the vote falls 3-2, that passes, and then the FCC has done what they said they were going to do in the run-up.

And again, I don't want to beat this to a pulp, but this essentially allows the ISPs to not be treated like a common carrier with a hands-off policy, but, if they choose to, to throttle traffic to enforce different sorts of plans, depending upon what it is you're downloading. In other words, arguably at the moment we're paying our ISP for our connection to the Internet; and the presumption is, and what has always been the case, you have access to the entire world equally. And that is a luxury that may be changing, sadly.

Leo: Yeah. Terrible.

Steve: And, you know, it's a problem, too, because, I mean, if you think, okay, well, what if in another two and a half years, or maybe three and a half years, oh, I guess about three years, we switch back to a Democratic executive. And if there are lots of seats moved back in that direction in the House and the Senate, and all this changes again. I mean, the problem is we're seeing a lot of this changing of policy based on who's in Washington, which is very expensive for businesses who need to have some sort of planning horizon that allow them to do what they want to do, regardless of which side they fall down on this. So I don't know. And then you guys were talking, and I've not followed what happens in localities. That is, for example, California is tending to be rather activist on the Democratic side. Well, can California enforce some change?

Leo: No. That was among the other things they're proposing with this is that the states - this supersedes all state and local laws. Which is appalling, frankly. These are the big "states rights" people. Oh, but only when we care.

Steve: I was just going to say that, unfortunately, the Republicans have been all into this federation and not having all the power concentrated in Washington.

Leo: Yeah. No, that's not true.

Steve: Except where they want it to be.

Leo: Only where, yes, only it comes with a few social issues.

Steve: Okay. Well, so that we don't descend into politics, let's talk about - and this is - the good news is there's some takeaways for our listeners, and I love it when listeners can do something. This is Intel's response to what we discussed last week, which is this horrific JTAG, I mean, it's just hard to even say the phrase "JTAG over USB."

We explained last week that what JTAG is, it's this universally supported, universally recognized, essentially serial debugging interface which, for embedded systems and Intel's IME, their management engine, is an embedded system by definition. What is going to be disclosed in more detail next month by the guys at Positive Technologies, the research firm who found this, is in more detail what this is.

And what's really interesting to me is was it a bug that allowed JTAG over USB? I don't understand how that can be. But maybe it was somehow the USB interface is abused to get access to JTAG. But it sounds more like a feature. Like, oh, yeah, we intended to export the JTAG serial debugging interface over the serial USB interface to make it easier, so you didn't have to open the box in order to access pins on the chip. It's like, okay.

But anyway, what has been revealed - and to their credit Intel scrambled in order to respond, although we also got a bunch of corporate speak. So Intel in their announcement of this, or sort of acknowledging announcement, says: "In response to issues identified by external researchers, Intel" - they, Intel - "has performed an in-depth comprehensive security review of our Intel Management Engine, Intel Server Platform Services, and Intel Trusted Execution Engine, with the objective of enhancing firmware resilience." And of course I'm thinking, wouldn't it have been nice if they'd done that in the beginning, rather than, like, now.

"As a result," they write, "Intel has identified" - Intel, right - "Intel has identified security vulnerabilities that could potentially place impacted platforms at risk." Uh-huh. So it turns out in their disclosure we're talking nine years of chipsets. That is, going back nine years the earlier chips are affected. So this is the ME firmware versions from 11.0, .5, .6, .7, .10, .20; SPS firmware 4.0; and TXE - that's the Trusted Execution Engine - 3.0 are impacted.

And then they said: "Based on the items identified through the comprehensive security review..." - which the disclosure, the public disclosure of these forced them to apparently do - they write: "...an attacker could gain unauthorized access to platform, Intel ME feature, and third-party secrets protected by the Intel Management Engine, Server Platform Service, and Trusted Execution Engine. This includes scenarios where a successful attacker could impersonate the ME/SPS/TXE, thereby impacting local security feature attestation validity; load and execute arbitrary code outside the visibility of the user and operating system" - in other words, as we know, down in the Ring -3 zone where nothing exists except this environment, which is then able, but which despite the fact that it's so low has access to everything going on above - "cause a system crash or system instability." And then they say: "For more information, please see this Intel Support article."

Now, they have in response to this fixed the problems. Unfortunately, you don't get Intel firmware from Intel. Your motherboard manufacturer, or laptop or desktop, whatever, manufacturer, they have the relationship with Intel. We as the customers have the relationship with them. So at this point only Lenovo is listed on Intel's documents and has very quickly provided comprehensive coverage. And Leo, there's a link in the show notes, and to our listeners, there's a link in the show notes that is very extensive, support.lenovo.com, which a firmware update listing by model number for laptops, for desktops, I mean, for everything. And for example, my Carbon X1 Lenovo is there.

But the other cool thing, and this is, as I was mentioning, there's a takeaway for our listeners, is that Intel has published a detection tool. Yay. The link is in the show notes. I can say it because it's not too gnarly, so it's downloadcenter.intel.com/download/27150. So once again, https://downloadcenter.intel.com/download/27150. It is available both for Windows and Linux. And there's a command line version and a graphical user interface version. Grab the one you want.

And this is very nice. This will allow you to determine whether the system you're running it on has these problems and therefore needs its firmware updated. And so understand these are serious motherboard firmware, both remote over the network and local, so this is fixing many things, only one of which - I think it's 5705 was the number. That's the CSV assigned. Only one is the USB JTAG bug. I can't wait to find out next month like what's the nature of this. Did they discover JTAG on USB? Did they abuse something in order to get it? Did Intel mean to publish JTAG out of the USB ports? I just - this is really interesting to me because it's so crazy to do that. Or is it fine, do they think, as long as there are no bugs? I just don't know.

So there are still a bunch of questions. But in the meantime, certainly for enterprises who have a large number of identical systems, as we were talking about last week, there is now a detection tool. So thank you, Intel, for that. And unfortunately it will be then up to the producers of the retail versions of these systems, like Lenovo that's using Intel chipsets, like Dell, well, like everybody essentially, to arrange to provide the similar patches as Lenovo has that people can run on their machines in order to fix them.

Leo: It's interesting because I got an update to the firmware, the ME Firmware 11.8 on November 1st. But I just got an update to that, 11.8 with a more recent number, 11.8.50.3425. I'm guessing that this is the fix. But I'm going to run the install. I'm showing it right now on the screen. And then I'll run the Intel. The Intel program will tell us whether we've been patched or just whether...

Steve: Correct.

Leo: Yeah, okay.

Steve: Yes. It will tell you whether you are running a version of firmware which they know to be vulnerable.

Leo: Okay. Okay.

Steve: And so it'll say yes, you're current, and you're not susceptible. So is that on a Lenovo?

Leo: That's the X1 Yoga, second generation.

Steve: Nice, yes.

Leo: Which was marked as vulnerable.

Steve: Yup. Good. Let us know.

Leo: That's one of the reasons I buy ThinkPads. I know you love your X1 Carbon.

Steve: I do.

Leo: I really love these laptops.

Steve: Yup. And I agree with you, too, that the keyboard really feels nice.

Leo: No, it's a good laptop, yeah.

Steve: They've not sacrificed it the way Apple unfortunately has in order to make it incredibly thin. It's like, okay, look.

Leo: I like the X1 Carbon, but I wanted the two-in-one screen, and I wanted the OLED screen. I really splurged on a new OLED screen. But, boy, it sure looks good. I use it for photography so it's okay.

Steve: So we have another case of an exposed Amazon Web Services bucket. Well, okay, three buckets, actually. What?

Leo: Well, just, you know, just another case, yeah.

Steve: I know. I know.

Leo: The last one was a government bucket, wasn't it?

Steve: Well, this one is. So it's Chris Vickery of UpGuard. He's the guy who has recently been discovering these open and publicly accessible AWS shares. He's found three buckets belonging to the, get this, the U.S. Department of Defense, the DOD.

Leo: Yeah, nice, mm-hmm.

Steve: Yeah. Both CENTCOM and PACOM, which are intelligence-gathering operations. The three S3 buckets were configured to allow anyone with an Amazon Web Services account, which of course anyone can get, to access them, and were labeled "centcom-backup" - okay, you don't have to guess much what that is. Lord - "centcom-archive," and "pacom-archive." In response to this, CENTCOM responded in their typical bureau speak, saying, "We determined that the data was accessed via unauthorized means by employing methods to circumvent security protocols..."

Leo: Yeah, the browser.

Steve: Exactly, yeah, "...said Major Josh Jacques, a spokesperson for U.S. Central Command. Once alerted to the unauthorized access, CENTCOM implemented additional security measures to prevent unauthorized access." Translation: We decided not to leave the password field blank. So, my lord, yeah. So UpGuard posted. They said: "The data exposed in one of the three buckets is estimated" - and get this - "to contain at least 1.8 billion posts of scraped Internet content over the past eight years, including content captured from news sites, comment sections, web forums, and social media sites such as Facebook, featuring multiple languages and originating from countries around the world. Among those are many apparently benign public Internet and social media posts by Americans, collected in an apparent Pentagon intelligence-gathering operation, which raises questions of privacy and civil liberties."

Leo: I think it just found their reddit porn folder and...

Steve: "While a cursory," UpGuard writes, "while a cursory examination of the data reveals loose correlations of some of the scraped data to regional U.S. security concerns, such as with posts concerning Iraqi and Pakistani politics, the apparently benign nature of the vast number of captured global posts, as well as the origination of many of them from within the U.S., raises serious concerns about the extent and legality of known Pentagon surveillance against U.S. citizens. In addition, it remains unclear why and for what reasons the data was accumulated, presenting the overwhelming likelihood that the majority of posts captured originate from law-abiding citizens across the world."

So, whoops. Caught with their pants down because they decided not to put a password on their 1.8 billion archive named "centcom-archive," thus not even a disguised name. Wow.

Leo: Obviously they collected it, and they just left it there, and they forgot about it.

Steve: Yeah. Yeah, although you have to wonder, like maybe it got collected, and then they transferred it to a bucket to make it available to other people, to sort of - because you would hope that they're not collecting it in an Amazon cloud bucket. It's like, okay, what? That's where I keep this podcast backed up. But I don't - you don't want the USDOD to be archiving its Internet scrapings there.

Leo: But do we care that they're scraping in Internet posts?

Steve: Nah, not really.

Leo: I guess they're public posts.

Steve: It's public, yeah, exactly. Okay. So Android was discovered to be, all of this year, starting in early January, to be sending tracking information back to Google without the user's permission. This was from some research that Quartz did. Throughout all of 2017, Android devices, with all location tracking permissions disabled and even without a carrier SIM card installed - so you'd kind of think, okay, I'm off the grid - they've been sending, it turns out, the IDs of all cell towers within range back to Google for some kind of analysis. Which Google acknowledged, after the guys at Quartz observed the behavior and said, "What's going on with this, Google?" when all location tracking has been disabled.

So in their reporting, Quartz wrote: "The cell tower addresses have been included in information sent to the system Google uses to manage push notifications and messages on Android phones for the past 11 months, which was acknowledged by a Google spokesperson." Google said they were never used or stored, and the company is now taking steps to end the practice after being contacted by Quartz. By the end of November, this November, this month, so like a week and a half, the company said, Google said: "Android phones will no longer send cell tower location data to Google, at least as part of this particular service, which consumers cannot disable."

Google explained to Quartz by email: "In January of this year we [Google] began looking into using Cell ID codes as an additional signal to further improve the speed and performance of message delivery. However, we never incorporated Cell ID into our network sync system, so that data was immediately discarded, and we updated it to no longer request Cell ID."

So I guess my feeling is, okay, not a big deal. There was some panting and huffing and puffing on the Internet about this. It's like, eh. I think there's a worthwhile debate about what truly is anyone's reasonable expectation of privacy when you're walking around with one of these deeply Internet-connected computers keeping your pocket warm. My own feeling is that anyone carrying a cell phone should have a very low expectation of privacy. By its very nature, it's a heavily Internet-connected computer that its user, its own user, did not design and build. So it could be doing anything. And we keep finding that indeed it is doing anything.

So if you don't wish to be tracked and monitored, leave the phone at home. It's just not realistic, I think, to imagine that it's possible to have it both ways. You've got a machine that you want all kinds of incredible features from. And take the battery out. That would probably do it, like the way they do in the movies.

Leo: I don't think Google has yet figured out how to keep it going with the battery out.

Steve: No.

Leo: Pretty sure they don't have that.

Steve: And unfortunately, we can't take the batteries out, actually.

Leo: Oh, then never mind. Oh. Now the conspiracy theories are really going crazy.

Steve: And the question is, when you shut it completely off, is it really off?

Leo: Uh-huh. Now the new Apple Macintosh apparently has an ARM chip that will continue to run even if the Intel chip is powered down.

Steve: Yeah, I just think, you know, we old fogies actually care about such things, but, yeah.

Leo: Nobody else does, yeah.

Steve: Nobody else does.

Leo: Just the people who listen to this show.

Steve: And speaking of old fogies caring about stuff...

Leo: The name of this - that's the subtext, subtitle of this show.

Steve: That's a great, yes, that is old fogies...

Leo: Two old fogies caring about stuff.

Steve: ...still caring, still giving a you know what, yes. We've talked about this before, but there's some interesting new research out of Princeton. Websites are logging keystrokes and mouse movements in real-time, whether or not their users ever click Submit. This has been named by Princeton, the researchers there, as "Session Replay Scripts." So researchers at Princeton's Center for Information Technology Policy have a cool site. We've talked about it from time to time through the years. Freedom to Tinker is their - it's freedom-to-tinker, with the words separated by hyphens, so freedom-to-tinker.com, where they blog about stuff that they've found. They call this their "no boundaries exfiltration of personal data by session replay scripts."

Last week they posted the result of their analysis of websites that are using JavaScript to monitor their visitors' actions and activities while people are present on the site, meaning that you just go to a website, and stuff starts happening before you press any buttons. And we've talked about this sort of behavior in the past, how once upon a time in the quaint old days when Windows was young, which is now 32 years ago, before a website, before contemporary websites could run code, powerful code in the web browsers of everyone who visited, what used to be the case is a static page was delivered, and it just sat there. And it was only when we went to another page on that site, clicking a link, or when we explicitly and deliberately filled out a form and then submitted it to the site, that it saw anything further from us.

But today's websites are not going to be content with that. They want to know what is going on at the other end to every degree possible. So we've talked about how even the position of the user's mouse can be monitored. We talked about it in the context of the latest version of Google's CAPTCHA, the I'm Not a Robot, where all you have to do is just click on the box, making that declaration. Yeah, I'm not a robot. And the point is that they use a number of signals, one being watching your mouse as it moves, as it kind of zeroes in on the checkbox, deciding if you're a computer calculating an arc, or if you seem human based on the way you moved the mouse.

The point is so there's sort of a, you could argue, kind of a benign use case for constant mouse position monitoring. And in fact over on the keystroke side I'm using a web page's ability to dynamically watch your keystrokes on GRC's Password Haystacks page because, if you think about it, when you go to the Password Haystacks, as you're entering test, as you're playing with it, putting in passwords, with every single keystroke I'm recomputing and showing you the size of the alphabet, the how long it would take to brute force the password. That's all being done with client-side script, which is being awakened every time the user hits a key to capture now the password in that field and then recompute everything.

So again, a nice, benign, friendly, user-facing benefit to being able to do this. But we now also have a technology known as WebSockets, which allows code running on a web page to silently initiate a connection back to the mothership, or in this case to the third-party provider of the JavaScript which is running on the page, and to dynamically send to stream in real-time everything that the page's user is doing without them having any awareness of it, having given permission or anything. I mean, the presumption is you go to a site, you're sort of turning yourself over to the site more and more as we move through the years with ads and videos playing. By the way, Ars Technica has started to do this a lot, and it's really annoying me.

Leo: I pay for a subscription, and I think that that helps.

Steve: Okay. Because, I mean, I'm looking at 47MB if I scroll into - it's like, holy crap, really? To show me a video that I don't want? Just really...

Leo: Really, that's way too prevalent now.

Steve: Yeah. So, okay. So the researchers at Princeton have named this practice "session reply" - or I'm sorry, "replay scripts," I have a typo in my show notes - "since the stream can be replayed at the hosting end to recreate the user's actions on the page." For example, did they scroll? So like the scroll position of the page is being streamed back. Where is the mouse? How far down did they maybe read? I mean, you can see how these could be useful analytics.

The researchers wrote: "The stated purpose of this data collection includes gathering insights into how users interact with websites and discovering broken or confusing pages. However, the extent of data collected by these services" - now, that's the other thing. This is not just someone like Ars, to pick on them again. I love Ars, so I don't mean to be picking on them. But, for example, these are typically not their own website coders doing this. Naturally there are third-party services where it just, oh, yeah, just like Coin Hive. Drop this little line of code on your page, and we'll send you all this really cool analytics stuff. So there are third-party services which are collecting data which far exceeds probably the typical user's expectations.

Text typed into forms, which is being collected before the user submits the form so that, for example, if the user decides, whoa, wait a minute, I didn't know I was going to be asked all this, and this is too intrusive, doesn't matter. What they've typed in has already left the browser. It's already been streamed back as they're entering keystrokes. The mouse movements are saved, which seems a little less invasive, but without any visual indication to the user. And some companies allow publishers to explicitly link the recordings to a user's real identity, so they're within this ecosystem. All of this is being deanonymized.

They write: "For this study they analyzed seven of the top session replay companies" - there are seven companies that are the top ones, and I saw an eighth one mentioned in the comments to the story - "based on their relative popularity in their measurements. The services studied are Yandex, FullStory, Hotjar, UserReplay, Smartlook, Clicktale, and SessionCam." So these are services available to websites. They wrote: "We found these services in use on 482 of the Alexa top 50,000." And these are not obscure. In the show notes under webtransparency.cs.princeton.edu is their raw data.

From the top of the list down we have WordPress, Microsoft, Adobe, GoDaddy, Outbrain, Spotify, Skype, and RT. So some very well-known, non-obscure sites are employing these third-party services to provide dynamic user monitoring metrics about how their pages are used. Again, not super invasive, although I wanted to mention it because I know that our users care about these kinds of things. And this is not behavior which you would expect or is obvious when you go to a site. And of course all brought to us now by JavaScript, which is you just can't use the 'Net without it any longer because so many sites are using it to enhance their functionality. So as everyone knows, I gave up on NoScript, and I switched to uBlock Origin to give us some control.

Oh, and speaking of which, after this reporting came out, the EasyList service was updated to include all these domains; and Gorhill's uBlock Origin, which is already pulling from EasyList, will automatically protect anyone using uBlock Origin from these third-party analytics systems. So anyone using EasyList - and I think Adblock Plus also uses EasyList. But uBlock Origin I know does, I checked as I was pulling all this together, already blocks this. So again, another reason to consider maybe using a third-party manager to get some additional privacy enhancement.

So we've many times discussed the fallout, the really amazing fallout from the Infineon crypto library, which was found - which is used in embedded devices and found to be producing factorable private keys, which should be an oxymoron. That's, I mean, the whole point is that you cannot factor a private key, which is the way this cool technology hides the public key within the private key, by multiplying these two primes, and you don't know how to pull them apart again.

So Estonia, which was very forward looking, we talked about a couple weeks ago, had "only" issued 760,000 Infineon-based identity cards. Turns out to be one of the smaller players. Get this: Spain has issued 60 million of these cards, 60 million now completely useless and vulnerable identity smartcards. And remember that here's the problem is that they've issued them, and they weren't very popular. I think I saw 0.02% of them were getting any use. But the problem with something like this is that, once they cannot be trusted, no one can use them. Even though they're in very low common usage, now that any of them can be attacked, because that's what this means, the private key generated by the cards can be factored.

And the price to do that, as we have also since discussed, keeps coming down. Now it's about two grand and a few hours in order to perform this factorization. That means that, if any card can be attacked, then none of the security assertions made by any card can be trusted. Therefore all 60 million are useless for their intended purpose, which in any given instance is that this card is producing an assertion that cannot be spoofed, cannot be later either created or broken. And we now know that's no longer the case. So, yes, 60 million for Spain.

And, boy, I don't know what kind of pain Infineon is in, what the contracts look like and so forth. But this has been a big disaster. And if nothing else we need better oversight. The problem that we have is that, for example, Spain purchased 60 million cards without requiring a third-party security review of the technology. Infineon would have balked, but so Spain could say, you want the contract or not? You're asking us to pay you a chunk of money. We're happy to do that. We want your product. But we're not taking your word for it. We need you to put some academic researcher under NDA and examine your technology and verify independently that we should trust you. That's the kind of thing that isn't happening yet because companies are still allowed to say, "Trust us, we're Infineon." And now this is what happens.

About 10 weeks ago, 10 episodes ago, in early September, we covered the BlueBorne attacks. Our listeners with a good memory will remember that there were a series of eight newly revealed zero-day vulnerabilities which affected virtually all Bluetooth-equipped devices, regardless of OS, both Windows and Linux and IoT devices, embedded OSes. Android was affected. It was big. Smartphones, laptops, TVs, IoT devices, you name it.

And these attacks which were enabled by these eight new zero-day vulnerabilities were extra potent, more so than usual, because no user interaction was required. No Bluetooth association was required. All that was required is that a malicious device leveraging the BlueBorne attacks would get within radio range, which is typically 10 meters, about 30 feet, of a device with a Bluetooth radio on. And that would allow, without any involvement from the target device's user, remote compromise.

So there was responsible disclosure. Lots of people scrambled around and updated their Bluetooth stacks in order to fix these eight newly discovered problems. Well, it turns out that another very popular pair of devices were also vulnerable. And that's Amazon's Echo and Google's Home, apparently numbering about 20 million, 15 for Amazon, 5 million for Google. The IoT security firm Armis, who was the original discoverer of the BlueBorne attacks, has disclosed that an additional estimated 20 million Amazon Echo and Google Home devices are also vulnerable to the attacks. They responsibly disclosed their discovery that the Echo and the Home were also vulnerable. Amazon and Google have both implemented and issued patches so that this news is only coming out after the fact.

And the good news is these devices are updating themselves. So it's unlike having an old Android phone where you're out in the dark at this point. There were two different CVE vulnerabilities affecting the Amazon, and one affecting the Google device. And there was also a denial-of-service attack. Basically it could just crash your Google Home, not that that would do the attacker much good. But in these devices the Bluetooth radios cannot be disabled. That was our advice 10 weeks ago. Until you were sure that you had an update, just turn off Bluetooth if you were likely to be targeted.

I mean, it wasn't - because it required physical proximity, some bad guy had to be within 30 feet of you, which seemed unlikely, or much less than an Internet-based attack that could be executed from anywhere else on the globe. They did responsibly notify, and the patches are out. So the problem has been dodged. But this does further reinforce one of the lessons we keep coming back to here, which is the absolute necessity of all interconnected devices, whether they be thousand-dollar smartphones or $7 light bulbs, to be remotely updatable.

We talked a couple weeks ago that there is a forthcoming, hopefully to be adopted, RFC describing a means for doing this, basically moving old-school big-iron Internet technology, the idea of having multiple firmware images from which you can boot either, to allow a device to be updated and then rebooted from the other image in order to bootstrap itself into improved code. But, boy, if anything is smart enough to have firmware, it needs to be smart enough to have that firmware updated. I think that's probably the perfect way of saying it. If anything is smart enough to contain firmware, that firmware needs to be smart enough to update itself, to arrange to keep itself current because, if it's smart and has firmware, what we're seeing is it can be abused.

And from the, as I said at the top of the show, "Well, that didn't take long" department, boy. We discussed with some skepticism on my part, I will say, a week or two ago, the release of Amazon's great new feature, Amazon Key, where as our listeners will remember a special Amazon camera would be positioned to look at the foyer, the entryway of a person's home, where there would be the back - it would be seeing the inside of the front door and the inside area. You would also have an Internet-enabled smart lock on the door that Amazon is able to control. So the idea would be this solves the problem which, Leo, you were explaining. I don't have the problem, but I guess it is a big problem in some areas, of Amazon wanting to deliver packages, and it just being unsafe from a security standpoint. Packages are disappearing, apparently, from people's front porch.

Leo: It's a big problem around here, yeah.

Steve: Okay.

Leo: And, by the way, another problem, which frankly doesn't encourage me, remember we said it was only in areas where Amazon Logistics did the delivery.

Steve: Ah, right.

Leo: Well, I've learned that Amazon Logistics often comes in unmarked cars by a person - they basically do it like Uber. So now you're letting somebody just who drives up your driveway in an unmarked car, walk in your house.

Steve: Yup. Yup. And in fact I got such a delivery last night. They couldn't - this is a person who didn't know where I was. And I got a phone call from area code 209. I thought, what the heck? You know? And I just ignored it the first time. Then it called again, and I thought, okay, well, that's unusual. So I picked it up, and it was a person saying, "Where's your house?" I said, "Who are you?"

Leo: Why do you want to know, buddy?

Steve: Yeah. Anyway, so I said, "Okay, I'll come out." And so I came out, there was just a big white van and a person who maybe was doing this to earn some extra cash.

Leo: They do it like Uber. They post them, and you take up delivery jobs just like Uber.

Steve: Wow.

Leo: And so that worries me quite a bit. And those are billed as Amazon Logistics. I don't know if that means that's who's going to have access to your house or not.

Steve: Well, so get this. So just to finish the scenario, you've got the camera looking at the inside of your front door. With this system, when you permit this to happen, the delivery person is able to push a button on their little handy keypad thing, or maybe on their own smartphone, probably just uses their own smartphone. The camera starts streaming. Your front door is then unlocked, allowing them under video monitoring to open the door, slide the packages in, wave hi to the camera, and then close the door. Then they press a button saying, okay, delivery completed, and the door relocks. And the idea being then that the video that was captured of this process is then sent to the owner of the home, I think a few seconds afterwards, along with a message saying your packages are safely inside your locked front door.

So the "Well, that didn't take long" is Rhino Labs discovered that a courier equipped with a simple program could use their laptop, and it'll be turned into an app before you know it, to fake a command from the house's WiFi router to cause the Amazon Cloud Cam to be disconnected from the WiFi network, which causes the camera to stop functioning and freezes the image at the last frame. So it's, I mean, you couldn't make this up. This is like a movie. It's literally the camera stops refreshing the image, so it shows no motion, no door opening, no one sneaking in the house.

So the point is that at this point the courier could reenter the house, and I guess they've already done this once. I don't know why it says "reenter." But the courier reenters, does whatever they wish with no monitoring, with no surveillance, then exits and reactivates the camera and locks the door as usual. "This reentry," they write, "would be undetectable by the resident and would appear as a normal delivery in Amazon's data."

So apparently this bug that they found, this glitch, essentially allows the first delivery to occur, and apparently all goes well, but then to immediately reopen the door with the frozen camera and sneak around. And so while, yes, there would be some connection between some mischief and the delivery, if they did something like ran upstairs and stole some jewelry out of a jewelry case or something and then got back out of the house quickly...

Leo: You'd never know.

Steve: Exactly. You would never know. If it was a few days later that you discovered, or a month or a week or something, it's like, "Wait a minute. Where are those earrings?" You have no way of associating those events. So Amazon said, their response to this was: "We currently notify customers if the camera is offline for an extended period." They said: "Later this week" - meaning this week - "we will deploy an update to more quickly provide notifications if the camera goes offline during delivery." So that's good.

However, the guys at Rhino Lab explaining, who know this in somewhat more detail, said while this could help Amazon Key - which is the name of the service - customers know when something is amiss, it doesn't prevent the event from happening. Ben Caudill, who's Rhino Labs founder, told Wired Magazine for their reporting of this that the only way to fully close the loophole would be to cache video locally, even when the camera is disconnected from the network. However, the Cloud Cam doesn't currently cache video locally, and doing so would require significantly more local storage than for the pure streaming which is all it does. So to robustly fix this, you would need a hardware upgrade, like maybe if it hasn't been 30 days considering returning the Cloud Cam. Do you remember how expensive they are?

Leo: It was several hundred dollars for the setup because you need the Cloud Cam and the special slick lock.

Steve: Right, right, right.

Leo: You need the whole setup. And I think it was $399 or $299. It was not cheap.

Steve: Ooh, boy.

Leo: To me, that was enough deterrent.

Steve: Yeah, just the cost.

Leo: Is the cost. Get a Ring Video Doorbell, they'll know.

Steve: This is you, Leo, we're talking about.

Leo: Yeah. I'm not letting these guys in.

Steve: No.

Leo: Plain, unmarked car? Come on.

Steve: Oh, yeah. What could possible go wrong?

Leo: Geez.

Steve: So I don't have any personal knowledge of this yet because Firefox 57 doesn't look like it runs on XP. And I'm still sitting in front of XP, although I've got Windows 7, and I've got 10. I mean, I have other Windows machines all over the place. I only have one XP machine. But it's going to be so much down time for me to switch that I just haven't wanted to yet. I haven't needed to. But the news is Firefox 57, which is the so-called "quantum" release, has apparently been clocking in at two times the speed that it used to be.

Leo: Yup. I'm using it now. I love it.

Steve: Oh, nice. The Hacker News guy said it is time to give Firefox another chance.

Leo: That's exactly what I was thinking. I did.

Steve: Wow.

Leo: Yeah, I'm a big Chrome user. And I know you love Firefox.

Steve: I do.

Leo: But this new Quantum is great. It's beautiful.

Steve: Yeah. So the Hacker News wrote: "The Mozilla Foundation today announced the release of its much-awaited Firefox 57, a.k.a. Quantum web browser for Windows, Mac, and Linux, which claims to defeat Google's Chrome." They wrote: "It is fast. Really fast. Firefox 57 is based on an entirely revamped design and overhauled core that includes a brand new next-generation CSS engine written in Mozilla's Rust programming language, called Stylo," that is, the CSS engine is called Stylo.

"Firefox 57 Quantum is the first web browser to utilize the full power of multicore processors and offers 2X times faster browsing experience while consuming 30% less memory than Chrome. Besides fast performance, Firefox Quantum, which Mozilla calls 'by far the biggest update since Firefox 1.0 in 2004' [so 13 years ago], also brings massive performance improvements with tab prioritization and significant visual changes." Now I'm super excited to try this, Leo.

Leo: It looks like Edge, which is weird. But it's a very clean UI.

Steve: Nice, nice. I will be using it tonight because I have Win7 on my Carbon X1. "With a completely redesigned," they write, "UI called Photon. This new version also adds in support for AMD's VP9 hardware video decoding" - as we know, VP9 is kind of moving toward becoming the standard codec for the 'Net - "during playback to reduce power consumption, and thus helping to prevent systems from draining their batteries. Firefox 57 also includes built-in screenshot functionality, improved tracker blocking, and support for WebVR to enable websites to take full advantage of VR headsets.

"Firefox has plans to speed things even further by leveraging modern GPUs in the near future. Firefox Quantum for the desktop version is available for download now on Firefox's official website, and all existing Firefox users should be able to upgrade to the new version automatically. The Android version is rolling out on Google Play shortly, in coming days, and its iOS version should eventually arrive on Apple's official App Store." So yow, and yay. Cool.

Leo: Yeah, I've been using it for a while on Windows, and I really like it, yeah.

Steve: Nice. So a tiny bit of miscellany. My Twitter account is creeping up towards 60,000 followers. I was put in mind of it because I think it was Simon Zerafa who caused me to take a look at ShieldsUP!. We are just shy of the 100 million mark of ShieldsUP! users. We're getting, I checked, about 4,500 new ShieldsUP! uses per day. So we are about 45,000 shy of 100 million. So in about 10 days from now, probably just shy of next week's podcast, we'll be crossing the "100 million served" mark.

And these are not multiply counted. If a user goes there and does several ShieldsUP! uses during one session, I only count them as one use. I maintain the most recent 4,096 IPs from which ShieldsUP! is being used in a most recently used, an MRU list putting anyone who's not found at the bottom and knocking the oldest off the top. So that's really all - I did that from day one, so that's a really solid, separate 100 million use mark. So very cool. And of course ShieldsUP! is what got me into the security business. It was when I discovered that people's Windows hard drives were mapped as c: on the public

Internet. I thought, okay, somebody's got to bring this to the world's attention. So I did.

And I love this tweet. This was a retweet from Simon Zerafa, who was retweeting Stephen Cole, who had just phrased things in a fun way. He said: "The U.S. dollar plummets to a new low of 0.00013 bitcoins." And in fact I checked yesterday when I put my first post of this into my own notes for the show, and then again this morning, $8,200 is where...

Leo: We've got to unlock our wallets.

Steve: I've got to go find mine.

Leo: You're a lot richer than I am. You have 50 of them.

Steve: Wow. I do. Yikes.

Leo: You're a wealthy man.

Steve: Indeed. And so two notes about SpinRite from a Benjamin Rose. He tweeted: "Steve. Could hard drive ECC" - that's the error correction - "be disabled on a hard disk? If so, could that be in the next SpinRite release?" He says: "I'd rather let DynaStat do the work instead of the controller's ECC. SpinRite does a better job." And so to Benjamin, he should know that's always been there. SpinRite always disables error correction as one of the things that it does. It does the recovery in multiple phases. But, for example, you would never want to be doing surface analysis and checking the surface for defects if error correction was in place because it would hide the defect.

And so SpinRite has always been capable of dynamically bypassing error correction when it wants to get to the truth about the disk. And then, however, if it is during the data recovery phase, if it is unable to recover the data to get a perfect read, then it will reenable error correction in order to deliberately bring the error correction back online in order to use its leverage, the error correction leverage, to get just one final last fully corrected read before it then frees the drive up to map that out as a bad and uncorrectable sector. Then it maps a new one in and replaces the recovered data back in place. So all this happens behind the scenes with the user just saying, wow, look at that thing spin. So this is what's going on with SpinRite. So a lot of technology that isn't normally seen.

And then a second tweet, Naruto Uzumaki, I hope I said that right, he said: "While more is fun, on a clearly failing 1GB Sony Microvault flash drive, a Level 4 scan using SpinRite is able to temporarily fix it." And, okay. So I just did want to mention that this tweet put me in mind of just saying to people that there's a limit to how much SpinRite can pull - how far back from the grave SpinRite can pull data. And so if someone just wants to play with SpinRite by watching it fix a drive that is really trying to die, it can keep doing so, but there's a limit. So I do want to encourage people not to confuse SpinRite's ability to in some cases perform miracles of data recovery with that also being a signal that it's time to swap the drive out. Like use SpinRite to make the drive readable, make an image, make a backup.

Leo: Don't keep using it. [Crosstalk].

Steve: But, yeah, just don't - yes, exactly. After your 50th recovery it's like, oh, yeah, it won't keep doing it forever. Sooner or later the drive is going to win this battle. So use SpinRite to pull it from the gray zone; but, once it gets too far into the black, you're really in trouble.

Leo: Yeah, yeah.

Steve: So don't push it too far.

Leo: So what is this, what is it, Quad 9? What is that?

Steve: Okay. So Quad 9 is a new DNS service.

Leo: Like OpenDNS, kind of?

Steve: Exactly. Very, very interesting. It was founded by a consortium of three groups: IBM Security; a group called the Packet Clearing House (PCH), they're providing the global Internet infrastructure; and then there's a group known as the Global Cyber Alliance, which is the third piece of this. So it's called Quad 9 because the DNS is 9.9.9.9.

Leo: Ah. Okay.

Steve: And the 9-dot network is IBM's Class A, their what is it, 24 million IPs that IBM has. So they decided they wanted to support this. They have a security group, the IBM X-Force. So curious about this, I fired up GRC's DNS Benchmark, and I've got a screenshot of it here in the show notes. I added 9.9.9.9 to the list of DNS servers that were already built in. The DNS benchmark uses the ones that my system has, which naturally it's Cox since I'm a Cox cable subscriber. And then a bunch of others that it knows about. And then I added 9.9.9.9 myself. It's possible people don't know this, but if you right-click on the little icon, the little starburst icon in the upper left, there's a big dropdown menu list, and you can do things like add servers. You can create an INI file to make those choices sticky and so forth.

Anyway, ran the Benchmark, and it is using - comparing 9.9.9.9 to all the, like, 50 other DNS servers. Not surprisingly, the ones for Cox were the fastest because they're absolutely the closest to me. They're at the other end of my cable. But the 9.9.9.9 was not noticeably slower. That is, it was at the very top of the list, despite just being a universally usable DNS server.

Now, they're doing this by using anycast routing, which automatically finds and uses the nearest DNS server to the user. So at the time of this launch, which is now, there are 70 points of presence, that is, 70 physical servers located in 40 countries, but that number is growing to 160 through 2018. So this service is launching with 70 and will be more than doubling to 160 servers. So people will be getting them even closer to them.

So what is it? So the focus is a privacy and security enhancing DNS service which is completely free to use. So it's recursive anycast DNS that provides end-users with robust security protections, high performance, and privacy. From their own description, under security they said: "Quad 9 blocks against known malicious domains, preventing," they wrote, "your computers and IoT devices from connecting to malware or phishing sites. Whenever a Quad 9 user clicks on a website link or types an address into a web browser," they write, "Quad 9 will check the site against the IBM X-Force threat intelligence database of over 40 billion analyzed web pages and images. Quad 9 also taps feeds from 18 additional threat intelligence partners to block a large portion of the threats that present risk to end users and businesses alike."

They said under Performance, and I verified this myself just from here: "Quad 9 systems are distributed worldwide in more than 70 locations at launch, with more than 160 locations in total on schedule for 2018. These servers are located primarily at Internet exchange points, meaning that the distance and time required to get answers is lower than almost any other solution." For example, OpenDNS is in the benchmark and didn't show up. I mean, they're way slower than, right now, than Quad 9 is. "These systems are distributed worldwide, not just in high-population areas, meaning users in less well-served areas can see significant improvements in speed on DNS lookups. The systems are anycast, meaning that queries will automatically be routed to the closest operational system."

Under Privacy: "No personally identifiable information is collected by the system. IP addresses of end-users are not stored on disk or distributed outside of the equipment answering the query in the local data center. Quad 9 is a nonprofit organization dedicated only to the operation of DNS services. There are no other secondary revenue streams for personally identifiable data; and the core charter of the organization is to provide secure, fast, private DNS."

And of course it's easy to use, and I'll be switching to it when we hang up the podcast. I'm switching completely to it. Administrators can easily configure endpoint devices to point to the Quad 9 DNS server at address 9.9.9.9. Oh, and I did some more digging about the technology because they mention also having a whitelist. They have a blacklist. They also maintain a whitelist of the top one million sites on the Internet to prevent inadvertent spoofing or blacklisting. So the known safe list of major sites like all of our good ones that are known, overrides the known bad list.

There's also one additional trick. So this 9.9.9.9, or 9.9.9.9, includes the block list. They also support DNSSEC. So their servers will use DNSSEC on origin DNS, which offers signed records, to prevent any spoofing between the authoritative server and the Quad 9 recursive DNS and additional security features. They don't yet support DNSCurve. I think I saw something about DNS over TLS, but I haven't had a chance to track it down. And I also saw a posting where someone from Quad 9 was saying they were in the process of supporting, I think it was DNSCrypt is on the way. So they will be shortly offering a means for us to encrypt our connections to them so that even our own ISP is not able to snoop on what we're doing.

But for me, especially for IoT devices, the idea would be we would configure our household router not to pick up DNS from DHCP from our ISP, but to override it to 9.9.9.9. In which case, all of the systems in our home, all of our PCs, our smartphones and IoT devices, would be using Quad 9 to resolve the IP address for the domain. And if there's anything phishy, okay, P-H-I-S-H-Y, or flaky, we'll just get back a "not found." There's no redirect page. They send back the NX domain error code which is "This domain does not exist."

So, and the other advantage is it's not just for the site you're visiting, as we know. But if your web browser, if the site you're visiting is trusted but has been compromised and is causing your browser to attempt to pull content from a sketchy site, you just won't be able to get the IP through Quad 9.

Oh, what I was also going to mention is that 9.9.9.10 is a variant, just sort of for testing or for people who want to experiment. No block list, no DNSSEC, and some other security features are disabled. I'm not clear on why they're doing it, maybe just in case somebody wants the speed but not the security. Because, I mean, again, these guys are fast.

You can also do a traceroute from your location, if you're curious. I did because I was. It hopped out of Cox, and then a couple other hops, and then got to Quad 9's server. So like about eight hops, and I had a 10-millisecond roundtrip, which is very respectable. And again, using my own benchmark, none of the well-known alternative DNS services came close. I should have looked to see where OpenDNS was. It's easy enough to run a benchmark. Just add 9.9.9.9 to GRC's DNS Benchmark and run it and see where it rates. So I'm very impressed.

Oh, and in their FAQ they ask themselves the question: Does Quad 9 share the DNS data that is generated with marketers? They said: "Quad 9 does not and never will share any of its data with marketers, nor will it use this data for demographic analysis. Our purpose is fighting cybercrime on the Internet and to enable individuals and entities to be more secure. We do this by increasing visibility into the threat landscape by providing generic telemetry to our security industry partners who contribute data for threat blocking."

In other words, and I know this from having read more deeply elsewhere, in return for these other 18 additional threat intelligence partners providing that data, they receive anonymized information if a Quad 9 user trips on one of the sources of information that the intelligence partner provides. So the threat intelligence partners get back essentially a ping when someone did attempt to access one of the records that they provided. So it's useful, but it's completely anonymized at the local server where the query is answered, with completely benign data going back to the partners.

So I just don't see a downside here. These guys seem to have the right motivation. They're providing a useful service. They are basically preventing our computers from getting malware because no known malware will be able to have a sketchy domain resolved. It'll quickly be, essentially, unless we were the absolute very first to encounter it and get infected before anybody else, before it came to the attention of any of these 19 organizations that would immediately blacklist it, we're safe. So, which is not anything that any of the existing generic ISP DNS is doing for us. And there's, for me at least, zero performance penalty. So I'll give it a shot.

Leo: Nice, nice.

Steve: And I'll let anybody know if I change my mind. Yes, 9.9.9.9.

Leo: Very, very nice. And I like the group of people behind it. If IBM says that, I'm going to trust them.

Steve: Yup. And they are.

Leo: And the others are going to keep them honest; right? That's the other good thing. It's not any one person, yeah.

Steve: Yeah. And again, it's - and I've got links in the show notes for anybody who wants to dig deeper or just put in Q-U-A-D and then numeral 9 because this has been making a splash, and I wanted to give it some good coverage and the advantage of taking a deeper look and also verifying that it is absolutely as fast as your own ISP's local DNS. I'll look forward over the next week, I'm sure I'll hear some tweets from people who have run GRC's DNS Benchmark against it, compared to their own DNS servers. And so I'll know more next week. But it sure looks good.

Leo: A breaking story that we will undoubtedly be covering next week on a breach at Uber, fairly significant breach at Uber, that Uber paid hackers $100,000 to cover up, apparently.

Steve: [Gasp] Boy, these guys just do not know how to behave, do they. My lord.

Leo: It's a good way to put it. It's just coming in, so I hope I'm not misstating that. But that's what the headline reads. And it is a fairly large amount of data. Uber has fired their Chief Security Officer as a result. It says, they just announced this today, that October 2016 hackers accessed seven million driver's licenses and names, emails, phone numbers of 50 million riders. And then paid hackers $100,000 - I'm sorry. Did I say 50? 57 million.

Steve: Wait, wait. They have the licenses of the riders?

Leo: Probably not of riders. I would guess that's from drivers.

Steve: Okay, yeah.

Leo: But they do have some information from riders, and they paid hackers, even worse, $100,000 to keep it hushed up. Chief Security Officer Joe Sullivan and another executive have been fired summarily as a result. Personal information of seven million drivers, that would include driver's license numbers, about 600,000 U.S. driver's license numbers, and of 50 million Uber riders. But, good news, no Social Security numbers, credit card information, tip location details, or other data were taken. But names, email addresses, and phone numbers. So if you've used Uber, somebody's got your phone number. Big deal; right?

Steve: But these clowns are Uber and not in a good way.

Leo: Yeah. Yeah, Goobers. We're going to call it Goober from now on.

Steve: Goober, yes.

Leo: At the time they were negotiating with U.S. regulators investigating other privacy violations, which is probably the reason they tried to cover it up at the time. They now say they had a legal obligation to report to regulators and drivers. Instead, they paid hackers to delete data and keep the breach quiet. Uber said it believes the information has not been used, but will not disclose the identities of the attackers. I'm not sure if they know it.

"None of this should have happened," said the new CEO, Dara Khosrowshahi, "and I will not make excuses for it. We are changing the way we do business." This is before, of course, he became the new CEO, back under Travis Kalanick. After the disclosure today, New York Attorney General Eric Schneiderman launched an investigation. I'm sure there'll be much more to talk about next week.

Steve: Wow. Good early info, my friend.

Leo: Yeah. I thought you'd want to hear that. Of course Steve, one of the things we love about Steve, he likes to do his homework. He doesn't talk off the cuff about anything. And so anything that breaks during the show or right before the show you'll have to just listen to his thoughts next week. And of course that means there'll be a lot more information next week about this, I am sure. Just another breach. Just another day, another breach.

Steve: Ho hum.

Leo: Ho hum. You'll find this show and all of Steve's good works at his website, GRC.com. That's where you'll find of course SpinRite, the world's best hard drive maintenance and recovery utility, but also stuff like SQRL, his current project; SpinRite, Shoot the Messenger, DCOMbobulator, ShieldsUP! - now, what, 100 million strong users.

Steve: Yay.

Leo: That's amazing. I've used it many times, of course. Really good way to test your router when you first install it, make sure everything's copacetic. And the podcast, audio versions of the podcast as well as transcripts, which are very valuable for searching, as well as a lot of people like to read while they listen, at GRC.com.

We also have audio and video of the show at our website, TWiT.tv/sn. And you'll find it everywhere you subscribe. Just look for, I mean, this is a show you want to not only listen to every show, but keep previous shows. Lots of people relisten. Sometimes people listen at high speed the first time and then slow it down if they need to. It's good to have a copy. Subscribe on Pocket Casts or Stitcher, iTunes or your podcast app on your phone, whatever you prefer.

We will be back here, as we are most Tuesdays, at about 1:30 Pacific, right after MacBreak Weekly, whenever that ends. It was 2:00 o'clock today. That's 4:30 Eastern time, 21:30 UTC if you want to watch live. If you do watch live, join us in the chatroom. They're great people in there. And it's a lot of fun, the behind-the-scenes conversation. And people who watch this show are smart, so there's always great people in there, lots of good question-and-answers going on during the background. And jokes, too. I usually steal all my material from the chatroom, irc.twit.tv.

Don't forget TWiT.tv/store. The holidays are coming, and now would be a good time for you to pick up that Steve Gibson moustache cup, TWiT.tv/store. This is a new vendor. We were using for a long time a vendor that had some weird countdown thing, like we've got to wait till we get a thousand ordered and all that stuff. Now you can just go and buy directly. $19 for a Security Now! mug. Love that mug.

Steve: Smug in the mug.

Leo: It's the smug mug of Steve Gibson on the coffee mug that says Security Now!. But we've got T-shirts. We also have a Security Now! T-shirt, the baseball tee. Show with pride your affiliation. Plus of course lots of other great TWiT and other shows. Oh, there's a hoodie. You know, when you're going to DEFCON next month, maybe you ought to wear the Security Now! hoodie. Mr. Robot would appreciate that. TWiT.tv/store.

Steve, always a thrill, always a pleasure. A little bit of a short show today. I hope we haven't disappointed our vast listening audience.

Steve: Lots of good news. I think everyone's going to be putting 9.9.9.9 into their browser. Yup.

Leo: I like it. Somebody said, well, they should also get the IPv6 version. Then it would be, what, nine nines?

Steve: They do support IPv6 already, so they're ready to go. Okay, buddy.

Leo: Thanks, Steve. We'll see you next time on Security Now!.

Steve: Bye.